Astronomy and Computing 56 (2026) 101116

Contents lists available at ScienceDirect

Astronomy and Computing

journal homepage: www.elsevier.com/locate/ascom

A dual AI framework for the automatic characterisation of low-interacting H-rich SNe within the ASTRAI (Advanced Supernova Transient Research with Artificial Intelligence) project☆

Andrea Claudio Grasso [a], Vincenzo Del Zoppo [a], Alessio Mezzina [a,c,*], Marco Cataldo [a], Giuseppe Puglisi [a], Fabio Spampinato [a], Stefano Pio Cosentino [b], Maria Letizia Pumo [b], Luca Naso [a]

[a] Koexai Srl, Via Josemaria Escriva 6, 95125, Catania, Italy
[b] Department of Physics and Astronomy "Ettore Majorana", University of Catania, Via Santa Sofia 64, Catania, 95123, Italy
[c] Department of Mathematics and Computer Science, University of Catania, Viale Andrea Doria 6, 95125, Catania, Italy

ARTICLE INFO

MSC:
68T07
62M45

Keywords:
Supernovae
Transient astrophysics
LSST
Light curves
Deep learning
GenAI
Synthetic data

ABSTRACT

Rapidly increasing data volumes from high cadence surveys, and the still larger streams anticipated from the Legacy Survey of Space and Time (LSST), pose a major challenge for the timely and accurate analysis of transient phenomena such as core-collapse Supernovae (CC SNe). Traditional approaches based on semi-analytical or hydrodynamical models, while physically grounded, hinder the large-scale characterisation of these transient events. We present a novel machine learning approach, focusing on low-interacting hydrogen-rich (H-rich) SNe. This approach encompasses two stages: a generative model produces synthetic light curves (LCs) to augment sparse or incomplete datasets, and deep learning models characterise LCs by regressing fundamental physical parameters. The characterisation stage uses an ensemble of neural networks, each specialised to regress a single parameter, improving interpretability, modularity and robustness. To account for observational systematics, we employ data augmentation strategies that mimic realistic noise and cadence conditions. Benchmarking against semi-analytical and hydrodynamical simulations demonstrates that the proposed generative method recovers LCs with error comparable to or smaller than observational noise, while achieving inference speeds orders of magnitude faster than traditional modelling pipelines (with a throughput beyond 10^5 LCs per second on a single GPU hardware). This could enable large-scale characterisation of H-rich SNe within modern and upcoming survey data streams, preparing the scientific community to handle the massive LCs flow expected from LSST.

1. Introduction

Core-collapse Supernovae (CC SNe) are among the most energetic events in the universe and play a crucial role in our understanding of the cosmos (Arnett, 1996). Indeed, in addition to representing essential tools for investigating unresolved questions in stellar evolution (Heger et al., 2003), they are relevant for different branches of modern astrophysical research, ranging from cosmology (Arnett, 1996; Scannapieco et al., 2008; Pumo and Zampieri, 2013) and astrobiology (Smith et al., 2004) to multimessenger astronomy (Janka, 2017; Radice et al., 2019; Ahumada et al., 2021; Cosentino et al., 2025c).

The study of CC SNe relies on the combined analysis of photometric and spectroscopic data. In particular, photometric observations enable us to trace the light curve (LC) evolution (i.e. the variation of brightness over time), which encodes key information about the explosion dynamics and the physical properties of the progenitor star at explosion, such as its mass, radius, and explosion energy.

To interpret these observations, astrophysical models are required to reproduce observed spectra and LCs under different physical conditions. These models, ranging from semi-analytical formulations to full radiation-hydrodynamical simulations, allow us to explore the vast parameter space governing stellar explosions (e.g., Cosentino, 2024).

However, such simulations require substantial computational resources, limiting their applicability when analysing large datasets produced by modern and upcoming transient surveys. Next-generation

☆ This article is part of a Special issue entitled: 'HPC in Cosmology and Astrophysics' published in Astronomy and Computing.

* Corresponding author at: Department of Mathematics and Computer Science, University of Catania, Viale Andrea Doria 6, 95125, Catania, Italy.
E-mail address: alessio.mezzina@phd.unict.it (A. Mezzina).

https://doi.org/10.1016/j.ascom.2026.101116
Received 2 November 2025; Received in revised form 14 March 2026; Accepted 30 March 2026
Available online 15 April 2026
2213-1337/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

astronomical surveys, notably the Legacy Survey of Space and Time (LSST) (Abell et al., 2009) at the Vera C. Rubin Observatory, will collect an unprecedented volume of photometric data. These surveys are expected to observe thousands of transients every night. Specifically, the LSST is expected to increase the number of known SNe by orders of magnitude compared to existing samples (Abell et al., 2009).

Under these conditions, traditional model-driven analyses become computationally prohibitive, and a systematic large-scale characterisation of SNe becomes impractical without new methodological paradigms. To overcome these challenges, recent efforts have focused on accelerated analysis frameworks based on semi-analytical models (e.g., Cosentino et al., 2025a) and, increasingly, on machine learning approaches that emulate the behaviour of complex physical models while drastically reducing computational time (Grassia et al., 2025; Sarin et al., 2025; Grassia et al., 2025b).

Recent advances show that AI can both accelerate SN modelling pipelines and learn representations that capture the physical dependencies encoded in LCs (Lochner et al., 2016; Möller and de Boissière, 2019a). Two complementary approaches are particularly relevant. The first involves generative models, which produce synthetic LCs that reproduce the physical variability of CC SNe, mitigating data sparsity and improving training robustness. The second involves regression-based characterisation, where neural networks directly infer physical parameters—such as progenitor radius, ejecta mass, explosion energy, and circumstellar medium properties—from the observed light-curve morphology. Together, these methods enable efficient inverse modelling while preserving physical interpretability.

In this paper, we propose an AI-based approach to automatically characterise LCs of low-interacting H-rich SNe. We focus on these SNe because, in addition to being one of the most common types of CC SNe in the universe (Perley et al., 2020), they play a crucial role in addressing several key open questions, including the physical conditions that drive the diversity of H-rich SNe (e.g. Smartt, 2009; Pumo and Zampieri, 2011, 2013; Nagao et al., 2025), the relationship between the circumstellar medium and the mass-loss history of massive stars in their final evolutionary stages (e.g. Smith, 2014), and the conditions under which SNe can produce high-energy radiation and neutrinos (e.g. Murase et al., 2011; Fang et al., 2020).

The approach we present comprises two stages: first, a generative model produces high-quality synthetic LCs datasets to augment sparse or incomplete observations; second, a characterisation model rapidly infers physical parameters from observed LCs.

The remainder of this paper is organised as follows: Section 2 reviews related work on both physical modelling and machine learning approaches to SN analysis, Section 3 illustrates the datasets used and the data curation pipeline, Section 5 details the methodology and AI architectures developed, Section 4 presents the data augmentation techniques implemented to ensure model robustness, Section 6 reports and discusses the results obtained. To conclude, Section 8 summarises the findings, limitations, and outlines future perspectives.

2. Related work

SNe characterisation has traditionally relied on a diverse range of theoretical models developed over decades to explain the complex physical phenomena governing these explosive events. These methodologies range from analytical and semi-analytical models to complex hydrodynamic simulations. Each approach is built on specific physical assumptions and is suited to describe particular types of SNe or distinct evolutionary phases (Cosentino, 2024). A foundational category of models explains the luminosity of SNe, particularly Type Ia, as powered by the radioactive decay of isotopes such as Nickel-56 (^{56}Ni) and Cobalt-56 (^{56}Co) produced during the explosion.

Arnett's pioneering works (Arnett, 1980, 1982) provided an analytical solution for the LCs based on the radiative diffusion of energy within the expanding ejecta. For H-rich SNe such as Type II-P, other models are prevalent. The Popov model (Popov, 1993), for instance, analytically describes the plateau phase of the LCs as the result of cooling and recombination of the hydrogen envelope. More complex scenarios are also modelled, such as those powered by a central engine, for example a magnetar or black hole, where energy is transported via radiative diffusion, as explored by Chatzopoulos et al. (2012). Other models focus on the interaction between the SNe ejecta and the surrounding circumstellar medium (CSM), a key mechanism for Type IIn SNe, as detailed in the work by Moriya et al. (2013). Recent advances have also introduced hybrid analytical approaches that aim to provide a more comprehensive description of H-rich SNe. The model presented by Pumo and Cosentino (2025) focuses on the intrinsic evolution of the ejecta, incorporating recombination processes, radioactive heating from ^{56}Ni, and the effect of a non-homologously expanding outer shell, without invoking interaction with an external circumstellar medium. The subsequent work by Cosentino et al. (2025c) extends this framework to explicitly include ejecta–CSM interaction, modelling the dynamical and radiative effects of a low-mass CSM, particularly relevant for potential high-energy neutrino production.

Although these physical models are fundamental to our understanding, they often require significant computational time to explore large ranges of SN explosive parameters, making them impractical for the statistical analysis demanded by modern large-scale sky surveys . This computational bottleneck has driven the adoption of AI, in particular deep learning, as a strategic solution to automate and accelerate the analysis of SNe LCs.

A crucial preliminary step in many SNe LC analysis is to address the nature of observational data, which is often irregularly sampled and contains gaps. To handle this, LCs approximation methods are used to create a continuous representation of the transient's evolution. Although different aforementioned approaches making use of physical models were able to reproduce some subsets of observational data faithfully, recent work has demonstrated the effectiveness of neural network-based approaches. For instance, Demianenko et al. (2023) explored multilayer perceptrons, Bayesian neural networks, and normalising flows as alternatives to physically driven models, finding them to offer a computationally faster approach for LC approximation compared to iterative approximation methods, while achieving comparable quality for downstream tasks such as classification and peak estimation.

Early applications of machine learning in this domain have largely focused on photometric classification. Works by Lochner et al. (2016), Mahabal et al. (2017), Möller and De Boissière (2019b) and Burhanudin et al. (2021) demonstrate the power of various algorithms to classify SNe types from their LCs, a critical task for large-scale surveys.

Similar to very recent works (Grassia et al., 2025; Sarin et al., 2025; Grassia et al., 2025b), we move beyond photometric classification by introducing ASTRAI, a two-stage pipeline that (i) performs physical characterisation by inferring the underlying parameters directly from observed light curves, and (ii) employs a generative AI model to produce realistic synthetic LC datasets.

3. Datasets and data curation pipeline

Our AI framework was built and tested using three distinct but interconnected datasets: two large datasets of synthetic LCs used for training the models and a curated dataset of real SNe observations used for validation and benchmarking. The pipeline is designed to ensure that the synthetic data realistically mirrors the properties of genuine observations. The model processes bolometric light curves as 1D tensors of representing the log luminosity at each epoch after the explosion. Before feeding, standard scaling and linear PCA is applied to the vector of the log luminosities at each epoch to normalise the luminosity range and reduce the number of features both at the input of the classifier and at the output of the generator. Errors are not included as direct inputs; instead, the model accounts for data variability through a custom data corruption explained in detail in the following paragraphs.

- The first dataset consists of 14,300 synthetic LCs generated using a first semi-analytic model better described below that accounts for 4 parameters - Progenitor's Radius ($R\star$), Mass of the ejecta (M_{ej}), Explosion Energy (E_k), and Nickel mass ($M_{^{56}Ni}$) - sampled each day for 420 days.
- The second dataset consists of 11,211 synthetic LCs generated using a newer semi-analytic model better described below that accounts for 7 parameters – the 4 aforementioned plus, Radius, Mass and Slope exponent of the CSM (R_{CSM}, M_{CSM}, S_{CSM}) – sampled uniformly each six hours for a period of 400 days, resulting in 1600 time points.
- The third dataset is a collection of 219 LCs obtained by processing real in-band observations from ZTF, ATLAS and other astronomical surveys. This last dataset has not been used for training and testing the models directly, but has been used to create and refine the parameters of the data corruption pipeline (see Figs. 2 and 3).

3.1. Synthetic training dataset

In this work, we generated synthetic bolometric LCs using semi-analytical models of Pumo and Cosentino (2025) and Cosentino et al. (2025c). In particular we use the submodel "EXP+IE" of Pumo and Cosentino (2025) which is the same adopted in Cosentino et al. (2025c) to evaluate the LC during the ejecta recombination and radioactive decay phases

The detailed description of the two analytical models is explained in the papers (Pumo and Cosentino, 2025; Cosentino et al., 2025c).

We first adopted the 4-parameter model, which was designed specifically to reproduce 87A-like LCs. Then we extended the method to the more complex 7-parameter model to better approximate the first days of the explosion, taking into account the interaction with the Circumstellar Medium (CSM). Fig. 1, Fig. 2 and Fig. 3 illustrate how the principal parameters of the two models shape the resulting LC morphology. The two synthetic datasets have been the main resource used to train and validate our AI models, teaching them to map the features of an LC back to its underlying physical parameters.

3.2. Real observational dataset

A key challenge is that real observations (photometric data across different colour band filters) and our synthetic data (bolometric LCs) are not directly comparable. To solve this, we created a data curation pipeline leveraging the SuperBol algorithm (Nicholl, 2018). We optimised and applied the SuperBol routine to the real photometric data to estimate the bolometric LC for each observed SN. This step enabled to transform in-band data into bolometric estimates accounting for time dilation, extinction, reddening, and blackbody correction allowing us to directly compare the output of our models with real events, validate that our synthetic data is a realistic representation of actual SNe, and benchmark the final performance of our AI framework on real-world test cases.

4. Data augmentation

AI models trained exclusively on "perfect" synthetic data often exhibit limited generalisation ability when applied to real observational data. Astronomical observations are inherently imperfect, characterised by instrumental noise and non-uniform temporal sampling (observation gaps). While flux errors are not provided as a dedicated input channel. The model's robustness to observational noise is addressed through a data augmentation and corruption pipeline. During training and prediction, we apply noise and simulate missing data (gaps) using a custom corruption module. This forces the Conditional VAE to learn a noise-resilient representation of the light curves, effectively handling the intrinsic uncertainties of real bolometric data. The data augmentation techniques are detailed below:

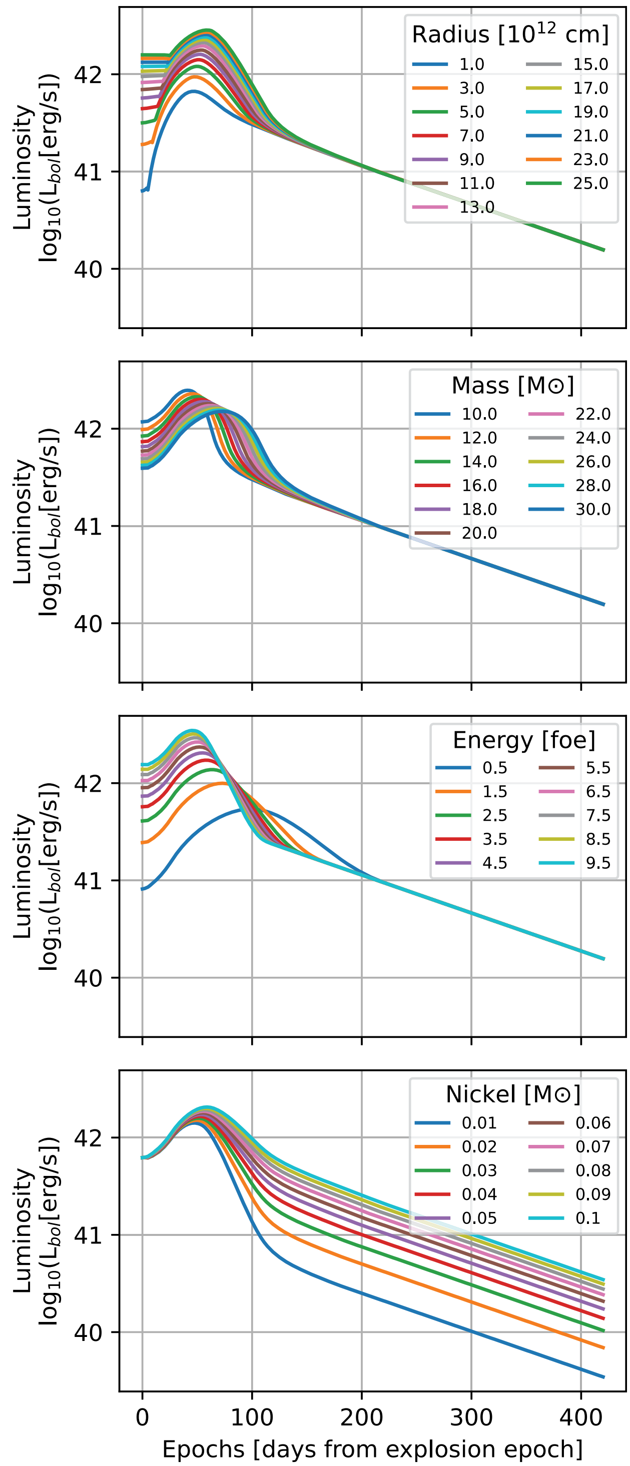

Fig. 1. Four panes illustrate how each parameter shapes light-curve morphology in the four-parameter synthetic dataset. Mass, Energy and Radius contribute to shaping the peak phase of the LC while the Nickel percentage acts on the tail part.

- Noise: The noise is generated independently for every time step in every LC, ensuring variability across both the temporal dimension and the batch. The SNR for a point source with flux F (in electrons) is calculated as:

$$SNR = \frac{F}{\sqrt{F + n_{pix}(\sigma_{sky}^2 + \sigma_{read}^2 + \sigma_{dark}^2)}} \tag{1}$$

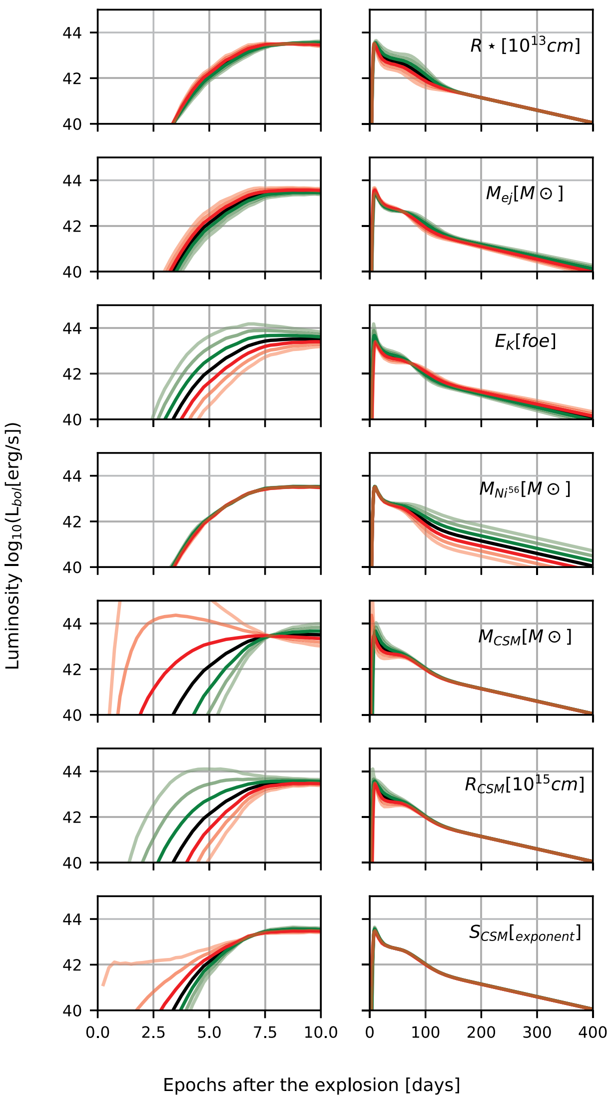

Fig. 2. Mean effect of the change of parameters on the LC over the epochs. The plots show the effect of the parameters at various epochs for the 7-parameter model on a linear scale. The effect is shown in green for a positive variation and in red for a negative variation. The three lines correspond to a change in the parameter 0.5, 1.0 and 1.5 standard deviation from the LC generated with the mean parameters. To enhance the visibility of the plot, a zoom has been added on the first 10 epochs after the explosion.

where:

- F: is the total number of photons from the source.
- n_{pix}: is the number of pixels in the integration aperture.
- σ^2_{sky}: is the variance from the sky background.
- σ^2_{read}: is the variance from the readout electronics.
- σ^2_{dark}: is the variance from the dark current.

However, the full formula can be simplified with a Poissonian component proportional to the source, plus a constant bias component from the background estimation.
Each LC is perturbed by adding random values drawn from a normal distribution with zero mean and a specified standard deviation proportional to the square root of the flux.

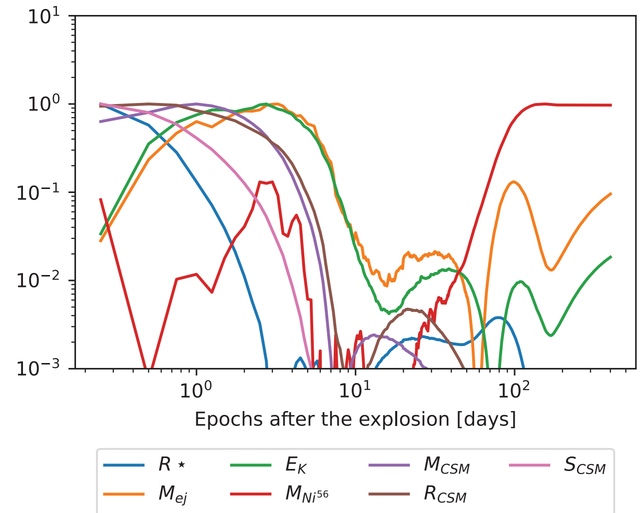

Fig. 3. Mean Variance of the effect of the parameters change on the LC over the Epochs in log–log scale. Graph showing the effect of the parameters at various epochs for the 7-parameter model. The variance is normalised to the maximum epoch. As it is also well known, the parameters of the progenitor can be inferred with good approximation also by only looking at the tail of the supernova. While the first week of data is essential for the inference of the CSM parameters.

This procedure mimics the stochastic nature of measurement errors, scaling their amplitude according to realistic photometric uncertainties.

- Observation Masking: "observation gaps" have been artificially introduced. Blocks of consecutive samples are replaced with NaNs to emulate the gaps between real observations from LSST. This emulates periods when the source is unobservable due to seasonal constraints like earth orientation, sun position, moon position and phase, and atmospheric masking. A plot showing the masking process is shown in the Fig. 4.
- Interpolation: Since MLP models can not be fed with NaN values, missing values introduced in the LCs were filled using linear interpolation. While this approach introduces a well-known bias on the peaks, this has been evaluated as being the most robust alternative, as other options like mean-filling, zero-filling or other non-linear interpolation methods introduced even stronger distortion, being affected by the Runge phenomenon.

Further sources of noise, like reddening or extinction, have not been considered for simplicity.

5. Methodology

5.1. Data splitting

To properly train and test our models and minimise eventual biases, the dataset was carefully divided into training, and test sets. This accurate partitioning was essential to improve the reproducibility of the model evaluation.

5.2. Generative modelling of LCs

A generative model (see Fig. 6) was developed to learn the mapping between SNe's physical parameters and their temporal LCs, enabling the rapid generation of large synthetic datasets.

Architecture: The generative model is based on a Multi-Layer Perceptron (MLP) enriched with residual blocks. Each block includes a

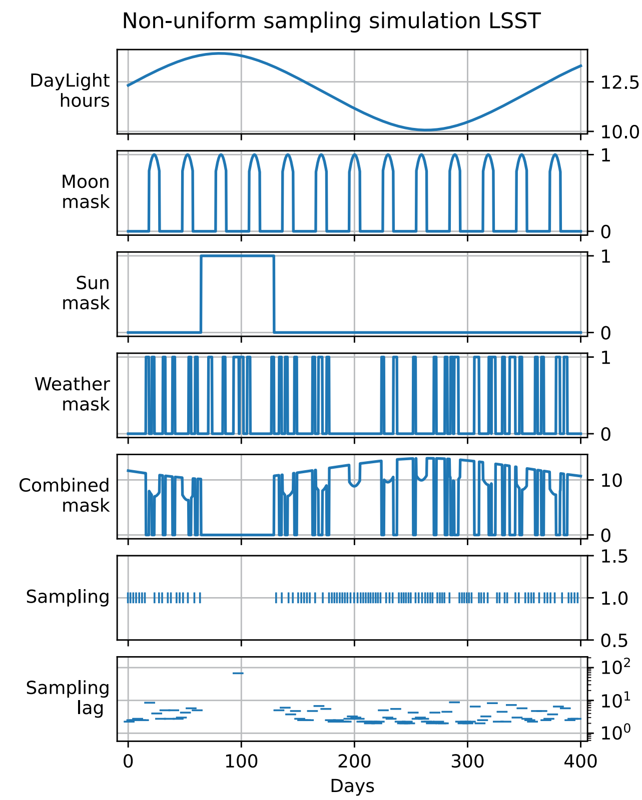

Fig. 4. Detailed explanation of how the masking is affecting the sampling: SN are initialised randomly in the southern hemisphere at the latitudes scanned with WDF cadence, while the sun and the moon are initialised randomly on the ecliptic plane.
Minimum angle between SN and the moon depends on its phase and ranges periodically (between 6.5 degrees for the new moon and 20 degrees for the full moon).
Minimum angle between SN and sun is fixed at 64 degrees, minimum angle between sun and Zenith should be 96 degrees to avoid atmosphere scattering, while maximum angle between the SN and the Zenith is 33 degrees. Given those periodic constraints, another 30% of the time has been removed, accounting for possible cloud coverage, adding randomly placed gaps of 2.3 days.
As a result of the combined masking, most of the lag between two samples is between 3 and 4 days on average, while, depending on the distance from the ecliptic, a big gap up to almost 3 months in the worst scenario may appear due to sun masking.

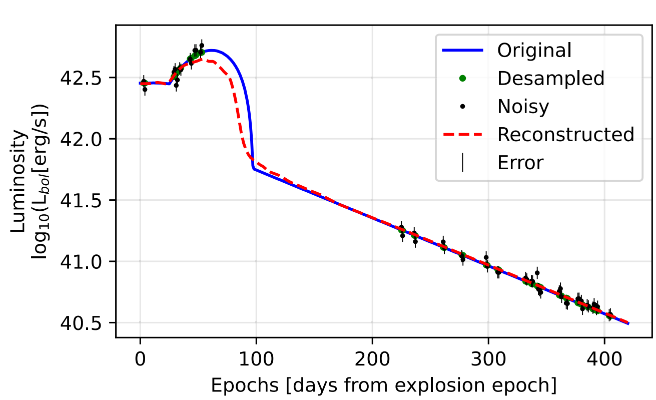

Fig. 5. Figure shows an example of the reconstruction process of a LC. The original synthetic curve gets undersampled and corrupted before being passed to the PPReg. The regressed parameters are passed to the LCGen to get back the reconstructed curve.

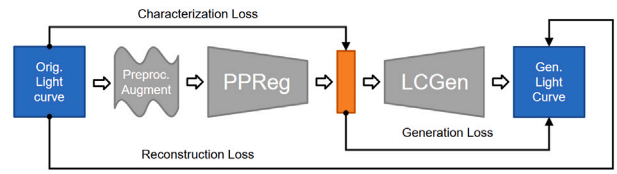

Fig. 6. The schema shows how the two models are trained using separated supervised losses for the characterisation and the generation, and how they are connected together with an additional unsupervised reconstruction loss. While the strong supervision of the separated blocks improves physical explainability thanks to simulated data, the reconstruction loss can be applied to unlabelled real data.

skip connection (He et al., 2016), a fully connected layer with Leaky-ReLU activation (Maas et al., 2013), batch normalisation (Ioffe and Szegedy, 2015) and a dropout (Srivastava et al., 2014) to stabilise and accelerate the training process. While Recurrent Neural Networks are a traditional choice for time-series data like light curves (LCs), the ASTRAI framework adopts an MLPs enriched with residual blocks to better address the specific demands of high-throughput. Compared to RNNs, which are often prone to vanishing gradient issues and higher computational latency due to their sequential nature, our MLP-based architecture offers superior scalability.

We considered recurrent architectures like RNNs, GRUs, and LSTMs; however, in our setting, each element of the generated sequence is directly conditioned on the input embedding, which encodes the relevant physical parameters, including the explosion properties. Although temporal correlations are present in the data, each timestep can, in principle, be predicted independently from an underlying analytical (albeit complex) formulation. Therefore, the sequence does not exhibit strong autoregressive dependencies that would require explicit recurrent modelling.

In this context, relying on recurrent propagation would mainly exploit information encoded in previous outputs, which already reflects the same physical parameters available in the embedding. This approach would introduce unnecessary error accumulation along the sequence and additional computational overhead at inference time, without clear theoretical or empirical advantages over a direct feed-forward formulation.

Moreover, this way, the proposed model can also be trained on inherently noisy observational data. In such conditions, recurrent architectures may further amplify noise through temporal propagation, whereas a direct prediction at each timestep promotes greater robustness and stability. For these reasons, we considered sequence-based solutions suboptimal for the present task and therefore focused on a fixed-length, embedding-conditioned MLP architecture.

To identify the optimal architecture, a hyperparameter tuning process was implemented through a systematic grid search. Several hyperparameters have been explored, including network depth (number of blocks), the width (number of neurons per layer), the number of training epochs, and the dropout rate, to minimise the root mean square error (RMSE) in the reconstruction of LCs. Different randomised runs on both datasets resulted in similar, although different optimal solutions; the final chosen architectures are shown in Fig. 7.

5.3. Characterisation model

The target of the Characterisation model is to solve the inverse problem, i.e. being capable of estimating the fundamental physical parameters of SNe from their sampled LCs. The initial approach involved a unified multi-target model, i.e. a single neural network trained to predict all the physical parameters simultaneously. However, the solution evolved into an ensemble of seven independent MLP models

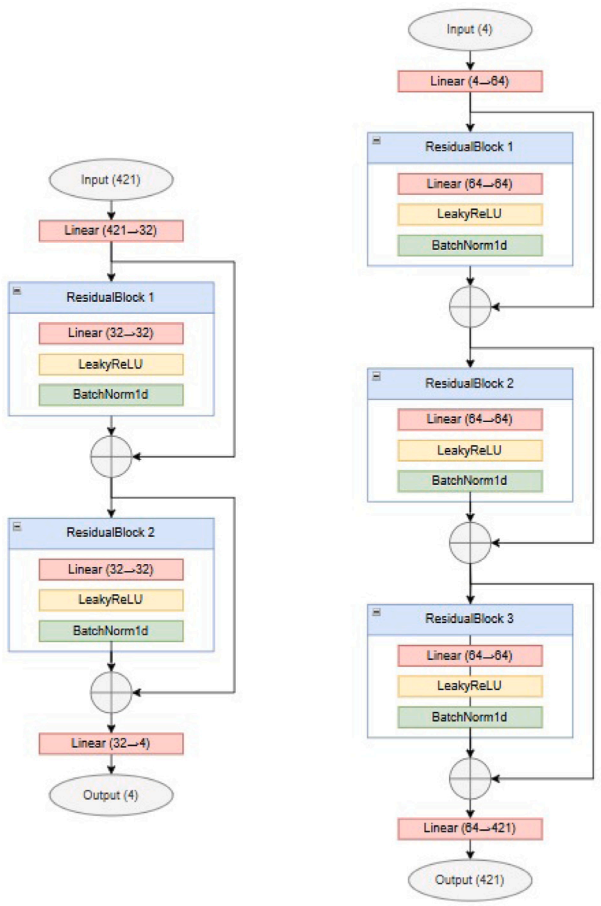

Fig. 7. Detailed explanation of the architectures employed. The blocks were designed modularly to allow for hyperparameter search. The final configuration of depth and width of the PPReg (on the left) and LCGen (on the Right) are shown.

as shown in Fig. 8. In this architecture, each model is specialised in the regression of a single physical parameter.

Using an ensemble of models instead of a single one, offers several advantages: (i) specialisation allows each network to learn the complex relationships between the LC and a specific parameter in a more targeted way; (ii) conflicts between the gradients generated by different loss functions during training are avoided; and finally (iii) the system acquires greater modularity, allowing the addition of new parameters without the need to retrain the entire framework. Furthermore, by systematically incorporating data augmentation to simulate observational noise and irregular sampling, ASTRAI is engineered for robustness and generalisation to real-world data, bridging the gap between detailed physical simulations and the high-throughput needs of next-generation transient astronomy.

For the training, the Mean Squared Error (MSE) loss function was used. Optimisation was performed with the Adam algorithm, while a Reduce_LR_On_Plateau scheduler dynamically managed the learning rate, reducing it when performance on the validation set stopped improving. This approach was intended to promote more stable convergence and help mitigate potential overfitting, rather than relying on a fixed learning rate schedule.

5.4. Data and code availability

All analyses were performed using a custom Python pipeline, publicly available at https://github.com/koexai/ASTRAI.

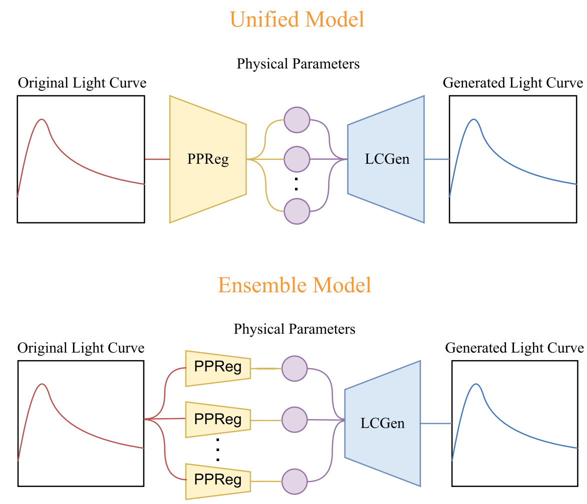

Fig. 8. A zoom-in on the differences between Unified model vs. Ensemble model Physical Parameter Regressor (PPReg) inside the full pipeline.

Table 1
Computational cost comparison of different modelling approaches.

Method	Time	Cores	Runs	Core/h
Hydrodynamical	48 h	8 CPU	180	70 K h
Semi-analytical	0.225 s	8 CPU	1 M	500 h
Our MLP Gen	0.009 ms	1 GPU	1 M	93 s

6. Results and discussion

The optimised generative model demonstrated high fidelity in reconstructing the synthetic LCs (see Fig. 5). Evaluation metrics, such as the Root Mean Squared Error (RMSE) on log luminosity, have been measured for each epoch, showing a similar decreasing trend for all the tested models. Fig. 9 shows a comparison between the generation and the reconstruction error on the clean and augmented test set. It shows that the model can capture the dependencies between physical parameters and morphological features of LCs with a variable accuracy across the epochs, being more accurate on the tail phase, while showing higher uncertainty in the peak phase.

The proposed generative method was tested by generating 1000, 10,000, 100,000, and 1,000,000 LCs, and measuring the generation time. R2Q17.1 A 10-fold cross-validation was performed for all tests, which were conducted on an NVIDIA GeForce GTX 1050 GPU, achieving a throughput of approximately 10^5 light curves per second (see Table 1). In contrast, hydrodynamical simulations can require several days to compute a single model, while semi-analytical models are faster but less accurate, typically requiring a few seconds per simulation (Cosentino et al., 2025b).

Data augmentation techniques transformed our synthetic dataset into a more realistic and diverse learning environment. By simulating observational noise, irregular sampling, and missing data, the model is exposed to the typical imperfections of real measurements. As a result, its ability to generalise to actual imperfect data is significantly improved, making it more robust and reliable for practical applications.

The accuracy of the characterisation models has been compared in the Tables 2 3. The discrepancy between the accuracy of the two models might be explained by the higher sparseness of the explored space, the higher dimensionality, and the randomisation of the parameters of the synthetic dataset, which does not account for meaningful physical correlation between the parameters explored.

The results indicate that the ASTRAI framework is accurate, fast, and robust enough for automated analysis of data from upcoming astronomical surveys. The ability to rapidly infer physical parameters will

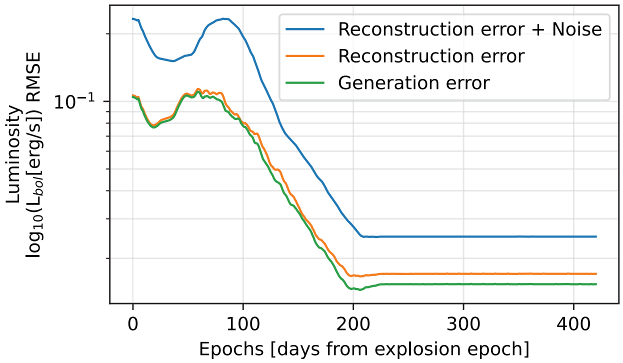

Fig. 9. RMSE on generated light curves. The figure illustrates the impact of the different stages of the algorithm on the reconstruction error. The green line shows the RMSE obtained when using the generation model alone, starting from the true parameters of the clean synthetic data. The orange line includes the propagation of parameter estimation errors by first passing the data through the characterisation model and then reconstructing the light curve, with the RMSE computed between the original and reconstructed curves. The blue line further incorporates the augmentation step, in which the light curve is corrupted with noise, downsampled, and interpolated before being processed by the models. Overall, the generated or reconstructed light curves are more accurate during the tail phase, whereas the peak phase exhibits a higher RMSE.

Table 2
Performance evaluation of the 4 parameter characterisation model.

Parameter	[Range]	(Unit)	MAE ↓	RMSE ↓	R^2 ↑	RRMSE ↓
$R\star$	[1–25]	(10^{12} cm)	0.04	0.05	0.996	0.02
M_{ej}	[10–30]	($M_\odot$)	0.04	0.05	0.98	0.017
E_K	[0.5–9.5]	(foe)	0.06	0.08	0.98	0.05
$M_{^{56}Ni}$	[0.01–0.1]	($M_\odot$)	0.0003	0.0005	0.9997	0.009

Table 3
Performance evaluation of the 7 parameter characterisation model.

Parameter	[Range]	(Unit)	MAE ↓	RMSE ↓	R^2 ↑	RRMSE ↓
$R\star$	[0.1–10]	(10^{13} cm)	0.4	0.6	0.8	0.18
M_{ej}	[5–25]	($M_\odot$)	0.11	0.15	0.6	0.06
E_K	[0.5–5]	(foe)	0.2	0.3	0.7	0.18
$M_{^{56}Ni}$	[0.01–0.1]	($M_\odot$)	0.02	0.04	0.8	0.4
M_{CSM}	[0.05–0.8]	($M_\odot$)	0.04	0.08	0.8	0.5
R_{CSM}	[1–10]	(10^{15} cm)	0.18	0.3	0.7	0.18
S_{CSM}	[0–3]	(exponent)	0.17	0.2	0.6	0.3

pave the way for large-scale population studies and targeted follow-up of the most interesting events. This study focused exclusively on some subtypes of hydrogen-rich SNe, and further work could extend the analysis to a broader set of SN subtypes to assess the generality of the proposed framework. Additionally, the current model was trained to reproduce bolometric light curves only; an interesting direction would be to adapt it for generating in-band light curves, as well as related physical quantities such as velocity and temperature profiles. From a modelling perspective, the generator architecture could be enhanced by incorporating convolutional components, which may better exploit the continuous nature of the data and improve computational efficiency. Finally, a more randomised and diversified version of the training dataset could improve robustness to noise and observational uncertainties present in real data Directions for future work include testing the framework on a larger and more diverse sample of real data, extending the model to include different types of SNe and exploring more advanced architectures for function modelling.

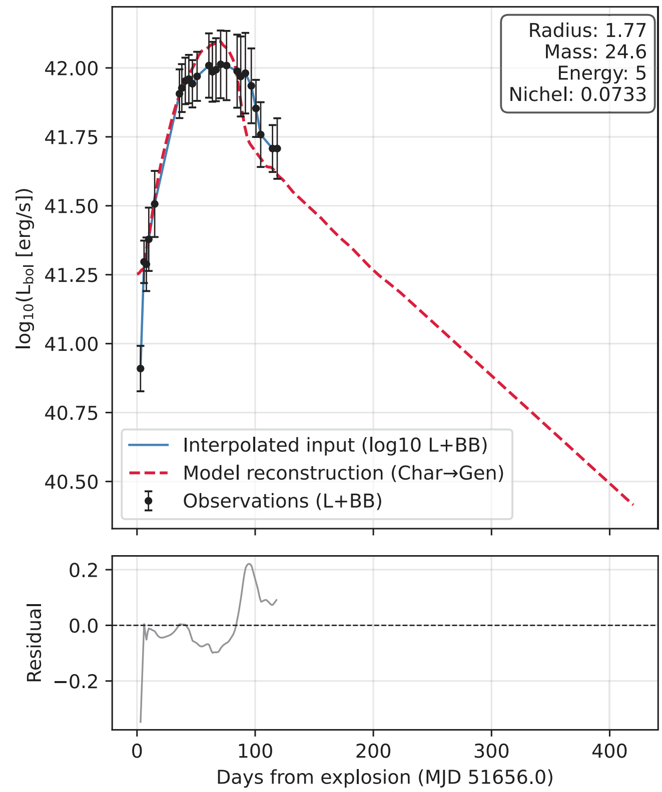

Fig. 10. SN2000cb.

7. Applicability on real data

In this section, we show some examples of characterisation and reconstruction of real bolometric LCs using the 4-parameter model to illustrate the model performance on some real SNe (see Figs. 10–17). Although the algorithm was trained exclusively on synthetic data generated under specific theoretical assumptions, which do not fully capture the complexity of real observations, the results are, in some cases, reasonably accurate. The discrepancies observed when applying the method to real data are more likely related to the physical assumptions underlying the theoretical models used to generate the synthetic training dataset, rather than to the machine learning method used for parameter regression. However, the scope of the current project is focused on the applicability and the scalability of the ML approximation rather than the physical accuracy of the theoretical models. The physical models provided have only been used as a POC to generate the synthetic dataset, while their applicability to real data has been widely discussed in the cited papers. In principle, the proposed ML framework is flexible and can be applied to any dataset, synthetic or real, provided that parameter labels or reliable estimates are available for training.

8. Conclusions

In this work, we present ASTRAI, a deep learning framework designed to address the challenges of SNe analysis in the era of big data. We developed and optimised a dual AI architecture capable of both generating high-fidelity synthetic LCs and rapidly characterising transient events by inferring their fundamental physical parameters. The key contributions of this project are the development of a comprehensive AI framework, covering the entire cycle from simulation to inference, the empirical demonstration of the capabilities of ML approaches and the creation of a robust and generalisable model through the systematic use of data augmentation techniques. In conclusion, the ASTRAI framework aims to be a powerful and scalable tool, capable

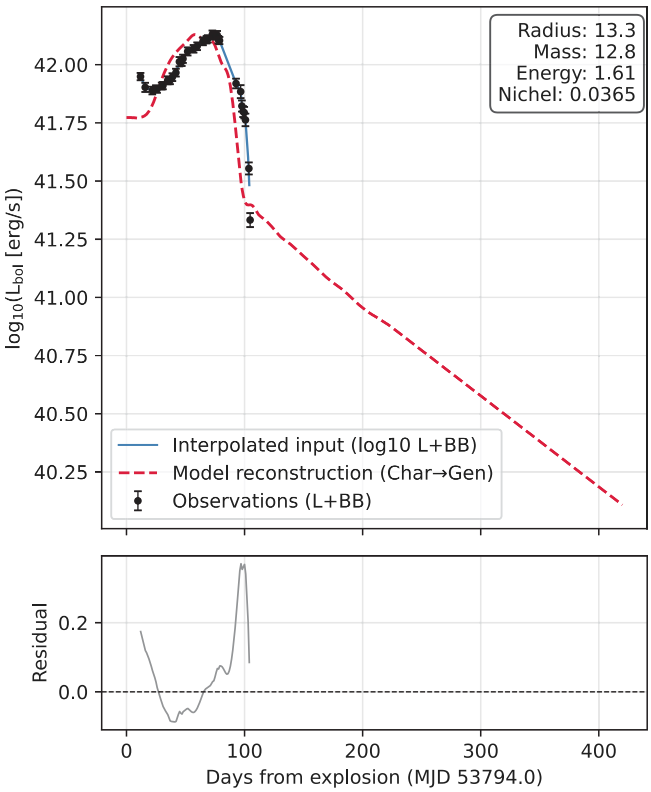

Fig. 11. SN2006au.

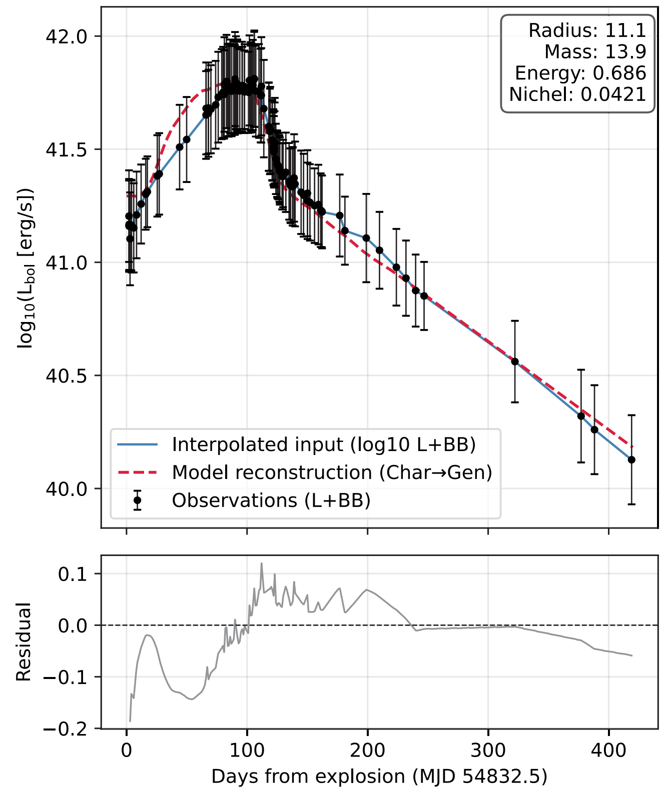

Fig. 13. SN2009E.

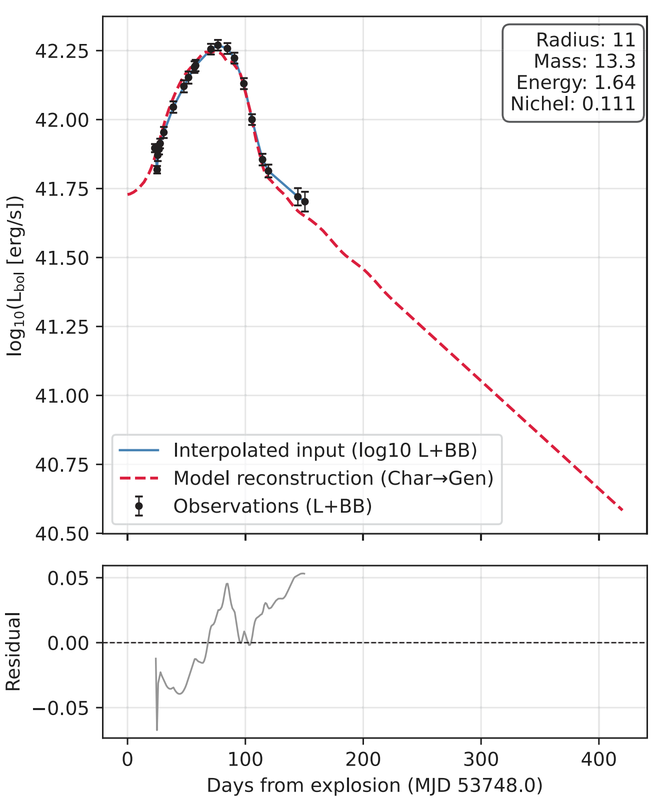

Fig. 12. SN2006V.

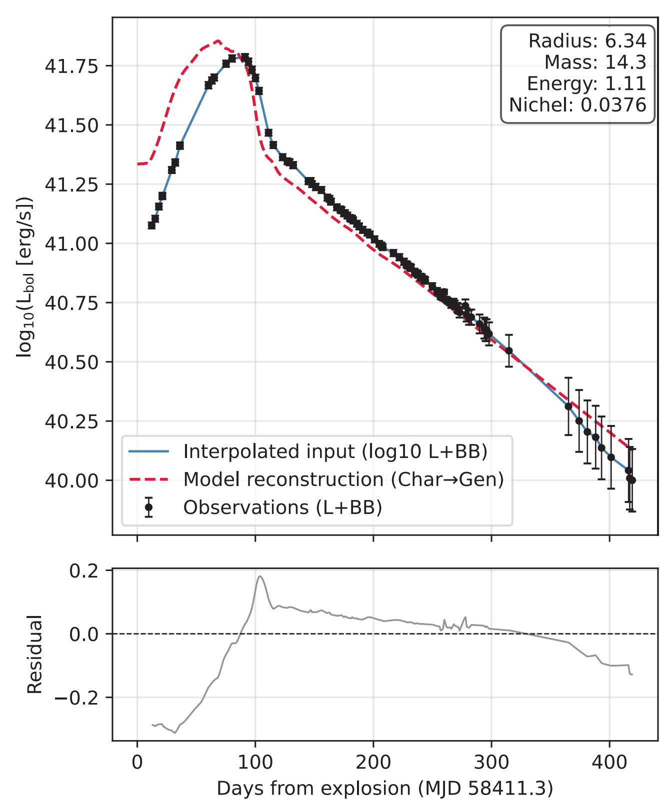

Fig. 14. SN2018hna.

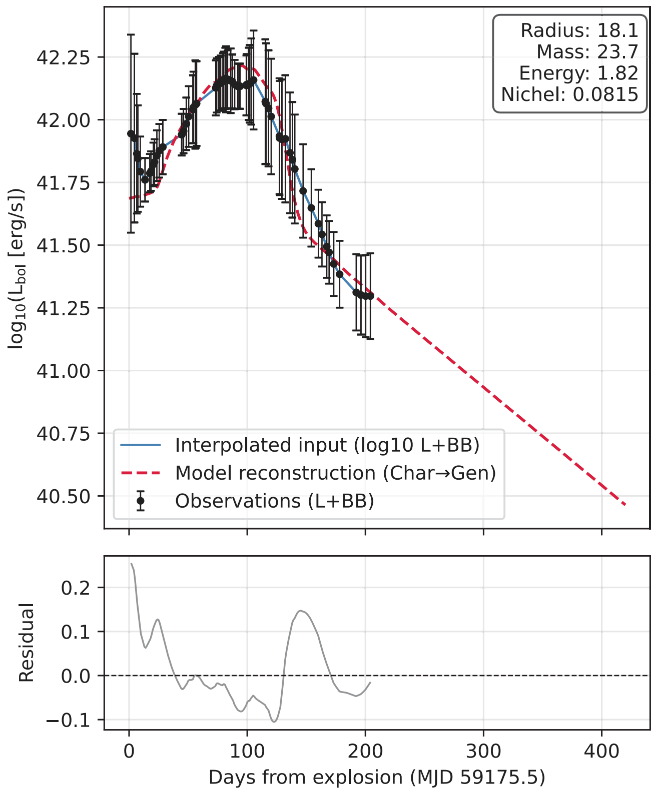

Fig. 15. SN2020bah.

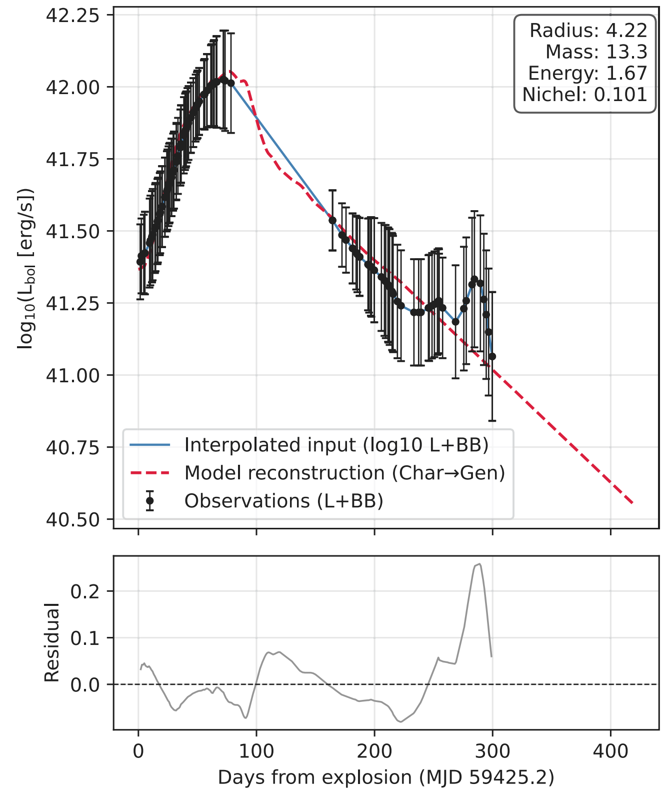

Fig. 17. SN2021wun.

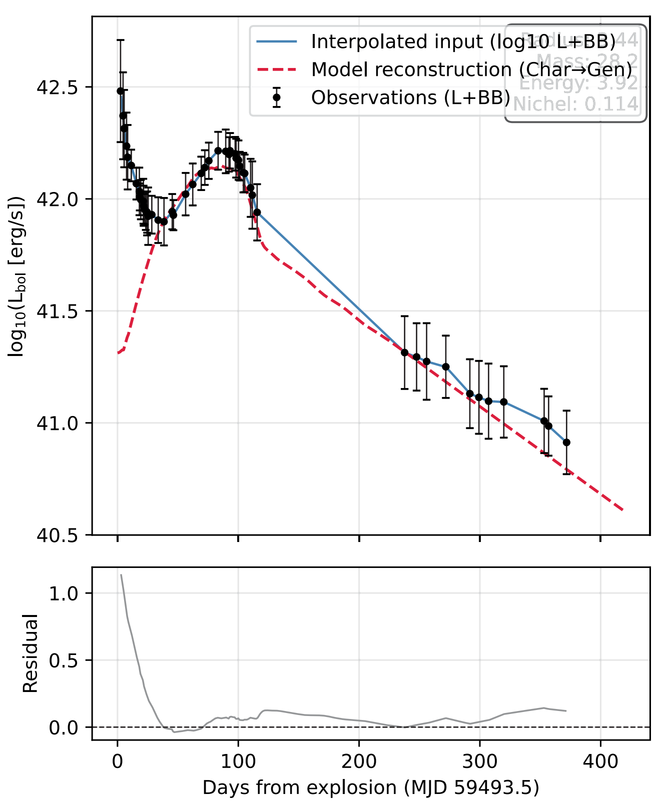

Fig. 16. SN2021aatd.

of supporting the scientific community in the analysis of future data streams, thus accelerating our understanding of transient phenomena in the universe.

CRediT authorship contribution statement

Andrea Claudio Grasso: Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Data curation. **Vincenzo Del Zoppo:** Writing – review & editing, Visualization, Supervision, Software, Methodology, Investigation, Conceptualization. **Alessio Mezzina:** Writing – review & editing, Validation. **Marco Cataldo:** Project administration. **Giuseppe Puglisi:** Software. **Fabio Spampinato:** Software, Writing – review & editing, Visualization. **Stefano Pio Cosentino:** Writing – review & editing, Writing – original draft, Validation, Methodology, Formal analysis. **Maria Letizia Pumo:** Writing – review & editing, Validation, Supervision. **Luca Naso:** Supervision, Funding acquisition.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors Vincenzo Del Zoppo and Andrea Claudio Grasso used ChatGPT4 and Grammarly solely to assist with language editing (grammar, spelling, and fluency). After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This paper is supported by the Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations, National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) "Italian Research Center on High-Performance Computing, Big Data and Quantum Computing" funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di "campioni nazionali di R&S (M4C2-19)" Next Generation EU (NGEU), and the work has been carried out within the ASTRAI project (Project ID CN_00000013). We are also grateful to the anonymous reviewers for their insightful feedback, which significantly improved the quality of this manuscript.

References

Abell, P.A., et al., 2009. LSST science book, version 2.0. doi:10.2172/1156415, arXiv:0912.0201.

Ahumada, T., Singer, L.P., Anand, S., Coughlin, M.W., Kasliwal, M.M., Ryan, G., Andreoni, I., Cenko, S.B., Fremling, C., Kumar, H., Pang, P.T.H., Burns, E., Cunningham, V., Dichiara, S., Dietrich, T., Svinkin, D.S., Almualla, M., Castro-Tirado, A.J., De, K., Dunwoody, R., Gatkine, P., Hammerstein, E., Iyyani, S., Mangan, J., Perley, D., Purkayastha, S., Bellm, E., Bhalerao, V., Bolin, B., Bulla, M., Cannella, C., Chandra, P., Duev, D.A., Frederiks, D., Gal-Yam, A., Graham, M., Ho, A.Y.Q., Hurley, K., Karambelkar, V., Kool, E.C., Kulkarni, S.R., Mahabal, A., Masci, F., McBreen, S., Pandey, S.B., Reusch, S., Ridnaia, A., Rosnet, P., Rusholme, B., Carracedo, A.S., Smith, R., Soumagnac, M., Stein, R., Troja, E., Tsvetkova, A., Walters, R., Valeev, A.F., 2021. Author Correction: Discovery and confirmation of the shortest gamma-ray burst from a collapsar. Nat. Astron. 5, 1179. doi:10.1038/s41550-021-01501-1.

Arnett, W.D., 1980. Analytic solutions for light curves of supernovae of Type II. Astrophys. J. 237, 541–549. doi:10.1086/157898.

Arnett, W.D., 1982. Type I supernovae. I - Analytic solutions for the early part of the light curve. Astrophys. J. 253, 785–797. doi:10.1086/159681.

Arnett, D., 1996. Supernovae and nucleosynthesis: An investigation of the history of matter from the big bang to the present.

Burhanudin, U.F., Maund, J.R., Killestein, T., Ackley, K., Dyer, M.J., Lyman, J., Ulaczyk, K., Cutter, R., Mong, Y.-L., Steeghs, D., Galloway, D.K., et al., 2021. Light curve classification with recurrent neural networks for GOTO: dealing with imbalanced data. Mon. Not. R. Astron. Soc. 1–16. doi:10.1093/mnras/stab1545.

Chatzopoulos, E., Wheeler, J.C., Vinko, J., 2012. Generalized semi-analytical models of supernova light curves. Astrophys. J. 746 (2), 121. doi:10.1088/0004-637X/746/2/121.

Cosentino, S.P., 2024. "Fast" Modeling Procedures for Core Collapse Supernovae and Similar Transient Objects (Ph.D. thesis). University of Catania, Department of Physics and Astronomy "Ettore Majorana", URL https://hdl.handle.net/20.500.11769/658009.

Cosentino, S.P., Inserra, C., Pumo, M.L., 2025a. Physical properties of long-rising type II Supernovae – Bayesian analytic modeling and spectrophotometric correlations. doi:10.48550/arXiv.2510.25018, ArXiv E-Prints arXiv:2510.25018.

Cosentino, S.P., Inserra, C., Pumo, M.L., 2025b. Physical properties of long-rising type II supernovae – Bayesian analytic modelling and spectrophotometric correlations. arXiv:2510.25018.

Cosentino, S.P., Pumo, M.L., Cherubini, S., 2025c. High-energy neutrinos by hydrogen-rich supernovae interacting with low-massive circumstellar medium: The case of SN 2023ixf. MNRAS 540 (4), 2894–2913. doi:10.1093/mnras/staf861, arXiv:2503.03699, URL https://ui.adsabs.harvard.edu/abs/2025MNRAS.540.2894C.

Demianenko, M., Malanchev, K., Samorodova, E., Sysak, M., Shiriaev, A., Derkach, D., Hushchyn, M., 2023. Understanding of the properties of neural network approaches for transient light curve approximations. Astron. Astrophys. 677, A16. doi:10.1051/0004-6361/202245189.

Fang, K., Metzger, B.D., Vurm, I., Aydi, E., Chomiuk, L., 2020. High-energy neutrinos and Gamma rays from nonrelativistic shock-powered transients. ApJ 904 (1), 4. doi:10.3847/1538-4357/abbc6e, arXiv:2007.15742.

Grassia, M., Cosentino, S., Pumo, M., Mangioni, G., 2025. Machine learning supernovae's progenitor characterization. pp. 439–442. doi:10.1109/PDP66500.2025.00068.

Grassia, M., Cosentino, S.P., Pumo, M.L., Mangioni, G., 2025b. Machine learning supernovae's progenitor characterization. doi:10.1109/PDP66500.2025.00068.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. CVPR, pp. 770–778. doi:10.1109/CVPR.2016.90.

Heger, A., Fryer, C.L., Woosley, S.E., Langer, N., Hartmann, D.H., 2003. How massive single stars end their life. ApJ 591 (1), 288–300. doi:10.1086/375341, arXiv:astro-ph/0212469.

Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning. PMLR, pp. 448–456, URL https://doi.org/10.48550/arXiv.1502.03167.

Janka, H.-T., 2017. Neutrino emission from supernovae. In: Alsabti, A.W., Murdin, P. (Eds.), Handbook of Supernovae. p. 1575. doi:10.1007/978-3-319-21846-5_4.

Lochner, M., McEwen, J.D., Peiris, H.V., Lahav, O., Winter, M.K., 2016. Photometric supernova classification with machine learning. Astrophys. J. Suppl. Ser. 225 (2), 31. doi:10.3847/0067-0049/225/2/31.

Maas, A.L., Hannun, A.Y., Ng, A.Y., et al., 2013. Rectifier nonlinearities improve neural network acoustic models. In: Proc. Icml. 30, (1), Atlanta, GA, p. 3, URL https://api.semanticscholar.org/CorpusID:16489696.

Mahabal, A., Sheth, K., Gieseke, F., Pai, A., Djorgovski, S.G., Drake, A., Graham, M., the CSS/CRTS/PTF Collaboration, 2017. Deep-learnt classification of light curves. In: 2017 IEEE Symposium Series on Computational Intelligence. SSCI, IEEE, p. 2757. doi:10.1109/SSCI.2017.8280984.

Möller, A., de Boissière, T., 2019a. SuperNNova: An open-source framework for Bayesian, neural network-based supernova classification. Mon. Not. R. Astron. Soc. 491 (3), 4277–4293. doi:10.1093/mnras/stz3312, arXiv:https://academic.oup.com/mnras/article-pdf/491/3/4277/31555240/stz3312.pdf.

Möller, A., De Boissière, T., 2019b. SuperNNova: An open-source framework for Bayesian, neural network based supernova classification. Mon. Not. R. Astron. Soc. 119. doi:10.1093/mnras/stz3312.

Moriya, T.J., Maeda, K., Taddia, F., Sollerman, J., Blinnikov, S.I., Sorokina, E.I., 2013. An analytic bolometric light curve model of interaction-powered supernovae and its application to Type IIn supernovae. Mon. Not. R. Astron. Soc. 435 (2), 1520–1535. doi:10.1093/mnras/stt1392.

Murase, K., Thompson, T.A., Lacki, B.C., Beacom, J.F., 2011. New class of high-energy transients from crashes of supernova ejecta with massive circumstellar material shells. Phys. Rev. D 84 (4), 043003. doi:10.1103/PhysRevD.84.043003, arXiv:1012.2834.

Nagao, T., Reynolds, T.M., Kuncarayakti, H., Cartier, R., Mattila, S., Maeda, K., Sollerman, J., Pessi, P.J., Anderson, J.P., Inserra, C., Chen, T.-W., Ferrari, L., Fraser, M., Young, D.R., Gromadzki, M., Gutiérrez, C.P., Lundqvist, P., Pignata, G., Müller-Bravo, T.E., Ragosta, F., Reguitti, A., Moran, S., González-Bañuelos, M., Kopsacheili, M., Petrushevska, T., 2025. Observational diversity of bright long-lived Type II supernovae. Astron. Astrophys. 699, A283. doi:10.1051/0004-6361/202554988, arXiv:2504.01427.

Nicholl, M., 2018. SuperBol: A user-friendly Python routine for bolometric light curves. Res. Notes the AAS 2 (4), 230. doi:10.3847/2515-5172/aaf799.

Perley, D.A., Fremling, C., Sollerman, J., Miller, A.A., Dahiwale, A.S., Sharma, Y., Bellm, E.C., Biswas, R., Brink, T.G., Bruch, R.J., De, K., Dekany, R., Drake, A.J., Duev, D.A., Filippenko, A.V., Gal-Yam, A., Goobar, A., Graham, M.J., Graham, M.L., Ho, A.Y.Q., Irani, I., Kasliwal, M.M., Kim, Y.-L., Kulkarni, S.R., Mahabal, A., Masci, F.J., Modak, S., Neill, J.D., Nordin, J., Riddle, R.L., Soumagnac, M.T., Strotjohann, N.L., Schulze, S., Taggart, K., Tzanidakis, A., Walters, R.S., Yan, L., 2020. The Zwicky transient facility bright transient survey. II. A public statistical sample for exploring supernova demographics*. Astrophys. J. 904 (1), 35. doi:10.3847/1538-4357/abbd98.

Popov, D.V., 1993. An analytical model for the plateau stage of type II supernovae. Astrophys. J. 414, 712–725. doi:10.1086/173117.

Pumo, M.L., Cosentino, S.P., 2025. Long-rising Type II supernovae resembling supernova 1987A - II. A new analytical model to describe these events. MNRAS 538 (1), 223–242. doi:10.1093/mnras/staf288, arXiv:2502.09752.

Pumo, M.L., Zampieri, L., 2011. Radiation-hydrodynamical modeling of core-collapse supernovae: Light curves and the evolution of photospheric velocity and temperature. ApJ 741 (1), 41. doi:10.1088/0004-637X/741/1/41, arXiv:1108.0688.

Pumo, M.L., Zampieri, L., 2013. Calibration relations for core-collapse supernovae. MNRAS 434 (4), 3445–3453. doi:10.1093/mnras/stt1256.

Radice, D., Morozova, V., Burrows, A., Vartanyan, D., Nagakura, H., 2019. Characterizing the gravitational wave signal from core-collapse supernovae. ApJ 876 (1), L9. doi:10.3847/2041-8213/ab191a, arXiv:1812.07703.

Sarin, N., Moriya, T.J., Singh, A., Gangopadhyay, A., Hinds, K.-R., Schulze, S., Omand, C.M.B., Das, K.K., 2025. Surrogate models for lightcurves and photosphere properties of Type II supernovae. Mon. Not. R. Astron. Soc. staf1808. doi:10.1093/mnras/staf1808.

Scannapieco, C., Tissera, P.B., White, S.D.M., Springel, V., 2008. Effects of supernova feedback on the formation of galaxy discs. MNRAS 389 (3), 1137–1149. doi:10.1111/j.1365-2966.2008.13678.x, arXiv:0804.3795.

Smartt, S.J., 2009. Progenitors of core-collapse supernovae. Annual Rev. Astron. Astrophysics 47 (Volume 47, 2009), 63–106. doi:10.1146/annurev-astro-082708-101737, URL https://www.annualreviews.org/content/journals/10.1146/annurev-astro-082708-101737.

Smith, N., 2014. Mass loss: Its effect on the evolution and fate of high-mass stars. ARA&A 52, 487–528. doi:10.1146/annurev-astro-081913-040025, arXiv:1402.1237.

Smith, D.S., Scalo, J., Wheeler, J.C., 2004. Transport of ionizing radiation in terrestrial-like exoplanet atmospheres. Icarus 171 (1), 229–253. doi:10.1016/j.icarus.2004.04.009, arXiv:astro-ph/0308311.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15 (1), 1929–1958, URL https://dl.acm.org/doi/10.5555/2627435.2670313.