

HOLISMOKES

XVII. Detecting strongly lensed type Ia supernovae from time series of multiband LSST-like imaging data

Satadru Bag^{1,2,*}, Raoul Cañameras^{3,2,1}, Sherry H. Suyu^{1,2}, Stefan Schuldt^{4,5}, Stefan Taubenberger^{1,2},
Irham Taufik Andika^{1,2}, and Alejandra Melo^{1,2}

¹ Technical University of Munich, TUM School of Natural Sciences, Physics Department, James-Franck-Straße 1, 85748 Garching, Germany

² Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Straße 1, 85748 Garching, Germany

³ Aix Marseille Univ, CNRS, CNES, LAM, Marseille, France

⁴ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy

⁵ INAF – IASF Milano, Via A. Corti 12, I-20133 Milano, Italy

Received 27 June 2025 / Accepted 1 November 2025

ABSTRACT

Strong gravitationally lensed supernovae (LSNe), though rare, are exceptionally valuable probes for cosmology and astrophysics. Upcoming time-domain surveys such as the Vera Rubin Observatory’s Legacy Survey of Space and Time (LSST) offer a major opportunity to discover large number of LSNe. Early identification is crucial for timely follow-up observations. We have developed a deep learning pipeline to detect LSNe using multiband, multi-epoch image cutouts. Our model is based on a 2D convolutional long short-term memory (ConvLSTM2D) architecture designed to capture both spatial and temporal correlations in time-series imaging data. Predictions are made after each observation in the time series, with accuracy expected to improve progressively as additional data are processed. We trained the model on realistic simulations derived from Hyper Suprime-Cam (HSC) data, which closely matches LSST in depth and filter characteristics. In this work, we focus exclusively on Type Ia supernovae (SNe Ia). LSNe Ia were injected into HSC luminous red galaxies (LRGs) at various phases of evolution to create positive examples of LSNe Ia time series. Negative examples include variable sources observed in the HSC Transient Survey (including unclassified transients) and simulated unlensed SNe Ia in LRG and spiral galaxies. Our multiband model shows rapid classification performance improvements during the initial few observations and quickly reaches a high detection efficiency: At a fixed false-positive rate (FPR) of 0.01%, the true-positive rate (TPR) reaches $\geq 60\%$ by the seventh observation and exceeds $\geq 70\%$ by the ninth observation. If we relax the FPR to 0.1%, the TPR reaches close to 60% as early as the fourth observation. Although the single-band analysis performs reasonably well in isolation, the multiband model significantly outperforms it, particularly in the early stages, by building a richer memory and leveraging color information. Among the negative examples, SNe in LRGs remain the primary source of FPR, as they can resemble their lensed counterparts under certain conditions. Additionally, the model detects quads more effectively than doubles, and it performs better on systems with larger image separations. Although we trained and tested the model on HSC-like data, our approach is applicable to any cadenced imaging survey – particularly LSST, where the higher expected cadence (five to ten times that of HSC) should further boost performance.

Key words. gravitational lensing: strong – supernovae: general

1. Introduction

Strong gravitational lensing occurs when the gravitational potential of a massive galaxy or galaxy cluster bends light from background sources and produces multiple magnified images, including elongation of extended sources such as another galaxy (see, e.g., Narayan & Bartelmann 1996 for a pedagogical introduction). Acting as a cosmic telescope, this effect enables detailed studies of the initial mass function (IMF), galaxy evolution, and the nature of dark matter (Mao & Schneider 1998; Metcalf & Madau 2001; Dalal & Kochanek 2002; Pooley et al. 2009; Oguri et al. 2014; Jiménez-Vicente & Mediavilla 2019; Shajib et al. 2024; Vegetti et al. 2024). Among its most significant applications is time-delay cosmography (TDC) (Refsdal & Bondi 1964; Refsdal 1964; Saha et al. 2006; Oguri

2007; Treu & Marshall 2016; Bonvin et al. 2017; Wong et al. 2020; Birrer et al. 2020; Birrer & Treu 2021; Treu et al. 2022; Suyu et al. 2024; TDCOSMO Collaboration 2025), which uses time delays between multiple images of variable sources to estimate cosmological distance ratios. These measurements directly constrain the Hubble constant (H_0), independent of local distance ladder methods – thus addressing the on-going H_0 tension (Aghanim et al. 2020; Riess et al. 2022; Di Valentino et al. 2021; Verde et al. 2024). They also provide crucial insights into dark energy and spatial curvature when there are multiple strongly lensed sources at different redshifts (e.g., Grillo et al. 2024) and when combined with other cosmological probes (e.g., Linder 2011; Suyu et al. 2012b; Jee et al. 2016). Time-delay measurements require time-variable sources, such as lensed quasars (QSOs) and supernovae (SNe), though other transients such as kilonovae (KNe), tidal disruption events (TDEs), gamma-ray bursts (GRBs), and fast radio bursts (FRBs) are

* Corresponding author: satadru.bag@tum.de

increasingly gaining attention. For further applications of strong lensing, see the review by [Treu \(2010\)](#).

Lensed supernovae (LSNe) and QSOs each offer distinct advantages for TDC. LSNe are extremely rare, with only a handful known to date ([Kelly et al. 2015](#); [Goobar et al. 2017](#); [Rodney et al. 2021](#); [Chen et al. 2022](#); [Goobar et al. 2023](#); [Frye et al. 2023](#); [Pierel et al. 2024](#)), but they are expected to become pivotal in the coming decade ([Suyu et al. 2024](#)). In contrast, lensed QSOs have been the mainstay of TDC due to their abundance. However, despite extensive monitoring by COSMO-GRAIL and TDCOSMO ([Millon et al. 2020a,b](#)), only a handful have enabled precise H_0 measurements. The H0LICOW team measured H_0 with 2.4% uncertainty using six lensed QSOs ([Wong et al. 2020](#)) based on two families of mass models (a power-law model and a baryon-plus-dark-matter model), where individual lenses constrained H_0 at the 5–10% level. This underscores the need for significantly larger samples, especially since the uncertainty increases when standard assumptions about the lens mass density profile are relaxed ([Birrer et al. 2020](#)). The recent TDCOSMO results from a sample of eight lensed quasars with relaxed assumptions on mass density profile yielded a $\sim 5\%$ uncertainty on H_0 (TDCOSMO Collaboration 2025). To achieve percent-level precision on H_0 through LSNe, independent of lens model assumptions, it is essential to discover a few tens to a hundred new LSNe, thereby building a gold sample of lens systems for cosmography ([Suyu et al. 2020](#); [Arendse et al. 2024](#)).

Lensed supernovae also offer a powerful way to probe the nature of their progenitors. For instance, thermonuclear supernovae (Type Ia, SNe Ia hereafter) are believed to be originated from the disruption of a carbon–oxygen white dwarf in a binary system, but the precise nature of their progenitor systems remains a topic of active debate (e.g., [Maoz et al. 2014](#)). A major limitation to probe the nature of the exploding star and its potential companion is the lack of data sufficiently close to explosion and at sufficiently blue wavelengths. LSNe provide a unique opportunity to overcome this challenge by offering critical early-time observational access. Once the first image of a LSN is detected, mass modeling can be used to predict the time delays and the appearance of future images, enabling coordinated follow-up efforts to capture the SN evolution within the first hours or days.

The Legacy Survey of Space and Time (LSST) of the Vera C. Rubin Observatory ([LSST Science Collaboration 2009](#); [Ivezić et al. 2019](#)), along with other wide-field surveys such as the Euclid space mission ([Euclid Collaboration: Mellier et al. 2025](#)) and the Roman Space Telescope ([Spergel et al. 2015](#)), is expected to significantly expand the sample of lensed transients. LSST alone is projected to detect hundreds of LSNe ([Sainz Murieta et al. 2023](#); [Arendse et al. 2024](#); [Bag et al. 2024](#)) and thousands of lensed QSOs ([Oguri & Marshall 2010](#)) over its 10-year survey. The scientific exploitation of these future promising samples of LSNe nonetheless requires a rapid identification, ideally before peak, to enable their mass modeling and follow-up on coordinated facilities. As an example, reaching the main goals of our Highly Optimized Lensing Investigations of Supernovae, Microlensing Objects, and Kinematics of Ellipticals and Spirals (HOLISMOKES, [Suyu et al. 2020](#)) program requires not only spectroscopic typing and high-cadence imaging for time-delay and H_0 measurements, but also rest-frame UV spectroscopy in the early phases for progenitor studies. The sheer volume of alerts from wide-field surveys (e.g., $\mathcal{O}(10^7)$ per night from LSST) makes automated LSNe identification essential, with machine learning (ML) as the most practical solution.

Lensed supernovae can be detected in multiple ways. The most straightforward method is based on image multiplicity, where one directly resolves multiple images of an LSN ([Oguri & Marshall 2010](#)). Another approach, particularly effective for SNe Ia due to their well-understood luminosity after standardization, relies on photometry. In this case, magnification-based selection is used to identify candidates whose apparent brightness exceeds what is expected at their redshifts, even when the angular resolution is insufficient to deblend multiple images (e.g., [Goobar et al. 2017](#); [Goldstein et al. 2019](#); [Sagués Carracedo et al. 2024](#)). Additionally, color–magnitude diagrams offer another selection tool for SNe Ia, as LSNe Ia tend to appear redder due to their typically higher redshifts compared to their unlensed counterparts ([Quimby et al. 2014](#)). Recently, efforts have also been made to identify unresolved LSNe using the shape of their blended light curves ([Bag et al. 2021](#); [Denissenya et al. 2022](#); [Denissenya & Linder 2022](#); [Bag et al. 2024](#)).

Another fundamentally different strategy is catalog cross-matching. In this approach, one assembles a large sample of static strong lenses and monitors them for appearance of transients at the positions of multiple lensed images of background galaxies (e.g., [Shu et al. 2018](#); [Craig et al. 2024](#)). In previous studies, we leveraged the capabilities of supervised CNNs in selecting static galaxy-scale strong lenses from deep, wide-field surveys (e.g., [Petrillo et al. 2017](#); [Jacobs et al. 2019](#); [Metcalf et al. 2019](#)), and developed automated methods to systematically identify these systems over large footprints, with minimal false-positive rates (FPRs) (e.g., [Cañameras et al. 2020](#); [Schuldt et al. 2025](#); [Cañameras et al. 2024](#)). Applying such algorithms to LSNe detection requires the SN host galaxy to be detectable and sufficiently deblended from the foreground lens galaxy.

Here, we aim to extend the discovery space of LSNe to include systems without detectable hosts which are expected to constitute a significant fraction (e.g., around half of LSNe Ia observed by LSST [Sainz Murieta et al. 2024](#)), and with smaller image separations, while pushing the detection as early as possible. Time-domain surveys such as LSST (and Roman in the future) will produce multiband 2D image sequences of transients, capturing both spatial features (e.g., multiple images of lensed transients and arcs from their lensed hosts) and temporal variations (following specific light curves). In this work, we develop a deep learning-based pipeline to detect LSNe directly from the alert stream of the upcoming time-domain surveys. Specifically, we employed Convolutional Long Short-Term Memory (ConvLSTM) architecture, which capture spatial and temporal correlations simultaneously across sequences of imaging data. Unlike previous approaches that convolve reference images and apply LSTMs to the corresponding light curves ([Morgan et al. 2022, 2023](#)), we employed ConvLSTM directly to the full image time series, capturing spatio-temporal signatures such as centroid shifts and extended residuals to distinguish LSNe from other transients.

[Kodi Ramanah et al. \(2022\)](#) previously explored ConvLSTM-based architectures on fully simulated, single-band datasets, demonstrating that using the time series of images instead of a single-epoch image significantly improves the classification performance. Our work introduces two key improvements. First, we extend the framework to multiband imaging, leveraging richer color and temporal information for improved early detection. Second, instead of relying solely on synthetic data, we train and validate the model on realistic simulations based on actual Hyper Suprime-Cam (HSC) observations ([Aihara et al. 2018](#)),




Fig. 1. Observation epochs of the HSC Transient Survey in the COSMOS field across different filters during MJD 57710–57875 (Nov. 2016–Apr. 2017), used to extract image time series of variable sources (HSC variables) with highest available cadence. The number of observations in the *griz* bands are 5, 9, 12, and 12, respectively. The observation details are taken from Table 1 of Yasuda et al. (2019).

which closely match LSST in depth and filter characteristics. In particular, we simulate both lensed and unlensed/normal SNe, inject them into HSC images of galaxies (LRGs and spirals), and construct the resulting time series. We further incorporate real transients from the HSC Transient Survey (Yasuda et al. 2019), drawing on a catalog of over 100 k variable sources identified in the same footprint by Chao et al. (2021) using the difference imaging pipeline of Chao et al. (2020). This strategy trains the model to identify LSNe early using only the first few observations, even amid diverse transients and realistic image artifacts. In this initial work, we focus on detecting LSNe Ia without host light, providing a proof-of-concept architecture. While our model is trained and evaluated on HSC-like data in this study, the methodology is general and transferable to any time-domain imaging survey. Its design is especially compatible with LSST, whose significantly higher expected cadence (five to ten times that of HSC) is expected to further enhance the performance. An immediate objective is to enable real-time identification from the LSST alert stream.

The paper is organized as follows. In the next section, we describe how the datasets are constructed and how our simulations of different components used for training are anchored in real HSC observations. Section 3 details the model architecture and data processing. Key results are presented in Section 4. Section 5 summarizes the main conclusions of this work, and discusses the assumptions made, current limitations, and outlines future directions, particularly with an eye toward application to LSST.

Throughout this study, we use cutout sizes of 59×59 pixels corresponding to $\sim 10'' \times 10''$ physical size. We adopt the flat concordant Λ CDM cosmology with $\Omega_M = 0.308$ and $\Omega_\Lambda = 1 - \Omega_M$ (Planck Collaboration XIII 2016) and with $H_0 = 72 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Bonvin et al. 2017).

2. Data acquisition and simulation methodology

2.1. The HSC Transient Survey

The HSC Transient Survey (Yasuda et al. 2019) is a cadenced survey conducted with the 8-m Subaru telescope, on the two ultra-deep fields of the HSC Subaru Strategic Program (HSC-SSP, Aihara et al. 2018), namely COSMOS (Scoville et al. 2007) and SXDS (Subaru/XMM-Newton Deep Survey, Furusawa et al. 2008). In this study, we focus on the sin-

gle pointing of 1.77 deg^2 over the COSMOS field, corresponding to tract 9813, which was observed between Nov. 2016 and Apr. 2017.

With exposure times of a few 10^3 sec per epoch, the HSC Transient Survey reaches median depth per epoch of 26.4, 26.3, 26.0, and 25.6 mag in *g*-, *r*-, *i*-, and *z*-bands, respectively (Yasuda et al. 2019). To artificially degrade the depth and get closer to LSST single-epoch image depth, we produced an alternative version of the data set based on a subset of exposures available per epoch. Images reduced and calibrated by the HSC pipeline (Bosch et al. 2018), divided into 9×9 patches, and with a pixel size of $0.168''$, were imported using the Data Archive System search form¹.

We used these “warped” images, which are saved on a common sky grid prior to stacking, to produce shallower images per epoch. For simplicity, we restricted our data set to patches covered systematically at all epochs from the HSC Transient Survey, thereby excluding a few patches around the tract borders due to dithering. We also focused on exposures taken at MJD between 57710 and 57875 with homogeneous sets of filters. Three individual exposures of 300 sec were stacked per epoch, in *griz* bands. Pixels having flag values >1000 in any of the three exposures were masked out during stacking to avoid introducing artifacts near bright stars. The resulting frames were stitched together with SWarp (Bertin et al. 2002) to produce the final single-epoch images over the complete footprint.

We obtained a total of 5, 9, 12, and 12 epochs of observation in *g*, *r*, *i*, and *z* bands, respectively, over the six-month duration of the survey as shown in Figure 1. As a result of using only three of the available exposures, larger regions between HSC CCDs are masked out compared to the images released in Yasuda et al. (2019). Effective exposure times per pixel are variable due to the masking process.

2.2. Positive examples: Simulation of mock LSNe Ia time series

To simulate galaxy-scale LSNe, we inject point-like SN images into real HSC images of LRGs and place them in blank regions of the HSC Transient Survey. We use LRGs as foreground deflectors, given their high lensing cross sections². In this work, we focus on Type Ia SNe, as their intrinsic luminosities and light curves are well understood and straightforward to model. Extending the pipeline to other SN types and transients is conceptually straightforward and will be pursued in later phases by incorporating the appropriate light curve models during injection.

We adopted the redshift distributions of lenses and sources based on the Oguri & Marshall (2010, hereafter OM10) simulation framework (Oguri & Marshall 2010; Oguri 2018). First, we ensured that the redshift distribution of our selected LRG sample, used as lenses, broadly aligns with the lens redshift distribution for LSNe Ia in OM10. Instead of following the realistic Einstein radius distribution as in OM10, we adopted a flat

¹ https://hsc-release.mtk.nao.ac.jp/das_search/pdr3/index.html

² We have not included spiral lenses, another class of galaxy deflector, for several reasons. They are not straightforward to model, for example via velocity dispersion. Compared to LRGs, they are typically less massive, thereby producing smaller image separations and magnifications, and are intrinsically rare. For these reasons, including them in the training set may not be worthwhile and could even increase false positives. Even without explicit training on spiral lenses, the model is expected to identify some SNe lensed by spirals using other indicators, such as lensing-like configurations, color, and temporal evolution.




Fig. 2. Einstein radius (θ_E) distributions. The blue histogram shows the $\sim 93\,000$ lens–source pairs with a flat θ_E distribution between $0.1''$ and $2.0''$. After removing artifacts, $\sim 80\,000$ clean systems (orange) remained that still follow a flat θ_E distribution. The green histogram shows the final $\sim 50\,000$ mock LSNe Ia passing our brightness selection and restricted to $\theta_E > 0.5''$, which delimits the region of interest in this study. Of these, $\sim 48\,000$ are used in the training, validation, and test sets with a balanced number of doubles and quads.

distribution between $0.1''$ and $2.0''$ to ensure a more balanced training set. Based on these choices, we construct a lens–source matched catalog by assigning a source redshift to each lens, which determines at what redshift the SN Ia will later be injected. In Figure 2, we show the Einstein radius (θ_E) distribution of the lens–source catalog before and after removing entries with artifacts in the LRG cutouts, represented by the blue and orange histograms, respectively. Both distributions follow the intended flat trend across the θ_E range. In this initial study, we focus on Einstein radii of $0.5'' \leq \theta_E \leq 2.0''$, thereby excluding low-separation systems below the LSST’s resolution limit. The range shown in the figure reflects this selection. Note that restricting to $\theta_E \geq 0.5''$ reduces the detectable LSNe Ia rates in LSST only modestly, by a factor of ~ 2 – 3 depending on the detection criteria (Wojtak et al. 2019; Arendse et al. 2024; Bag et al. 2024).

The next step involves determining the optimal lensing configuration (image positions θ , time delays Δt , and magnifications μ) for each lens–source pair. We then inject the SN Ia images at different phases of the explosion, placing the SN at the assigned source redshift and applying the lensing configuration accordingly to generate a time series capturing its evolution. The top panels of Figure 3 show time series of two examples: a quadruply lensed system on the left and a doubly lensed system on the right. More examples, systematically arranged by different θ_E , are shown in Figure A.2 in Appendix A.

A detailed description of the simulation process is provided in Appendix A. The following are a few key points:

- In rare cases, we fail to find a viable lensing configuration within a reasonable number of trials that produces sufficiently bright images for detection; such lens–source pairs are discarded.
- Here, we also restrict ourselves to mock LSNe Ia with sufficient image brightness for detection, exceeding $5\sigma_B$, where σ_B is the standard deviation of the background noise. This selection criterion, combined with the previously mentioned cut on angular separation ($0.5'' \leq \theta_E \leq 2.0''$), yields a total

of approximately 50,000 successful mock LSNe Ia, whose θ_E distribution is shown by the green histogram in Figure 2. Since this sample contains a slight mismatch in the number of doubles and quads, we include a subset of $\sim 48,000$ systems with equal representation in our final datasets, as shown in Table 1. Note that systems with higher θ_E tend to correspond to sources at higher redshifts, which are more likely to fail the brightness threshold and hence are slightly underrepresented in the selected sample. This explains the decline in the higher θ_E bins. However, this bias is not necessarily unfavorable. As we discuss in the results section, systems with smaller θ_E are intrinsically harder to detect, so their slight over-representation may help the model learn to better identify these challenging cases.

- We do not include the SN host galaxy at this stage, regardless of whether it would be lensed or detectable. This reflects a significant fraction of cases (e.g., roughly half of LSST cases) where the host galaxy of an LSNe Ia is either not strongly lensed or too faint to be observed in the surveys. In future work, we will systematically evaluate the impact of detectable lensed hosts by examining different SN-to-host center offset distributions and detectability criteria.

2.3. Adding blanks to the cutouts at different epochs

For the synthetic time series (i.e., for mock LSNe and normal unlensed SNe exploding in galaxies, explained later in this section), the images vary over time due to the injected LSN images or the SN itself. The outer pixels remain unchanged, leading to identical noise across epochs. To address this, we add white noise (based on the sensitivity of the telescope) at each epoch. For this, we also collect single-epoch cutouts (referred to as “blanks” in Table 1) from the empty regions in the HSC COSMOS field at different epochs, separately for the four bands of interest. These blanks are added to all epochs of the time series for all components, including the fully observed time series of HSC variables, to ensure consistent noise characteristics across the entire dataset.

In reality, there would be Poisson noise associated with the measurements. In our case, the added white noise at each time step is of the order of Poisson noise near the bright central part and dominates at the outskirts. Therefore, for simplicity, we do not add Poisson noise to the coadded layer, and leave this for future work.

Since the standard deviation of the sum of N normal distributions is $\sigma_{\text{sum}} = \sqrt{\sum_i^N \sigma_i^2}$, by adding the blanks at different epochs, we reduce the S/N by a factor of $\sqrt{2}$. This reduction is unavoidable if we want to ensure that the noise in the synthetic time series (of lensed and normal SNe) varies across different epochs without separating the signal from the noise at the pixel level. We stress on the fact that adding blanks is important because, without this step, the cutouts would have identical noise pattern in the outskirts at different time steps, which is unphysical and could cause the model to learn from this repeated noise rather than solely from the underlying signal.

We consider different types of objects as the negative examples. These components are described in the next two sections.

2.4. Negative examples: Variable sources observed in the HSC Transient Survey

Chao et al. (2021) compiled a catalog of 101 353 variable sources based on multiband photometry from the HSC Transient

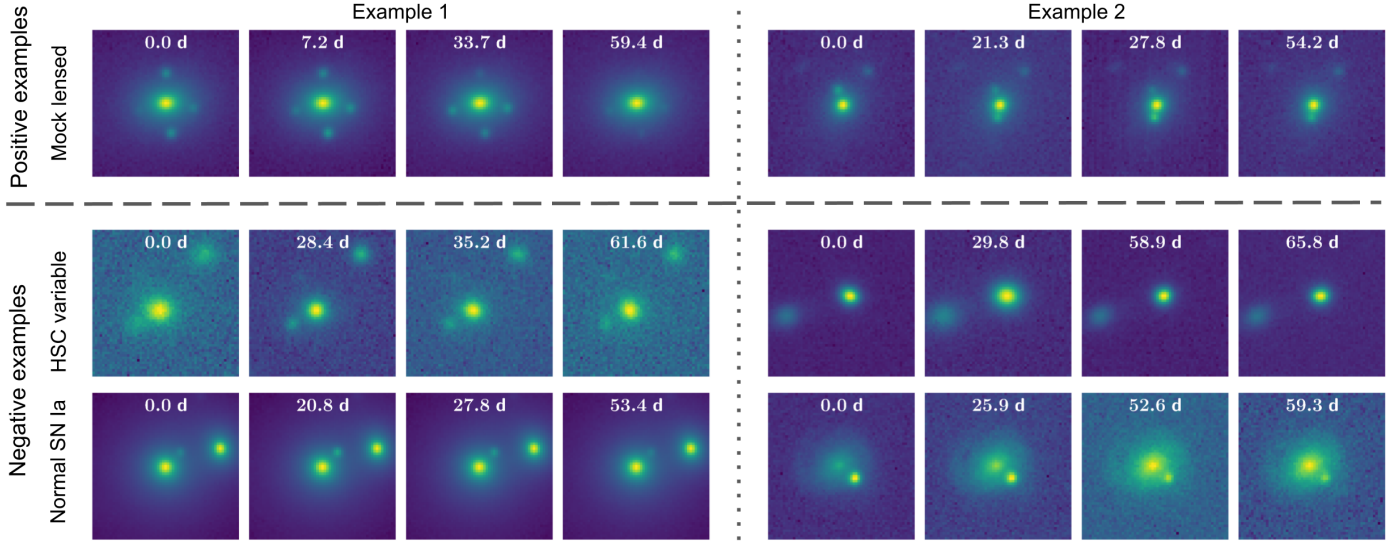


Fig. 3. Example *i*-band time series of different components used for training: mock LSNe Ia (top row, forming the positive class), and variable sources and normal SNe Ia (middle and bottom rows, forming the negative class). Each row shows two representative samples. The top row includes a quadruply LSN Ia (left) and a doubly lensed one (right). The middle row shows two unrelated HSC variables, where the central objects exhibit variability over time. The bottom row shows an SN Ia in an LRG (left) and in a spiral galaxy (right). Note that, only the time series of the HSC variables in the middle row include PSF variation over time. Timestamps in each frame indicate days since the first detection, which is always set to zero. The classification task is binary: distinguishing LSNe Ia from all other types of transients and bogus detections. For illustration, we show only single-band (*i*-band) time series here. The multiband time series corresponding to the quadruply LSN Ia shown in the top-left panel is presented in Figure 5.

Table 1. Ingredients from HSC observations.

Objects	from	# Available	# used in
LRGs	Wide layer (co-added)	~80 000 (after augmentation) (<i>griz</i>)	LSNe Ia: ~48 000 negatives: SN in LRGs: ~28 000
Blanks	HSC Transient Survey in COSMOS field (single epoch)	<i>g</i> : ~ 1200 (5 epochs) <i>r</i> : ~ 1200 (9 epochs) <i>i</i> : ~ 1400 (12 epochs) <i>z</i> : ~ 1500 (12 epochs) see Figure 1 for the epoch details	all in LSNe Ia & negatives
HSC variables from Chao et al. (2021)	HSC Transient Survey in COSMOS field (single epoch)	~72 000 (after augmentation) with ‘good’ light curves (<i>griz</i>)	negatives: ~28 000
Spirals	wide layer (co-added)	~40 000 (unique) (<i>griz</i>)	negatives: SN in spirals: ~16 000

Survey to run their lensed QSO search algorithm. These sources, referred to as ‘HSC variables’ hereafter, are unclassified and may include a diverse mix of real transients (e.g., variable stars, SNe, GRBs, KNe, TDEs, and rare lensed events) as well as bogus or non-astrophysical detections. We aim to put these HSC variables into our negative samples. However, as we need multiple epochs from HSC Transient Survey, not all of them could be used. We collected all the variables that are present in our HSC warps. Figure 1 shows the HSC Transient Survey observation dates in different bands. Details of the extraction process are provided in Section 2.1. The key points are summarized below:

- We concentrate on HSC variables from the 2017 window, where the cadence is most favorable, though still not ideal. This yields a maximum of 5, 9, 12, and 12 usable epochs

in the *g*, *r*, *i*, and *z* bands, respectively. However, not all HSC variables have good-quality cutouts in every epoch. We retain only those with sufficiently dense sampling, i.e., “good cadence”, as defined in Appendix B.

- After applying these selection criteria, we obtain approximately 9000 HSC variables with good multiband cadence. To further increase the dataset size, we apply data augmentation by performing four 90° rotations and two mirror flips for each rotated instance, resulting in an 8× augmentation per original sample. This gives us a final dataset of roughly 72 000 time series for the HSC variables, sufficient for our training, validation, and test sets.
- As noted in Section 2.3, we also add blanks (primarily white noise) to all epochs of the HSC variable time series to maintain

consistency across all components. The i -band time series of two unclassified HSC variables are shown in the middle row of Figure 3 as examples, where the centrally located objects exhibit variability over time.

2.5. Negative examples: Simulated normal SNe Ia exploding in LRGs and spirals

Since the HSC variables discussed above are unlabeled, i.e., they are not further classified into specific subcategories, they could potentially include lensed and unlensed (normal) SNe, among other types of transients. However, one of the most significant sources of confusion in distinguishing LSNe Ia arises from normal SNe. To improve the model’s training and enhance its ability to differentiate these cases, we increase the number of normal SNe in the negative samples by explicitly incorporating systems where an SN explodes in LRGs and spiral galaxies.

2.5.1. SNe Ia in LRGs

We have approximately 28 000 LRGs that are not used in creating the mock LSNe. We use these as hosts for unlensed SN Ia. To place the supernovae on the galaxy cutouts, we require: (i) the galaxy redshifts, which are also used as the supernova redshifts, (ii) point spread functions (PSFs), (iii) the distance from the galactic center, and (iv) the corresponding light curve. Since these LRGs are already cross-matched with the SDSS catalog, we have their spectroscopic redshifts. We also have their HSC PSFs, which will be used to paint the point-like supernovae onto the galaxy cutouts. For simplicity, we use the Hsiao template (Hsiao et al. 2007) to vary the supernova brightness over time for a given redshift and filter. Regarding the distance from the galactic center, we distribute the supernovae uniformly within two times the Kron radii of the LRGs. We adopt a uniform distribution in SN-host center offset to ensure equal representation across the training set, analogous to the uniform sampling in Einstein radius (θ_E) used for the mock LSNe samples, while spanning the range of realistic SN locations within host galaxies³. Additionally, we clip the distance at $2''$ to maintain consistency with the mock LSNe samples, where the maximum Einstein radius (θ_E) is $2''$, and to avoid placing supernovae too far from the galactic center. Just like the LSNe Ia, we generate the time series of cutouts at different phases of the SN Ia evolution.

2.5.2. SNe Ia in spirals

We also collect around 40 000 spiral galaxies from the HSC wide layer, which are used exclusively to host SN Ia. We follow the same procedure as described above for the LRGs, with a key difference: we use photometric redshifts instead of spectroscopic redshifts (as the latter is not available for all spirals).

³ We sample SN-host center offsets uniformly to enable the model to recognize SNe across all offsets without introducing bias. While astrophysically motivated priors based on stellar or light density profiles exist (Lokken et al. 2023), their exact dependence on galaxy type, redshift, and other factors is not fully established. Moreover, our primary goal at this stage is to demonstrate lensed SN detection among all alerts, rather than to optimize classification between lensed and normal SNe. In the final pipeline, however, significant false alarms are expected to arise from cases where the system is a genuine strong lens but the transient occurs in the foreground lens galaxy rather than in the background lensed host. Mitigating this will require fine-tuning with more realistic offset distributions.

As explained in Section 2.3 above, we add blanks (primarily white noise) to the images across epochs to introduce realistic noise variation over time for both of SN in LRGs and spiral galaxies. The bottom row of Figure 3 shows i -band time series of normal SNe Ia exploding in an LRG (left) and in a spiral galaxy (right), as illustrative examples.

Note that in the synthetic time series (for both lensed and normal SNe), we do not incorporate seeing variations for simplicity. Such variations are naturally present in the fully observed time series of HSC variables. As PSF variations are expected to have minimal impact on our method, we defer their inclusion to future work.

2.6. Matching data quality

As outlined earlier, our dataset consists of two fundamentally different types of samples: (i) simulated time series, including mock LSNe Ia and normal SNe Ia exploding in LRGs and spiral galaxies, and (ii) real, fully observed time series of variables from the HSC Transient Survey. To ensure consistency across different components, we match the data quality in terms of depth, zero-point, and cadence. For the simulated samples, we have complete control over the temporal sampling, allowing us to freely define the cadence. However, for the observed HSC variables, the cadence is fixed by the actual observing schedule, effectively “etched in stone”, leaving no flexibility.

To ensure homogeneity between the simulated and observed time series in the dataset, it is essential that their cadences match. This means we must adjust the sampling of the simulated time series to mirror the cadence of the HSC Transient Survey. In practice, we adopt the HSC Transient Survey cadence from the 2017 observing season as our reference and enforce this distribution across all time series used in training and testing, including the simulations. This step is crucial for aligning the temporal resolution of inputs across the dataset and ensuring that the model does not inadvertently learn spurious features associated with differences in cadence. It is important to emphasize that this choice of cadence is driven by practical constraints imposed by the observed data, and it differs significantly from the cadence anticipated for LSST.

The details of how we match the cadence distributions for all components (simulated mock LSNe Ia, normal SNe Ia exploding in galaxies, and observed HSC variables) are explained in Appendix B. Overall, the cadence distributions for all components match well across bands, as seen in Figure B.1, with a slight difference in the LSNe samples due to the additional criterion for capturing at least the second image.

2.7. First detection

In the final training, validation, and test sets, the first detection of the simulated time series (i.e., for the synthetic lensed and normal SNe Ia time series) is randomly selected within $(-20, 0)$ days relative to the peak of the first image for lensed cases or the SN light curve for the normal cases. In contrast, the first detection of the HSC variables, derived from observations in the HSC Transient Survey, is determined by the data and cannot be manually adjusted. Since the phase of the SN at a given time is unknown in real observations, we ensure that the timestamps (i.e., temporal values for each frame in the time series) do not inadvertently reveal the identity of any component. To achieve this, we set the timestamp of the first detection for each sample (i.e., the starting time of the time series) to zero, with all subsequent timestamps measured relative to this first detection.

3. Modeling framework : ConvLSTM2D

In this study, our primary goal is to distinguish between LSNe Ia and other transients (and bogus detections) using time series of 2D images. We aim to make predictions after each observation, with the expectation that the accuracy will improve over time as more data are processed throughout the time series. For this purpose, we employ a ConvLSTM2D, a specialized version of a Recurrent Neural Network (RNN) that extends Long Short-Term Memory (LSTM) networks to handle 2D data with both spatial and temporal correlations. For a more detailed explanation of LSTM and ConvLSTM2D, we refer to Appendix C.

In brief, ConvLSTM2D, just like LSTM, utilizes two types of memory: long-term memory (commonly referred to as the ‘cell state’) and short-term memory (referred to as the ‘hidden state’). The key difference between ConvLSTM2D and LSTM is that ConvLSTM2D applies convolutional operations to both the input images and the hidden states, as opposed to LSTM’s use of matrix multiplications (which is analogous to a fully connected layer). As a result, the two memory states in ConvLSTM2D retain their 2D structure, preserving spatial correlations throughout the process. At each observation epoch, the ConvLSTM2D produces an output based on the input image and the two memory states from the previous epoch. This output is passed forward as the hidden state for the next observation epoch, and the cell state is updated based on the current input and previous memory states.

Similar to convolutional layers, multiple ConvLSTM2D layers can be stacked together to capture spatial and temporal information more efficiently. Figure 4 illustrates our model architecture, which consists of two parallel branches. The first branch processes image cutouts, where the input is passed through three ConvLSTM2D layers with max-pooling and batch-normalization applied in between to capture both spatial features and their temporal evolution. The second branch processes the associated timestamps using an LSTM layer, allowing the model to learn the characteristic timescales of lensed and unlensed SNe Ia.

The outputs from both branches are concatenated and passed through a dense layer, followed by a final sigmoid-activated node that predicts a number, $\mathcal{P} \in (0, 1)$, representing the probability that the sample is a LSN Ia (ideally $\mathcal{P} \rightarrow 1$) or not (ideally $\mathcal{P} \rightarrow 0$). At each observation epoch, the model makes decisions based on the 2D cutout and timestamp for that epoch, as well as the memories (cell and hidden states) from past observations, if any. Performance is expected to improve progressively over time as more data are processed.

We test two types of networks: separate models for each band in single-band analysis, and a single model for multiband analysis that processes multiband time series data.

3.1. Arranging data for multiband analysis

Observations in different bands will generally not be synchronized, meaning that we will not have observations in all bands at the same time⁴. Since LSNe Ia are typically at higher redshifts than their unlensed counterparts, they tend to appear “redder” in optical filters. As a result, the color information available in the multiband scenario can be highly informative for distinguishing LSNe Ia.

The multiband time series of a quadruple LSNe Ia is shown in Figure 5 as an illustrative example. For a balanced analysis

⁴ LSST aims to revisit a sky location in two bands per night, though the specific bands are not fixed. However, to build a generic pipeline, we assume completely asynchronous observations.

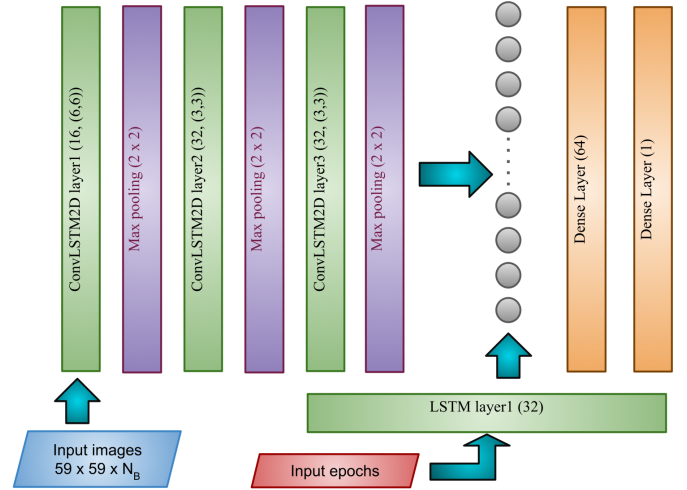


Fig. 4. The figure illustrates the architecture of our deep learning model. At each observation epoch, the image cutout is processed through a ConvLSTM2D channel, while the corresponding time value is fed into a standard LSTM channel. The outputs from both channels are then concatenated and passed through a dense layer. The final layer, consisting of a single node with sigmoid activation, outputs a value $\mathcal{P} \in (0, 1)$, representing the probability that the sample is lensed at that particular observation epoch. We use the binary-cross-entropy loss function for binary classification. Note the number of bands is $N_B = 4$ and 1 for the multiband and single-band analyses, respectively.

and without loss of generality, we ensure that each sample contains exactly 14 observation epochs in total, distributed across the *griz* bands. We denote these observation epochs as $N_{\text{obs},m-b}$, which also represents the cumulative count of multiband observations up to a given point in time; thus, $N_{\text{obs},m-b} \in 1, 2, \dots, 14$, as marked above the top panels. The corresponding timestamps (in days, MJD to be more precise) relative to the first detection are given just below the $N_{\text{obs},m-b}$ labels, enclosed in parenthesis.

At each observation epoch, i.e., in each column of the time series shown in Figure 5, data are available in only one band. The temporal order in which these observations occur in different bands varies from sample to sample. To enable a fair comparison with single-band analyses, we ensure that each sample contains a fixed number of observations per band: specifically, $N_{\text{obs}}^{\text{tot}} = 2, 4, 4, \text{ and } 4$ for the *g*, *r*, *i*, and *z* bands, respectively, resulting in a total of 14 observation epochs per sample. This setup ensures that the same set of samples can be used to compare model performance at every observation step, whether indexed by $N_{\text{obs},m-b}$ for multiband analysis or by $N_{\text{obs},X}$, where $X \in g, r, i, z$, for single-band analyses. Since our models generate predictions after each new observation, this constraint does not limit generality but instead enables a consistent comparison. For our example time series in Figure 5, $N_{\text{obs},X}$ in single band is also marked in each observed frame.

The asynchronous time sampling across bands can be addressed in several ways. As a first step, we adopt a simple, brute-force strategy: missing band images are zero-padded to construct artificial multiband inputs at each observation epoch, enabling the network to exploit color information when available. This uniformly structured input, as shown in Figure 5, is then fed into the multiband model, where we expect the network to learn to ignore the padded entries and focus on the real data. We experimented with different padding values and normalization schemes, and found that padding with zeros, while ensuring the range of real images in the time series remains

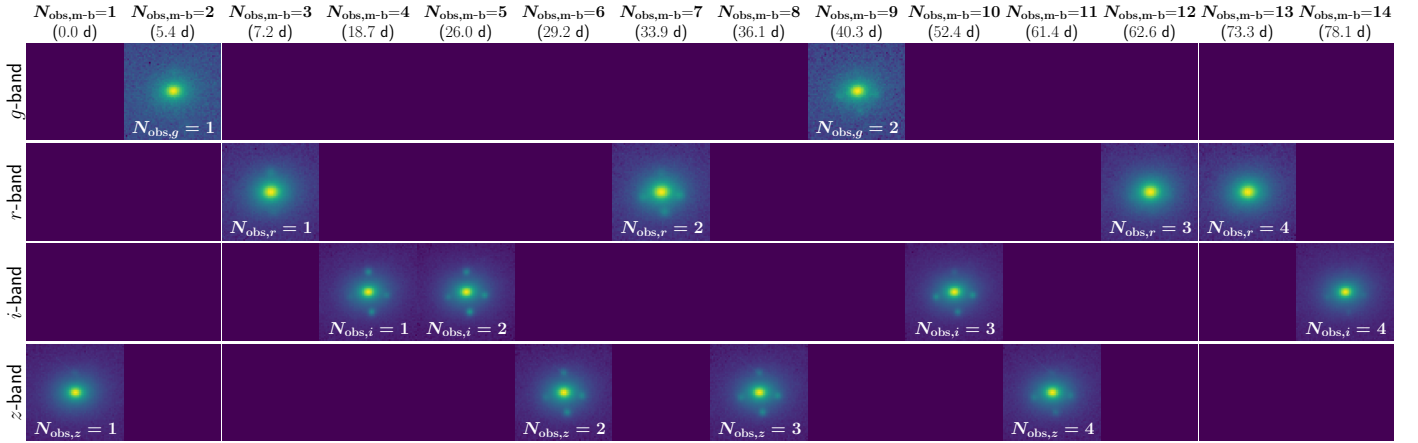


Fig. 5. Example multiband time series of a quadruply LSN Ia, the same system whose i -band time series is shown in the top-left panel of Figure 3. At each observation epoch ($N_{\text{obs},m-b}$), data are available in only one band; missing-band frames are zero-padded to maintain a uniform input format. The timestamps corresponding to each $N_{\text{obs},m-b}$ is shown above the top frames in parentheses (in days, relative to the first detection, which is always set to zero). The band-specific cumulative observation count ($N_{\text{obs},X}$, where $X \in g, r, i, z$) is also marked within each observed frame. In the multiband analysis, the full time series is input as a 4-channel sequence to the multiband model, while in the single-band case, the observation sequence in the respective band is provided to the corresponding single-band model without the need of any padding. Note that we follow the HSC Transient Survey cadence here, and the multiband time series for other classes (HSC variables, SNe in galaxies) have statistically similar time sampling.

positive after normalization, led to the fastest convergence. This approach enables the model to effectively distinguish real data from padded placeholder data, resulting in better classification performance with fewer training iterations.

3.2. Normalization strategy

Normalization of the input data becomes non-trivial in realistic application scenarios, particularly since predictions are made at each observation epoch starting from the first, i.e., from $N_{\text{obs},m-b} = 1$ to 14. One possible approach is to avoid normalization altogether, but we found that normalization can enhance model performance. To address this, we normalize based on the first observation, which can be in any band, depending on the sample. After testing several normalization schemes, we found that the following strategy works best.

1. Each image observed in the time series ($\{I_t\}_{t=1}^{14}$, where $I_t \in \mathbb{R}^{59 \times 59}$, a real-valued (59×59) matrix) is first shifted to have non-negative pixel values by subtracting its minimum pixel value. We then take the square root to enhance the dynamic range:

$$I'_t = \sqrt{I_t - \min(I_t)} \in \mathbb{R}_{\geq 0}^{59 \times 59}. \quad (1)$$

2. For normalization, the first image in the sequence is min-max scaled to the $[0, 1]$ range. The same maximum value (from the first image) is then used to scale all subsequent images:

$$\tilde{I}_t = \frac{I'_t}{\max(I'_1)}. \quad (2)$$

After normalization, the resulting image time series spans a range from zero to order-unity, without intersecting the padding values, which are set to zero exactly. This ensures that the model can more easily disregard the padded missing band images while still retaining the color information. Moreover, this approach leverages the benefits of normalization to improve the model's ability to focus on the actual data.

Table 2. Number of samples of different components in the training, validation, and test sets.

Set	Training	Validation	Test
LSNe Ia	40 000	1893	5811
HSC variables	12 000	4081	12 631
SN in LRGs	20 000	1980	6374
SN in spirals	8000	2046	6308
Total	80 000	10 000	31 124

3.3. Composition of training, validation, and test sets

Table 2 summarizes the distribution of samples across different components in the training, validation, and test sets. In the training set, we maintain a balanced representation of positive and negative samples. However, initial tests revealed that normal SNe Ia in LRGs are most frequently confused with LSNe Ia, contributing the highest FPR⁵. As a result, we give these samples more weight in the final training set compared to the other two negative components – HSC variables and SNe in spirals. As a result, the negative samples in the training set are composed of 30% HSC variables, 50% SNe in LRGs, and 20% SNe in spirals. In the validation and test sets, which share the same component ratios, we increase the number of negative samples to evaluate performance at lower FPRs (note that, in reality, negative samples will vastly outnumber positives). For these two sets, the negative samples are approximately distributed as 50% HSC variables, 25% SNe in LRGs, and 25% SNe in spirals.

⁵ The lack of analogous confusion from normal SNe in spiral hosts could stem from not including spiral lenses in training for reasons described in footnote 2. Having not seen examples of SNe lensed by spiral galaxies, the model could be trivially associating a spiral galaxy with a normal SN. Although difficult to test explicitly, the model is expected to detect some SNe Ia lensed by spirals using lensing features such as image configurations, color, and temporal evolution.

4. Results

4.1. Multiband analysis

As mentioned in Section 3.1, without any loss of generality, we constructed each sample to consist of 14 observations in random order across 4 bands. To assess the classification performance and how it evolves with additional observations, Figure 6 shows the Receiver Operating Characteristic (ROC) curves computed after each observation, yielding 14 ROC curves in total. After each observation, i.e., at $N_{\text{obs},m-b} \in 1, 2, \dots, 14$, we collect the predicted scores (\mathcal{P}) for all samples and compute the corresponding ROC curve, i.e., the true-positive rate (TPR) as a function of the false-positive rate (FPR). This sequence of 14 ROC curves illustrates the progressive improvement in classification performance as the model is exposed to more data over time.

The area under the ROC curve (AUC), shown in the legend of Figure 6, serves as a useful metric for evaluating classification performance. Initially, after the first observation (i.e. at $N_{\text{obs},m-b} = 1$), the AUC is approximately 0.885, reflecting sub-optimal performance. However, as additional observations are incorporated, the ROC curve improves rapidly. At $N_{\text{obs},m-b} = 7$, the model achieves over 60% TPR at the expense of FPR of 0.01%, and by 9th observation, the TPR exceeds 70% with the AUC reaching ~ 0.999 . The corresponding median times for these $N_{\text{obs},m-b}$, relative to the peak of the first image, are approximately 17 and 25 days, respectively. These timescales, however, are strongly influenced by our choice to follow the coarse HSC cadence and by the ad hoc prescription we used for setting the time of first detection. We expect that both the timing and the model’s performance could improve under a more optimized cadence and detection strategy. If we relax FPR to 0.1%, TPR reaches $\sim 60\%$ even earlier at $N_{\text{obs},m-b} = 4$. Beyond 9th observation, the ROC curve saturates, suggesting that further observations do not lead to significant improvements in classification performance. Note that while we used diverse classes of negative examples, in a real classification setup the class proportions will inevitably differ, which will affect performance metrics such as ROC curves and their AUC values.

In Figure 7, we show the distribution of predicted scores (\mathcal{P}) for four different components – LSNe Ia, HSC variables, SNe Ia exploding in LRGs, and spiral galaxies – across four panels. Each panel presents the \mathcal{P} -distribution after the 1st, 3rd, 7th, and 14th epochs of observation, illustrating how the score distribution evolves as the model processes more data. Among the negative components, SNe in LRGs most frequently receive high scores in the $0.9 < \mathcal{P} \leq 1.0$ bin, indicating they are most easily confused with the positives (LSNe Ia). Interestingly, as more observations are available, the model becomes slightly overconfident, increasingly assigning extreme scores. For instance, it predicts a higher number of LSNe Ia samples with scores near unity at later epoch of observation, but this also comes at the cost of a few false positives (negatives) receiving similarly extreme high (low) scores.

Next, we explicitly assess the contribution of different negative components to the FPR. In Figure 8, we show the FPR as a function of score threshold (\mathcal{P}_{th}) for HSC variables, SNe Ia exploding in LRGs, and spiral galaxies, shown in blue, orange, and green, respectively. These curves represent the cumulative distributions, i.e., the fraction of each component with scores exceeding a given threshold, $\mathcal{P} > \mathcal{P}_{\text{th}}$. The top x-axis indicates the corresponding TPR values at different score thresholds \mathcal{P}_{th} . The four panels display results after the 1st, 3rd, 7th, and 14th observation epochs. As expected, the FPR for all components decreases as the model sees more data, i.e., with increasing

$N_{\text{obs},m-b}$. However, in all panels, the FPR associated with SNe Ia in LRGs consistently dominates over the other negative components, indicating that they are more frequently confused with the positive class (LSNe Ia) and are thus the primary contributors to the overall FPR.

To better understand the sources of confusion, we investigate two extreme cases: mock lensed systems that receive very low scores ($\mathcal{P} < 0.01$) which turn out to be mostly doubles, and SNe Ia in LRGs that receive very high scores ($\mathcal{P} > 0.99$). We identify two main scenarios in which LSNe Ia can resemble SNe Ia in LRGs:

- In some occasions, when an image of a doubly LSN appears within the bright region of the lens galaxy, its light curve is heavily dominated by the lens light. As a result, the image may become difficult for the model to detect making the event resemble an unlensed SN embedded in a bright host galaxy.
- When the fainter image(s) of a LSN Ia are not sampled near their peaks due to poor cadence, they may fall below the background noise and go undetected in the image sequence. This again makes the system appear more like an unlensed SN Ia.

Both types of confusion are more common among doubles than quads. In principle, color information can still provide a useful cue for distinguishing between the two classes, as LSNe Ia statistically have higher source redshifts (z_s) compared to the unlensed SNe Ia. However, this signal is not always reliable. The first scenario could potentially be mitigated by adopting difference imaging techniques, which better isolate variable sources from static galaxy light. The second issue would be alleviated with a finer cadence, such as that expected from LSST, improving the chances of capturing the fainter images near their peaks.

Additionally the Appendix D compares classification performance between doubles and quads, and examines how image separation affects the recovery accuracy. Quads are generally easier to detect, and while larger separations improve the performance, the gain diminishes beyond $\sim 1''$.

4.2. Multiband vs. single band results

Instead of using multiband data, one can consider training separate models for each individual band, a setup we refer to as the “single-band analysis”. This approach avoids the complications that arise from asynchronous time sampling across different filters in the multiband case. Due to their higher redshifts, LSNe Ia appear redder, making color information in multiband data particularly informative. In this section, we evaluate whether, and to what extent, the multiband scenario offers improved classification performance over the single-band models.

We remind the reader that, for a fair comparison between single- and multiband analyses, each sample is constrained to have exactly $N_{\text{obs}}^{\text{tot}} = 2, 4, 4, \text{ and } 4$ observations in the $g, r, i,$ and z bands, respectively. Also, $N_{\text{obs},X}$ denotes the observation epochs in band X , i.e. the cumulative number of observations up to a given point in the sequence. Since the order of observations across bands varies randomly from sample to sample, a fixed value of $N_{\text{obs},X}$ corresponds to different multiband indices $N_{\text{obs},m-b}$. Nevertheless, all observations in each band are guaranteed to be present somewhere in the sequence. This setup enables a consistent comparison between single- and multiband models by evaluating their performance after the same observations.

To do this systematically, we track the predictions from multiband analyses for each band separately. Consider a generic




Fig. 6. Receiver operating characteristic (ROC) curves for the multiband classification results obtained from 14 observations per sample. The corresponding area under the curve (AUC), which approaches unity for a perfect classification, is indicated in the legend. Note that the ROC curve improves rapidly over the first few observations, suggesting increased model accuracy with early observations, and begins to saturate when $N_{\text{obs},m-b} \gtrsim 9$.




Fig. 7. Distribution of model predicted scores for four different components (LSNe Ia, HSC variables, SNe Ia in LRGs and spiral galaxies) – are plotted in four panels. Each panel shows the distributions obtained after 1st, 3rd, 7th, and 14th epoch of observation ($N_{\text{obs},m-b}$). Among the negative components, SNe in LRGs contributes the most to the highest score bin indicates this is responsible for most FPR as indicated explicitly in Fig. 8.

band $X \in \{r, i, z\}$, which contains four observations per sample. We exclude the g -band from the single-band analysis due to its limited cadence in the HSC Transient Survey, coarse cadence anticipated from LSST, and the lower signal-to-noise expected for LSNe Ia in this band. To construct the ROC curve after the j -th observation (where $j \in \{1, 2, 3, 4\}$) in X -band, we proceed as follows: for the single-band analysis, we collect the X -band model's predictions after the j -th observation across all samples, i.e. at $N_{\text{obs},X} = j$. For the multiband analysis, we extract the multiband model's predictions after it has seen the same j -th observation in X -band for each sample, regardless of where that observation occurs within the full multiband sequence. The key difference in setup is that, unlike single-band models, the multiband model is exposed not only to prior X -band observations but also to all preceding observations across other bands. This allows the multiband model to develop a richer contextual understand-

ing and incorporate color information, potentially enhancing its ability to learn from the data.

Figure 9 compares the classification performance between multiband (solid) and single-band (dashed) analyses using ROC curves for the r , i , and z bands, shown in three separate panels. In each panel, ROC curves corresponding to predictions after different observation epochs in the concerned band are shown in different colors. The plot legends also indicate the AUC values of these ROC curves, listed first for the multiband and then for the single-band analyses at a given epoch of observation.

As expected, the multiband ROC curves consistently outperform their single-band counterparts in all bands. Notably, the advantage of the multiband analysis is more significant at earlier epochs of observation. For instance, in the i -band at a fixed FPR of 0.1%, the multiband TPR reaches $\sim 40\%$ and $\sim 75\%$ for $N_{\text{obs},i} = 1$ and 2, respectively – substantially higher than the




Fig. 8. FPR as a function of score threshold, separately for different negative components, for results obtained after 1st, 3rd, 7th, and 14th epochs of observation. TPR values corresponding to different score thresholds are marked on the top x -axis. The FPR for all three negative components – HSC variables, SN exploding in LRGs and spiral galaxies – gradually decreases with additional observations. However, it is clearly evident that SNe exploding in LRGs dominate the FPR budget, as they are more likely to be confused with the positive class.

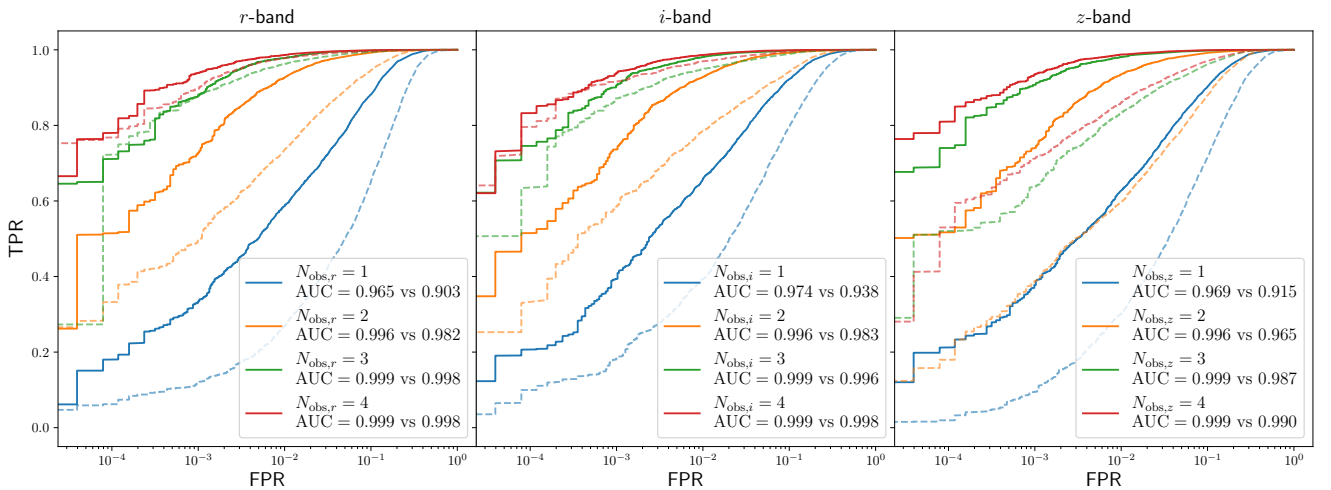


Fig. 9. Comparison of classification performance between multiband (solid) and single-band (dashed) analyses using ROC curves for the r , i , and z bands, respectively shown in three panels from the left to right. In each panel, ROC curves corresponding to different observation epochs ($N_{\text{obs},X}$ where $X \in \{r, i, z\}$) are shown in different colors. Legends indicate the AUC values for each curve, with multiband listed first, followed by single-band for the same epoch.

corresponding TPRs of $\sim 20\%$ and $\sim 60\%$ from the single-band analysis. This is intuitive, as in the earlier part of the sequence, the multiband model benefits more from the accumulated information from all previous observations across all bands and color information. At later epochs of observation, however, the multiband model’s ROC curves from different bands tend to converge in performance, as reflected in similar AUC values. This suggests that the classifier becomes increasingly confident with more observations, making the specific band less critical to the final decision.

One may also be tempted to compare the performance across different single-band analyses. However, we caution that the results may be affected by the cadence distributions, which differ across bands because we matched each band’s cadence to that of the selected HSC variables. It can be seen that, after the first observation, the ROC curve performs best in the i -band, followed by the z - and r -bands, as reflected in their respective AUCs. This is expected, given that LSNe Ia typically exhibit the highest signal-to-noise ratios in the i and z bands. Interestingly, the single-band ROC curves for the r -band show the most significant improvement at later epochs of observation. This may be

partly attributed to the more homogeneous cadence in the r -band compared to the others.

5. Conclusions and discussion

We aim to develop a pipeline for detecting LSNe from other transients and bogus detections using multiband imaging data obtained over multiple epochs, as expected from upcoming time-domain surveys like the LSST. In other words, we construct time-series of 2D image cutouts in multiple filters from these surveys’ alert streams, with the goal of identifying LSNe as early as possible to enable timely follow-up observations. To achieve this, we employ the ConvLSTM2D model, a specialized recurrent neural network (RNN) designed to understand both spatial and temporal correlations in data, making it well-suited to handle time-series imaging data.

Given that we have to rely on simulations to train our model, we aim for as much realism as possible, ensuring that the simulations closely mirror the complexities and challenges encountered in real observational scenarios. To do so, we base our simulations on real observations, particularly from the HSC, as the LSST is expected to match HSC’s data quality in terms of

depths and filter throughputs. Specifically, we use LRGs from the HSC wide-layer as lenses and inject LSN images at various time points corresponding to different phases of SN evolution. Although we focus on SNe Ia in this work, the analysis can be extended to other types of SNe or transients. Negative samples are drawn from variable sources observed in the HSC Transient Survey (HSC variables) in the COSMOS field across multiple epochs. Additionally, we simulate time series of normal SNe Ia exploding in HSC LRGs and spiral galaxies to further enhance the training on these particular cases.

While multiband data provide critical color information, temporally asynchronous observations across different bands presents a challenge in building a seamless network. As an initial approach, we pad the missing-band images to create a consistent input shape, effectively generating synthetic placeholders where data are absent. This method has already shown promising results, as outlined below.

- The ROC curves for the multiband analysis show a rapid improvement in classification performance as the model processes the first few observations across different bands. Keeping the FPR fixed at 0.01%, the TPR reaches $\geq 60\%$ by the 7th observation ($N_{\text{obs,m-b}} = 7$) and exceeds $\geq 70\%$ by the 9th observation ($N_{\text{obs,m-b}} = 9$)⁶. If the FPR is relaxed to 0.1%, the TPR reaches $\sim 60\%$ as early as $N_{\text{obs,m-b}} = 4$. All these demonstrate the model’s strong potential to identify LSNe Ia early in their evolution, enabling timely follow-up observations. Beyond $N_{\text{obs,m-b}} = 9$, the ROC curve begins to plateau, indicating that the model has already gathered sufficient information to make confident predictions, and additional observations offer diminishing returns.
- Among the various negative classes, normal (unlensed) SNe Ia occurring in LRGs contribute the most to the overall FPR, as LSNe Ia can sometimes mimic these events under certain conditions, making them more prone to misclassification compared to other negative classes. For example, when one of the images from a doubly LSN falls within the bright central region of the lens galaxy, or when the dimmer counterimages are poorly sampled and missed around their peak brightness due to cadence limitations, the resulting time series may resemble that of an unlensed SN Ia in an LRG. In such cases, even though color information can help, it is not always sufficient or reliable for distinguishing between the two.
- As expected, quad systems are generally easier to detect than doubles. Likewise, systems with smaller image separations ($\theta_E < 1.0''$) tend to be more challenging to identify compared to those with larger separations ($\theta_E > 1.0''$).
- To avoid dealing with asynchronous observations across different bands, one can instead train separate models for each band and apply them individually to the corresponding single-band data, a strategy we refer to as single-band analysis. While this approach performs reasonably well, it is consistently outperformed by the full multiband analysis. The performance gap is most pronounced at observation epochs, where the multiband model benefits from a richer memory formed by preceding observations across multiple bands. This not only helps establish a stronger temporal context but also leverages valuable color information early on.

⁶ The corresponding median times for these $N_{\text{obs,m-b}}$, measured from the peak of the first image, are approximately 17 and 25 days. These timescales are shaped by our use of the coarse HSC cadence and an ad hoc definition of the first detection. Both the timing and model performance could improve with a more optimized cadence and earlier detection. Notably, if the first detection had occurred earlier in the SN light curve, these timescales would have shifted accordingly.

In contrast, the single-band model’s memory is limited to the past observations in that specific band, lacking both broader temporal context and color cues that are crucial for distinguishing LSNe Ia.

- Our model tends to become overconfident in certain scenarios. As it processes more observations, it increasingly assigns very high scores to more positive samples; however, this comes at the cost of misclassifying a few negative examples, primarily normal SNe Ia occurring in LRGs, with similarly high confidence.

While our model is trained and tested on HSC-like data here, the method can be applied to any cadenced imaging survey. It is especially suitable for the upcoming LSST, where a cadence 5–10 times better than HSC is expected to further enhance detection performance. The promising results highlight the potential of our method for real-time detection of LSNe from the LSST alert stream.

Below, we outline the assumptions and caveats of our current method to maintain clarity and avoid confusion, along with potential improvements and remaining tasks.

- In the first step, this work focuses exclusively on Type Ia SNe, both lensed and unlensed (normal), as their light curves are well understood, and hence easy to implement. However, this method can easily be extended to other types of SNe and transients by simply incorporating their respective light curves. We have not yet considered microlensing, but we expect it to have minimal impact on the performance of our technique, as it primarily affects the brightness of the LSN images, not their positions. Additionally, chromatic microlensing occurs later in the SN Ia evolution, whereas our goal is to detect lensed cases early in their evolution. Therefore, not accounting for microlensing is not a critical issue at this stage. Nevertheless, we plan to include microlensing in a more comprehensive pipeline in the future.
- We have not included host galaxies in the mock LSNe Ia images. In many cases, the host galaxies will also be lensed and visible. However, the choice of SN–host offset distribution and host detectability threshold play a critical role in determining when lensed hosts appear in the data and how prominently. If not treated carefully, these factors can introduce biases into the model. To avoid such premature assumptions, we defer the inclusion of host galaxies to future work, where we will address these aspects systematically. Importantly, our current simulations still reflect a substantial subset of realistic scenarios (roughly half of LSST LSNe Ia) in which the host galaxy is either not strongly lensed or too faint to be detected in the surveys. Since we intend to apply this tool directly to LSST transient alert cutouts, incorporating host galaxies is important as their arc-like features will enhance the performance of our method. Additionally, although we generate synthetic time series for small-separation systems ($\theta_E < 0.5''$), we exclude them here to focus on cases with more easily distinguishable spatial features. These challenging systems will be considered in future work, in conjunction with the inclusion of lensed host galaxies.
- One limitation of the current setup is the simplified treatment of observational effects in the synthetic time series. In particular, we do not model seeing variations or Poisson noise over time—both of which are naturally present in real observations, such as those of HSC variables. While time-varying PSFs are expected to have minimal impact on our method, accurately modeling them is nontrivial and therefore deferred to future work. Similarly, adding Poisson noise to coadded layers is technically challenging, and is also left

for future development. In our current setup, we add white noise at each time step, which is of the same order as Poisson noise near the bright central regions and dominates in the outskirts. Given this, we expect that the absence of explicit Poisson noise modeling is unlikely to significantly affect the results. Nonetheless, incorporating more realistic observational conditions remains an important step for improving the robustness of our approach.

- We use real galaxies and HSC variables to create our datasets, treating HSC as a stand-in for LSST. While this helps us address key conceptual challenges in advance, the upcoming LSST early data will offer a valuable opportunity to build a more representative training set using real LRGs and transients.
- In this work, we have so far matched the cadence of the HSC Transient Survey, which, while useful, is not fully representative of LSST’s cadence. With LSST’s superior cadence, we expect to achieve further improved results, especially enabling earlier detection of LSNe Ia.
- Above, we demonstrated the inclusion of multiband data using a brute-force approach. However, a more sophisticated alternative would involve a composite architecture, where separate ConvLSTM branches process each band in parallel. After each observation, the model could make a decision based on the aggregated outputs from all ConvLSTM branches up to that point. Furthermore, entirely different architectures, such as transformers, may be inherently better suited to handling asynchronous multiband data, and we are exploring this possibility as well.

Acknowledgements. We thank Christian Vogl, Tim Meinhardt and Laura Leal-Taixe for many useful discussions. SB acknowledges the funding provided by the Alexander von Humboldt Foundation. This work is partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1258 – 283604770. SHS thanks the Max Planck Society for support through the Max Planck Fellowship. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (LENSNOVA: grant agreement No 771776). This work is supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2094 – 390783311. SS has received funding from the European Union’s Horizon 2022 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101105167 – FASTIDIoUS.

References

- Abolfathi, B., Aguado, D. S., Aguilar, G., et al. 2018, *ApJS*, 235, 42
- Aghanim, N., Aghanim, N., Akrami, Y., et al. 2020, *A&A*, 641, A6
- Aihara, H., Arimoto, N., Armstrong, R., et al. 2018, *PASJ*, 70, S4
- Aihara, H., AlSayyad, Y., Ando, M., et al. 2019, *PASJ*, 71, 114
- Arendse, N., Dhawan, S., Sagués Carracedo, A., et al. 2024, *MNRAS*, 531, 3509
- Bag, S., Kim, A. G., Linder, E. V., & Shafieloo, A. 2021, *ApJ*, 910, 65
- Bag, S., Huber, S., Suyu, S. H., et al. 2024, *A&A*, 691, A100
- Barbary, K., Barclay, T., Biswas, R., et al. 2016, *Astrophysics Source Code Library* [record ascl:1611.017]
- Bautista, J. E., Vargas-Magaña, M., Dawson, K. S., et al. 2018, *ApJ*, 863, 110
- Bertin, E., Mellier, Y., Radovich, M., et al. 2002, *ASP Conf. Ser.*, 281, 228
- Birrer, S., Shajib, A. J., Galan, A., et al. 2020, *A&A*, 643, A165
- Birrer, S., & Treu, T. 2021, *A&A*, 649, A61
- Bonvin, V., Courbin, F., Suyu, S. H., et al. 2017, *MNRAS*, 465, 4914
- Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, *PASJ*, 70, S5
- Cañameras, R., Schuldt, S., Shu, Y., et al. 2024, *A&A*, 692, A72
- Cañameras, R., Schuldt, S., Suyu, S. H., et al. 2020, *A&A*, 644, A163
- Chao, D. C. Y., Chan, J. H. H., Suyu, S. H., et al. 2020, *A&A*, 640, A88
- Chao, D. C. Y., Chan, J. H. H., Suyu, S. H., et al. 2021, *A&A*, 655, A114
- Chen, W., Kelly, P. L., Oguri, M., et al. 2022, *Nature*, 611, 256
- Craig, P., O’Connor, K., Chakrabarti, S., et al. 2024, *MNRAS*, 534, 1077
- Dalal, N., & Kochanek, C. S. 2002, *ApJ*, 572, 25
- Denissenya, M., & Linder, E. V. 2022, *MNRAS*, 515, 977
- Denissenya, M., Bag, S., Kim, A. G., Linder, E. V., & Shafieloo, A. 2022, *MNRAS*, 511, 1210
- Di Valentino, E., Mena, O., Pan, S., et al. 2021, *Class. Quant. Grav.*, 38, 153001
- Euclid Collaboration (Mellier, Y., et al.) 2025, *A&A*, 697, A1
- Frye, B., Pascale, M., Cohen, S., et al. 2023, *Transient Name Server AstroNote*, 96, 1
- Furusawa, H., Kosugi, G., Akiyama, M., et al. 2008, *ApJS*, 176, 1
- Goldstein, D. A., Nugent, P. E., & Goobar, A. 2019, *ApJS*, 243, 6
- Goobar, A., Amanullah, R., Kulkarni, S. R., et al. 2017, *Science*, 356, 291
- Goobar, A., Johansson, J., Schulze, S., et al. 2023, *Nat. Astron.*, 7, 1098
- Grillo, C., Pagano, L., Rosati, P., & Suyu, S. H. 2024, *A&A*, 684, L23
- Hsiao, E. Y., Conley, A., Howell, D. A., et al. 2007, *ApJ*, 663, 1187
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
- Jacobs, C., Collett, T., Glazebrook, K., et al. 2019, *ApJS*, 243, 17
- Lee, I., Komatsu, E., Suyu, S. H., & Huterer, D. 2016, *JCAP*, 2016, 031
- Jiménez-Vicente, J., & Mediavilla, E. 2019, *ApJ*, 885, 75
- Kelly, P. L., Rodney, S. A., Treu, T., et al. 2015, *Science*, 347, 1123
- Kodi Ramanah, D., Arendse, N., & Wojtak, R. 2022, *MNRAS*, 512, 5404
- Linder, E. V. 2011, *Phys. Rev. D*, 84, 123529
- Lokken, M., Gagliano, A., Narayan, G., et al. 2023, *MNRAS*, 520, 2887
- LSST Science Collaboration (Abell, P. A., et al.) 2009, *ArXiv e-prints* [arXiv:0912.0201]
- Mao, S.-D., & Schneider, P. 1998, *MNRAS*, 295, 587
- Maoz, D., Mannucci, F., & Nelemans, G. 2014, *ARA&A*, 52, 107
- Metcalf, R. B., & Madau, P. 2001, *ApJ*, 563, 9
- Metcalf, R. B., Meneghetti, M., Aveztruz, C., et al. 2019, *A&A*, 625, A119
- Millon, M., Courbin, F., Bonvin, V., et al. 2020a, *A&A*, 640, A105
- Millon, M., Galan, A., Courbin, F., et al. 2020b, *A&A*, 639, A101
- Morgan, R., Nord, B., Bechtol, K., et al. 2022, *ApJ*, 927, 109
- Morgan, R., Nord, B., Bechtol, K., et al. 2023, *ApJ*, 943, 19
- Narayan, R., & Bartelmann, M. 1996, *ArXiv e-prints* [arXiv:astro-ph/9606001]
- Naz, F., She, L., Sinan, M., & Shao, J. 2024, *Sensors*, 24
- Oguri, M. 2007, *ApJ*, 660, 1
- Oguri, M. 2018, *MNRAS*, 480, 3842
- Oguri, M., & Marshall, P. J. 2010, *MNRAS*, 405, 2579
- Oguri, M., Rusu, C. E., & Falco, E. E. 2014, *MNRAS*, 439, 2494
- Petrillo, C. E., Tortora, C., Chatterjee, S., et al. 2017, *MNRAS*, 472, 1129
- Pierel, J. D. R., Newman, A. B., Dhawan, S., et al. 2024, *ApJ*, 967, L37
- Planck Collaboration XIII 2016, *A&A*, 594, A13
- Pooley, D., Rappaport, S., Blackburne, J., et al. 2009, *ApJ*, 697, 1892
- Quimby, R. M., Oguri, M., More, A., et al. 2014, *Science*, 344, 396
- Refsdal, S. 1964, *MNRAS*, 128, 307
- Refsdal, S., & Bondi, H. 1964, *MNRAS*, 128, 295
- Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, *ApJ*, 934, L7
- Rodney, S. A., Brammer, G. B., Pierel, J. D. R., et al. 2021, *Nat. Astron.*, 5, 1118
- Sagués Carracedo, A., Goobar, A., Mörtzell, E., et al. 2024, *A&A*, submitted [arXiv:2406.00052]
- Saha, P., Coles, J., Macciò, A. V., & Williams, L. L. R. 2006, *ApJ*, 650, L17
- Sainz Murieta, A., Collett, T. E., Magee, M. R., et al. 2023, *MNRAS*, 526, 4296
- Sainz Murieta, A., Collett, T. E., Magee, M. R., et al. 2024, *MNRAS*, 535, 2523
- Schuldt, S., Suyu, S. H., Meinhardt, T., et al. 2021, *A&A*, 646, A126
- Schuldt, S., Cañameras, R., Andika, I. T., et al. 2025, *A&A*, 693, A291
- Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, 172, 1
- Shajib, A. J., Vernardos, G., Collett, T. E., et al. 2024, *Space Sci. Rev.*, 220, 87
- Shi, X., Chen, Z., Wang, H., et al. 2015, *ArXiv e-prints* [arXiv:1506.04214]
- Shu, Y., Bolton, A. S., Mao, S., et al. 2018, *ApJ*, 864, 91
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, *ArXiv e-prints* [arXiv:1503.03757]
- Suyu, S. H., & Halkola, A. 2010, *A&A*, 524, A94
- Suyu, S. H., Hensel, S. W., McKean, J. P., et al. 2012a, *ApJ*, 750, 10
- Suyu, S. H., Treu, T., Blandford, R. D., et al. 2012b, *ArXiv e-prints* [arXiv:1202.4459]
- Suyu, S. H., Huber, S., Cañameras, R., et al. 2020, *A&A*, 644, A162
- Suyu, S. H., Goobar, A., Collett, T., More, A., & Vernardos, G. 2024, *Space Sci. Rev.*, 220, 13
- TDCOSMO Collaboration (Birrer, S., et al.) 2025, *A&A*, 704, A63
- Treu, T. 2010, *ARA&A*, 48, 87
- Treu, T., & Marshall, P. J. 2016, *A&ARv*, 24, 11
- Treu, T., Suyu, S. H., & Marshall, P. J. 2022, *Astron. Astrophys. Rev.*, 30, 8
- Tripp, R. 1998, *A&A*, 331, 815
- van Dokkum, P. G. 2001, *PASP*, 113, 1420
- Vegetti, S., Birrer, S., Despali, G., et al. 2024, *Space Sci. Rev.*, 220, 58
- Verde, L., Schöneberg, N., & Gil-Marín, H. 2024, *ARA&A*, 62, 287
- Wojtak, R., Hjorth, J., & Gall, C. 2019, *MNRAS*, 487, 3342
- Wong, K. C., Suyu, S. H., Chen, G. C. F., et al. 2020, *MNRAS*, 498, 1420
- Yasuda, N., Tanaka, M., Tominaga, N., et al. 2019, *PASJ*, 71, 74

Appendix A: Simulating mock LSNe time series

A.1. Preparing LRG catalog

Galaxy-scale LSNe Ia are simulated by injecting point-like LSNe onto cutouts of real HSC galaxy images and embedding them into empty regions of the HSC Transient Survey. As population of foreground deflectors, we use samples of luminous red galaxies (LRGs) which have the highest lensing cross-section. Given the dearth of massive galaxies present in the COSMOS field, the sample of LRGs is obtained from a larger footprint, namely the Wide layer of the HSC-SSP second public data release (PDR2, Aihara et al. 2019) that covers nearly 800 deg^2 with depth of 26.6, 26.2, 26.2, and 25.3 mag in *griz* filters. We selected LRGs with spectroscopic redshift and velocity dispersion measurements available from DR14 of the SDSS-IV survey (Abolfathi et al. 2018; Bautista et al. 2018), in order to get robust proxy of the total galaxy masses for computing realistic light deflection angles. An upper limit on the velocity dispersion uncertainty, $\delta v_{\text{disp}} < 100 \text{ km s}^{-1}$, is set at this stage.

To focus on clean LRG cutouts without artifacts, we used the HSC tables to discard galaxies and their neighbor within $1''$ that were flagged by the pipeline. We request that the following flags are set to False in all *g*, *r*, *i*, and *z*-band: `cmodel_flux_flag`, `pixelflags_edge`, `pixelflags_interpolatedcenter`, `pixelflags_saturatedcenter`, `pixelflags_crcenter`, `pixelflags_bad`, `sdsscentroid_flag`. The cutouts passing these standardised HSC image quality flags are further analyzed to exclude remaining artifacts. We identify cutouts with repeated pixel values indicating partial coverage in PDR2, and cutouts with few-pixel-wide artifacts detected with LACosmic due to their sharp edges (van Dokkum 2001). This results in a sample of 36138 unique LRGs. Their cutouts were imported with the Data Archive System service⁷ together with the relevant PSF frames⁸ in order to produce the LSN simulations. However, to ensure sufficient sample size, we augment the LRG cutouts with up to four rotations.

A.2. The simulation pipeline

Here, we detail the key steps for simulating time series data for mock LSNe. Our focus is on type Ia SN, which are standardizable candles with well-constrained intrinsic luminosities. However, this pipeline can be easily extended to other types of SNe or transient sources.

We base the redshift distributions of both lenses and sources on the Oguri & Marshall (2010, hereafter OM10) simulation (Oguri & Marshall 2010; Oguri 2018). We selected our final LRG sample such that its redshift distribution broadly matches that of the lens population for LSNe Ia in OM10. Rather than adopting the realistic Einstein radius distribution from OM10, we use a flat distribution between $0.1''$ and $2.0''$ to ensure a more balanced training set. With these choices, we construct a matched lens-source catalog by assigning each lens a source redshift, at which the SN Ia will be injected.

We assume a Singular Isothermal Ellipsoid (SIE) + shear model to represent lens galaxies and create the mock lensed images using a pipeline (Cañameras et al. 2020; Schuldt et al. 2021) based on the GLEE software (Suyu & Halkola 2010; Suyu et al. 2012a). The complete pipeline is schematically illustrated in Figure A.1, and its key components are summarized below.

⁷ https://hsc-release.mtk.nao.ac.jp/das_cutout/pdr2/

⁸ <https://hsc-release.mtk.nao.ac.jp/psf/pdr2/>

- For each LRG, we have its spectroscopic redshift, velocity dispersion, and the PSF of its field. We first compute the SIE parameters. Using the source redshift given in the lens-source matched catalog, we obtain the Einstein radius (θ_E) under a Λ CDM background cosmology. A point source is then randomly placed on the source plane within $1.5 \times \theta_E$ to maximize the likelihood of strong lensing. At each random position, we assume a random shear value. Together with the SIE parameters, we compute the lensing configuration and derive the image positions (θ), magnifications (μ), and time delays (Δt) using the GLEE package.

– Lens configuration acceptance criteria:

- The major criterion is whether the faintest image is detectable against the background noise. We validate this in the *i*-band, where the brightness of most LSNe Ia is expected to peak. The background noise level (σ_B) for each LRG is estimated from the standard deviation of pixels in the outskirts.
- We focus on LSNe Ia, which are standardizable candles with well-constrained intrinsic luminosities. For simplicity, we assume a fixed peak absolute magnitude in the *B*-band, $M_B = -19.3$.⁹
- Placing the SN Ia at the source redshift, we compute the peak brightness of all images in the *i*-band, accounting for the PSF associated with the LRG and using the Hsiao template (Hsiao et al. 2007) from the SNCosmo package Barbary et al. (2016). We check whether the peak flux of the faintest image ($F_{0,\text{faint}}$) satisfies the detectability condition $F_{0,\text{faint}} > 5\sigma_B$. If the condition is not met, we calculate the necessary boost in absolute magnitude to achieve a 5σ detection as

$$\Delta M_{\text{faint}} \equiv 2.5 \times \log_{10} \left(\frac{F_{0,\text{faint}}}{5\sigma_B} \right) \Big|_{i\text{-band}}. \quad (\text{A.1})$$

Note that, only when $\Delta M_{\text{faint}} < 0$ we need a magnitude boost in the source, since $\Delta M_{\text{faint}} > 0$ implies the faintest image is already brighter than five times the background noise. We then impose a threshold, $\Delta M_{\text{th}} \leq 0$, and accept the source position only if it satisfies $\Delta M_{\text{faint}} > \Delta M_{\text{th}}$. Otherwise, we try another random source position.

- Condition on the time delays: Between the arrival of the first and second images, the frames may resemble samples of normal SNe exploding in LRGs, where the only differentiating factor could be the color information. The model may confuse these two cases, as color information alone may not be sufficient to distinguish them. Therefore, we require that the second image does not arrive too late, ensuring it can be accommodated within the time series of reasonable length. Thus, we imposed an upper limit on the time delay between the first and second images.
- Lastly, if the SN is placed too close to the caustic curves due to random positioning, the resulting magnification can become extremely high ($\mathcal{O}(10^3)$ or even greater). This can cause one of the LSN images to become excessively bright, potentially dominating the entire mock frames and creating issues during training. Such extreme magnifications are not observed in real lensed systems.

⁹ In reality, M_B varies and can be standardized using Tripp's relation (Tripp 1998). However, for computational efficiency, we assume a fixed value while allowing for a potential boost in M_B if required for detectability. The effect of allowing Tripp's relation-based variations is expected to be minimal on our exercise.

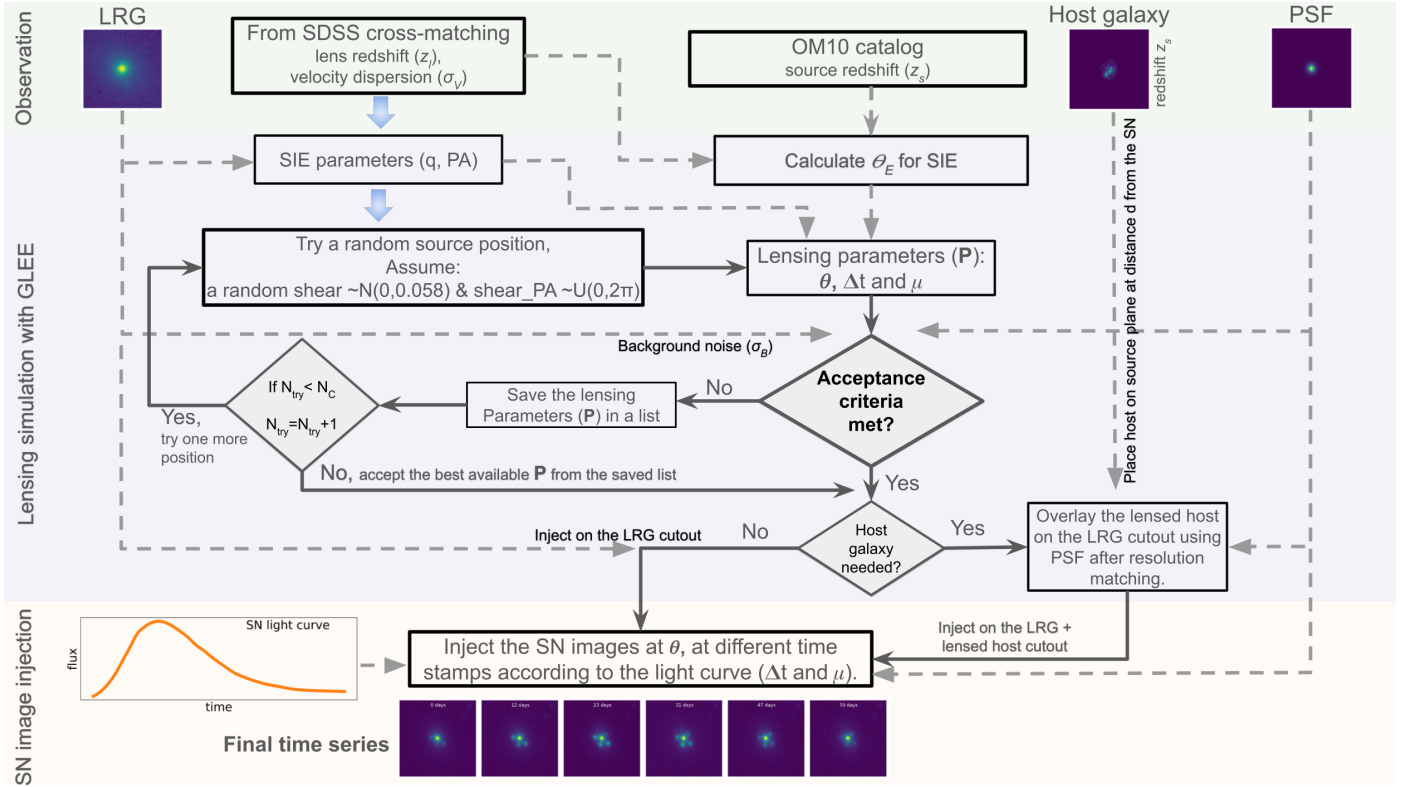


Fig. A.1. Schematic illustration of the simulation pipeline used to create time series of LSNe Ia. Starting from HSC Wide layer observations, LRG samples are cross-matched with SDSS, and lensed SN images are injected onto the LRG cutouts at different epochs, with or without the host galaxy as needed. While the host-inclusion feature is demonstrated for completeness, host galaxies are not included in the simulations used in this work.

To prevent unrealistic scenarios with extreme magnifications, we imposed a limit on the maximum magnification, requiring $|\mu|_{\max} < \mu_{\text{th}}$, where μ_{th} is set to 100.

- In summary, we accept a random position of the SN Ia on the source plane if it satisfies the above three conditions that are recast more generically below,

$$\Delta M_{\text{faint}} > \Delta M_{\text{th}}, \text{ and } \Delta t_{12} < \Delta t_{\text{th}},$$

$$\text{and } |\mu|_{\max} < \mu_{\text{th}} = 100, \quad (\text{A.2})$$

where $\Delta M_{\text{th}} \leq 0$ by construction.

- If the lensing parameters derived from a trial random source position fail to satisfy the acceptance criteria outlined above, we attempt another random source position with new random values for the shear parameters. During this process, we record the source position and corresponding lensing parameters for each trail. We continue this procedure until the number of trials reaches a maximum threshold, N_C , which we set to 400 to balance computational efficiency with scientific objectives. If no acceptable solution is found by then, we systematically relax the magnitude boost (ΔM_{faint}) and time-delay (Δt_{th}) thresholds to identify the most suitable lensing parameters from the pool of recorded attempts, as long as the resulting configurations remain physically plausible. In rare occasions when no realistic set of lensing parameters could be found within the finite (N_C) number of trails we discard that lens-source pair. This ensures that we obtain a reasonable lensing system for almost all lens-source pairs.

We note that in this initial work, we only consider systems with $\Delta t_{\text{th}} = 50$ days and $\Delta M_{\text{th}} = 0$, meaning only cases that require no magnitude boost are included. Notably, the major-

ity of lens-source pairs (approximately 60%) yield acceptable lensing configurations without any magnitude boost and with $\Delta t_{12} < 50$ days. This subset forms the primary focus of our study. Nevertheless, we present the full pipeline here, which can accommodate fainter sources requiring a magnitude boost and allow larger time delays if necessary, attempting to minimize the required magnitude boost and large Δt_{12} whenever possible.

- Once the lensing parameters (θ , Δt , μ) are determined for each lens-source pair based on the *i*-band data alone, we use these parameters to generate SN Ia images as point sources across all bands. Specifically, we place the multiple SN images at the corresponding positions on the LRG cutouts using the respective PSFs, applying the appropriate magnitude boost (ΔM) if necessary. The time delays (Δt) and magnification factors (μ) are incorporated consistently when placing the images in all bands. To build the multiband time series, we inject SN images at various epochs representing different phases of the SN evolution, modeling the SN light curves using the Hsiao template from SNCosmo. Additionally, a random positional shift of up to $0.85''$ (5 pixels) is applied in any direction for each sample (or time series) to mimic observational variations.

The final cutouts have a size of $\sim(10'' \times 10'')$, corresponding to (59×59) pixels at a given time step across all bands. Consequently, the resulting time series have a shape of $(N_t \times 59 \times 59 \times 4)$ where N_t is the number of time steps and 4 is the number of bands.

The pipeline described above enables the creation of multiband time series of mock LSNe Ia at arbitrary numbers of timestamps (N_t) with freely chosen values. We emphasize that,

in our training, we generate asynchronous time sampling across multiple bands, i.e., observations occur in only one band at each given time, using this pipeline.

Some examples of time series of LSNe Ia are shown in the top row of Figure 3 and Figure A.2 for i -band and, in Figure 5 for multiband.

Appendix B: Cadence matching

In Section 2.6, we explain the importance of matching the cadence of the synthetic time series (for LSNe Ia and normal SNe Ia) to that of the HSC variable time series to ensure consistent time sampling. Random sampling from the available observation epochs of the HSC variables, as shown in Figure 1, often results in very coarse cadences and consequently long, sparsely sampled light curves. Such cadences are typically not well-suited for the efficient detection of LSNe. To mitigate this, we define a set of selection criteria to enforce more suitable cadences for the HSC variables. As mentioned earlier, we fix the number of observations per band across all samples: specifically, $N_{\text{obs}}^{\text{tot}} = 2, 4, 4,$ and 4 for the $g, r, i,$ and z bands, respectively. The g -band is excluded when defining our cadence selection criteria due to its very sparse observations in the HSC Transient Survey. Consequently, the criteria are based on the $r, i,$ and z bands, each having four observations.

- For each HSC variable, and for each of the bands, we randomly draw a combination of $N_{\text{obs}}^{\text{tot}}$ observing epochs and check whether the following two conditions are satisfied.
 1. The gap between any two successive observations in a given band must not exceed a predefined threshold: $\delta t \leq \delta t_{\text{th}}$, chosen here as 30 days for riz bands.
 2. The total duration of the light curve in that band must remain below a specified limit: $T_{\text{tot}} \leq T_{\text{th}}$, chosen here as 80 days for riz bands and 50 days for the g -band. This prevents the light curves from becoming excessively long, which would otherwise reduce their relevance for SN-like variability timescales and hinder early detection.
- For each HSC variable and each band, we store all combinations of observing epochs that satisfy the above criteria. If no such valid combination is found in any of the bands, we discard that HSC variable from further consideration.
- While constructing the training set, we randomly select a “good” combination of observing epochs from the saved cadence database for each band for every HSC variable.
- Similarly, while simulating the time series for SNe exploding in galaxies (either LRGs or spirals), we again randomly select a “good” combination of observing epochs from our saved database for each band.
- We apply the same procedure for the LSNe Ia, but with one additional criterion: by the last observation in any band (except for g), the peak of the second arriving image must be included. This ensures that even for long time delay lensed systems (most of which are doubles), the model sees the second image somewhere in the time series. For each band, we randomly select a “good” combination of observing epochs from our saved cadence database until this additional condition is met.

The resulting cadence distribution for the different components in our training and test sets is shown in Figure B.1 for each band. The cadence distributions for HSC variables and SNe exploding in galaxies match exactly in all bands due to the formulation. However, because of the additional criterion for the LSNe Ia, their cadence distribution slightly favors slightly larger gaps between successive observations (i.e. somewhat coarse

cadence). Despite this, the difference is minimal, and the distributions across all components generally match well, as intended.

Appendix C: Details of ConvLSTM2D

Although the details of Long Short-Term Memory (LSTM) and Convolutional LSTM (ConvLSTM) are widely available in the literature, we provide an overview in this section to offer a more comprehensive and complete understanding. We begin with LSTM, which was designed to address the limitations of simple *Recurrent Neural Networks* (RNNs), specifically the vanishing and/or exploding gradient problem. Particularly, LSTM introduces the concept of short and long term memories.

C.1. An LSTM layer

An LSTM layer comprises three gating mechanisms that collectively carry out four core operations at each time step, as outlined below:

$$\mathbf{f}_t = \sigma(\mathbf{W}_x^f \cdot \mathbf{x}_t + \mathbf{W}_h^f \cdot \mathbf{h}_{t-1} + \mathbf{b}^f), \quad (\text{C.1})$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_x^i \cdot \mathbf{x}_t + \mathbf{W}_h^i \cdot \mathbf{h}_{t-1} + \mathbf{b}^i), \quad (\text{C.2})$$

$$\tilde{\mathbf{C}}_t = \tanh(\mathbf{W}_x^c \cdot \mathbf{x}_t + \mathbf{W}_h^c \cdot \mathbf{h}_{t-1} + \mathbf{b}^c), \quad (\text{C.3})$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_x^o \cdot \mathbf{x}_t + \mathbf{W}_h^o \cdot \mathbf{h}_{t-1} + \mathbf{b}^o). \quad (\text{C.4})$$

Here, σ and \tanh represent the sigmoid and hyperbolic tangent activation functions, respectively. \mathbf{x}_t is the input data at time step t , and \mathbf{h}_{t-1} is the hidden state (short-term memory) output from the previous time step. The symbol \cdot denotes the inner product. The superscripts ‘f’, ‘i’, ‘C’ and ‘o’ correspond to the forget gate, input gate, cell state, and output gate, respectively. Meanwhile, $\tilde{\mathbf{C}}_t$ denotes the candidate cell state, which represents the new information proposed to update the cell state. Next, the cell state and hidden state are updated as follows:

$$\mathbf{C}_t = \mathbf{f}_t \odot \mathbf{C}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{C}}_t, \quad (\text{C.5})$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{C}_t), \quad (\text{C.6})$$

where \odot indicates element-wise (Hadamard) multiplication. The hidden state \mathbf{h}_t serves as the output of the LSTM layer at the current time step t .

Analogous to how a dense layer processes fixed feature sets without modeling internal correlations, an LSTM layer efficiently models sequences of feature vectors, such as time series of temperature, humidity, and air pressure, where the input at each time step, \mathbf{x}_t , is one-dimensional and contains no spatial structure or internal correlations.

Similar to multiple nodes in a dense layer, an LSTM layer can have multiple units, denoted by N_h (which can be invoked by $\text{LSTM}(N_h)$ in TensorFlow), producing an output \mathbf{h}_t of shape $(N_h, 1)$ assuming the input \mathbf{x}_t is one-dimensional. This is illustrated in the left panel of Figure C.1, where, for clarity, we simplify the sum of the two matrix multiplications appearing in Equations (C.1)–(C.4) as

$$\mathbf{W}_x \cdot \mathbf{x}_t + \mathbf{W}_h \cdot \mathbf{h}_{t-1} = [\mathbf{W}_x \mathbf{W}_h] \cdot \begin{bmatrix} \mathbf{x}_t \\ \mathbf{h}_{t-1} \end{bmatrix}, \quad (\text{C.7})$$

given that the number of rows of \mathbf{W}_x and \mathbf{W}_h are the same (N_h) and \mathbf{x}_t and \mathbf{h}_{t-1} , being 1D, have matching numbers of columns.

All terms on the left-hand side of equations (C.1)–(C.6) share this same shape, $(N_h, 1)$. Therefore, N_h can be considered the




Fig. A.2. Examples of time series of mock LSNe Ia systems in *i*-band. The three rows correspond to systems with smaller (top), intermediate (middle), and larger (bottom) Einstein radii, with the specific values indicated at the first time step of each series. Time stamps, relative to the first observation, are displayed at the top of each frame. The time sampling follows the HSC Transient Survey cadence, as discussed in Section 2.6 and detailed further in Appendix B. In each row, two systems are presented: a quad on the left and a double on the right. For visual clarity, we show in this figure systems with prominent lensing features; however, many samples in the training set exhibit more subtle and less easily identifiable signals. Still, some examples shown here include foreground or background contaminants, reflecting the observational complexities incorporated in the training data.




Fig. B.1. The cadence distributions for different components across four bands are shown in four panels. HSC variables are selected with “good” cadences according to the criteria described in Appendix B. We then match the cadence distributions of our simulated time series to those of the selected observed HSC variables, separately for each band. Due to the additional requirement of capturing at least the second trailing image for the LSNe Ia, their cadence distribution tends to slightly favor coarse cadence, i.e. larger gaps between successive observations.

dimensionality of the LSTM layer. For example, consider a one-dimensional input vector of length N_x . The weight matrices for the input, all four \mathbf{W}_x , then have the shape (N_h, N_x) , while the weight matrices for the hidden state, all four \mathbf{W}_h , have the shape (N_h, N_h) . Each bias term will have the shape N_h . Consequently, the total number of parameters in the LSTM layer is given by:

$$N_{\text{PLSTM}} = 4 \times ((N_h N_x + N_h^2 + N_h)) = 4 \times N_h (N_x + N_h + 1).$$

C.2. A ConvLSTM2D layer

In Shi et al. (2015), the LSTM architecture is extended to capture spatial correlations by replacing inner products (\cdot) with convolutional operations ($*$). This transformation, which mirrors the shift from dense to convolutional layers, leads to the ConvLSTM model. The gating mechanisms remain unchanged, so equations (C.1)-(C.6) still hold, with the inner product (\cdot)

replaced by convolution ($*$), as shown by the right panel of Figure C.1 which is adapted from Naz et al. (2024). When we are interested in 2D convolutions, typically needed to capture spatial correlations in 2D data, this is called ConvLSTM2D. As a result, both the memory and output at each time step retain a 2D structure, analogous to that produced by a standard convolutional layer.

Suppose the input sample data have the shape (N_x, N_y, N_c) (N_c = no. of channels in the input data), and the convolution kernel has the spatial dimension (K_x, K_y) . With ConvLSTM2D(N_k , kernel_size= (K_x, K_y)) the following properties hold:

- The shape of all input kernels (\mathbf{W}_x) = (K_x, K_y, N_c) .
- The shape of all kernels for the hidden state (\mathbf{W}_h) = (K_x, K_y, N_k) .
- The shape of the hidden state (\mathbf{h}) = (N_x, N_y, N_k) when using the ‘same’ padding, and $(N_x - K_x + 1, N_y - K_y + 1, N_k)$ for

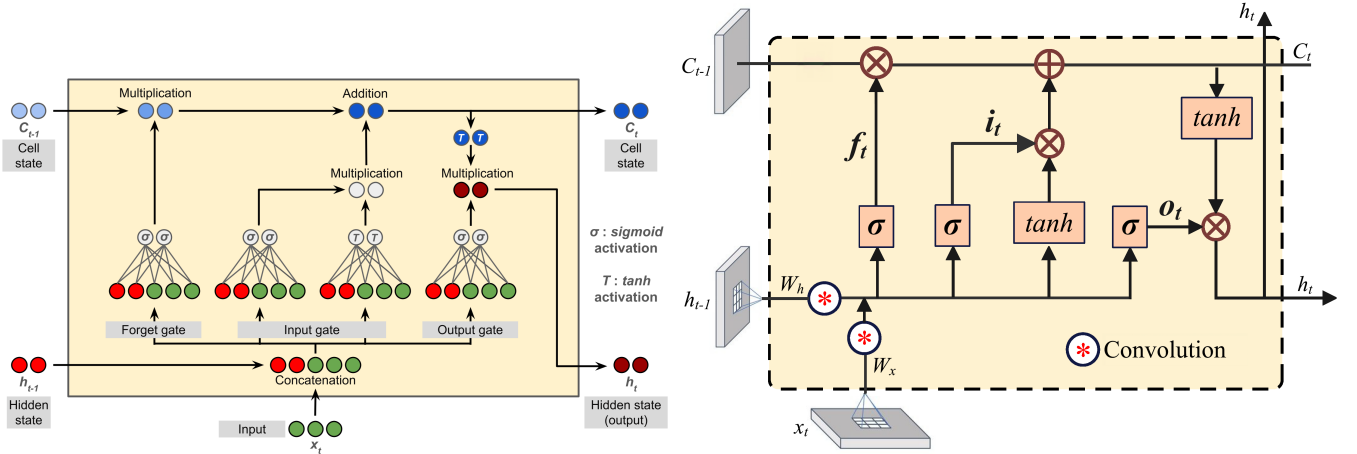


Fig. C.1. The left panel demonstrates how a two-unit LSTM layer processes a sequence with three features at each time step. For better visual illustration, we represent the sum of two matrix multiplications, in Equations (C.1)–(C.4), as a single matrix multiplication between a concatenated weight matrix of shape (2, 5) and a stacked vector of \mathbf{x}_t and \mathbf{h}_{t-1} of shape (5, 1), as described in Equation (C.7). The output h_t has a length of 2. The right panel, adapted from Naz et al. (2024), illustrates the generalization to ConvLSTM2D, where matrix multiplication (similar to fully connected operations) is replaced by convolution operations (denoted by *), while handling both the short-term memory and input data, all of which are 2D at each time step. While the gating mechanism remains the same between an LSTM and a ConvLSTM cell, the two panels depict them differently to enhance clarity.

the ‘valid’ padding.¹⁰ This will be the shape of the output (at each time step), the cell state etc, basically all the LHS in the Equations (C.1)–(C.6) above.

- Considering four core operations in Equations (C.1)–(C.4), the total number of parameters in a ConvLSTM2D layer is given by:

$$\begin{aligned} N_{P, \text{ConvLSTM2D}} &= 4 \times [N_k K_x K_y N_c + N_k K_x K_y N_k + N_k] \\ &= 4 \times N_k [K_x K_y (N_c + N_k) + 1]. \end{aligned} \quad (\text{C.8})$$

Note that, by construction, when $(K_x, K_y) = (1, 1)$, i.e., in the limit where the 2D filter reduces to a scalar, ConvLSTM2D simplifies to a standard LSTM, where the number of convolutional filters N_k plays the role of the hidden dimension N_h . Another interesting limit arises at the beginning of the sequence, where ConvLSTM behaves similarly to a standard convolution layer. This occurs when the initial short-term memory (hidden state) is zero-initialized, as is typically the default case, while the cell state is set to a constant-valued (non-zero). Therefore, ConvLSTM initially operates like a convolution layer, but as the sequence progresses and the memory mechanism becomes active, it progressively improves its predictions.

Appendix D: Recovery Performance by LSNe Ia Properties

D.1. Doubles versus quads

We first assess the classification performance between doubles and quads among the LSNe Ia samples. Figure D.1 presents the distribution of predicted scores (\mathcal{P}) for all LSNe Ia, as well as separately for doubles and quads, in four panels corresponding to results after the 1st, 3rd, 7th, and 14th epochs of observation. The left panel of Figure D.2 displays the corresponding ROC

¹⁰ Note that for convolution over the hidden state, i.e., for all terms like $\mathbf{W}_h * \mathbf{h}_{t-1}$ in the core operations in Equations C.1–C.4, ‘same’ padding is always used to ensure consistent hidden state dimensions across time steps.

curves – different colors represent the four observation epochs, while solid, dotted, and dashed curves (within each color) denote all mock LSNe Ia, doubles, and quads, respectively.

It is clearly evident from Figure D.1 that quads generally receive higher scores than doubles. This is expected, as quads with four images exhibit more profound spatial features compared to doubles with only two. Additionally, quads typically have shorter time delays between images, causing multiple images to appear closer together in time. Together, these factors result in stronger spatial and temporal correlations in quads, making them easier for the model to classify correctly. The relatively easier recoverability of quads compared to doubles is also evident in the ROC curves shown in the left panel of Figure D.2.

D.2. Separation of multiple SN images

In Figure D.3, we examine how the classification accuracy varies with the image separation of LSNe Ia. To do this, we divide the LSNe Ia samples into three bins based on their Einstein radius: $0.5'' \leq \theta_E < 1.0''$, $1.0'' \leq \theta_E < 1.5''$, and $1.5'' \leq \theta_E < 2.0''$. The four panels of Figure D.3 show the distributions of predicted scores (\mathcal{P}) for these three subsets after the 1st, 3rd, 7th, and 14th observation epochs, respectively. The right panel of Figure D.2 presents the corresponding ROC curves – colors indicate the observation epoch, while solid, dotted, and dashed lines represent the three θ_E bins at each $N_{\text{obs}, m-b}$.

We find that systems with lower separations ($0.5'' \leq \theta_E < 1.0''$) generally receive lower scores compared to systems with larger separations, and are therefore relatively harder to recover as evident from the ROC curves in the right panel of Figure D.2. This is also expected, as smaller separations increase the chance of multiple images blending together due to seeing/PSF effects, or merging with the light from the lens galaxy, which makes them harder to detect and classify confidently.

Interestingly, there is no significant difference in score distributions between the intermediate and large separation bins ($1.0'' \leq \theta_E < 1.5''$ and $1.5'' \leq \theta_E < 2.0''$). A plausible explanation is that, although higher separation generally makes the system easier to resolve, it can also come with challenges: such




Fig. D.1. Distributions of model-predicted scores (\mathcal{P}) for all LSNe Ia, as well as separately for doubles and quads, shown across four panels corresponding to the results after the 1st, 3rd, 7th, and 14th observation epochs. It is evident that quads generally receive higher scores than doubles, indicating that they are, on average, easier to recover.

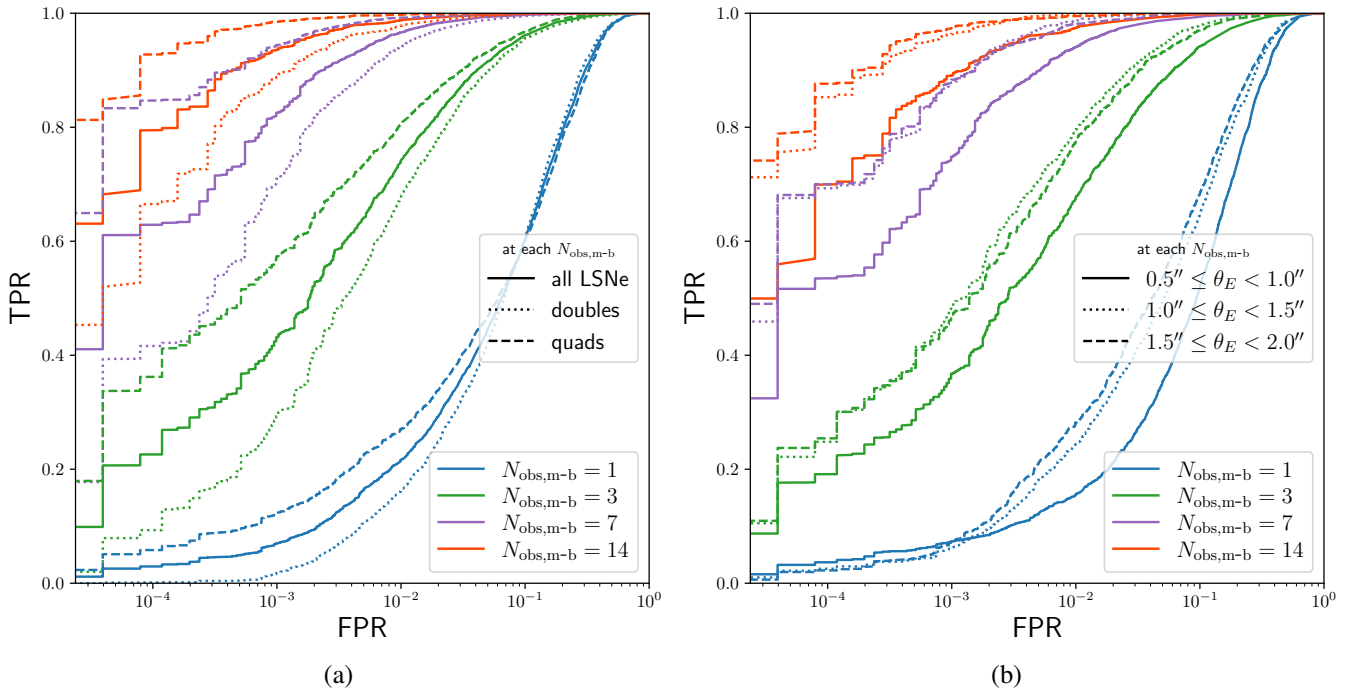


Fig. D.2. Comparison of classification performance across different sub-classes of LSNe Ia systems using ROC curves. The left panel compares doubles and quads, while the right panel shows results for systems grouped by angular separation, highlighting the impact of image separation. Results at the 1st, 3rd, 7th, and 14th observation epochs are shown in blue, green, purple, and orange, respectively, in both panels. In the left panel, solid, dotted, and dashed curves of a given color represent all LSNe Ia, doubles, and quads, respectively; in the right panel, the same line styles denote the three θ_E bins at the corresponding observation epoch.

systems often correspond to higher source redshifts, leading to fainter images, and very large separations can sometimes place images far enough from the lens to be affected by nearby bright objects in the field. These competing effects may balance out, resulting in comparable classification performance for the two higher-separation bins as evident in Figure D.3 and the right panel of Figure D.2.

D.3. Time-delay distributions of high-score mock lenses

The score threshold (\mathcal{P}_{th}) corresponding to a given FPR varies with the observation epoch ($N_{\text{obs},m-b}$) as the scores of individual samples evolve. In this initial work, we do not fix a detection threshold, since doing so would require careful calibration

of the false alarm rate, which depends on the true proportions of positive and negative samples. Our training set will also be refined in future work. At this stage, our goal is to demonstrate the method and explore how model architecture and data pre-processing affect the performance, rather than to benchmark it under fully realistic conditions. Nevertheless, we examine the time-delay distribution of mock lensed systems receiving high scores in this section, focusing on subsets defined by thresholds $\mathcal{P}_{\text{th}} = 0.75$ and 0.99 .

The left panel of Figure D.4 shows the cumulative distribution of the time delay between first and second arriving images, Δt_{12} . The filled cyan distribution represents the full test sample of 5811 mock lensed systems. The blue, green, purple, and red curves correspond to the distributions after the 1st, 3rd, 7th,




Fig. D.3. To investigate the impact of image separation on classification, we divide the LSNe Ia systems into three bins based on Einstein radius: $0.5'' \leq \theta_E < 1.0''$, $1.0'' \leq \theta_E < 1.5''$, and $1.5'' \leq \theta_E < 2.0''$. The four panels show the distributions of predicted scores (\mathcal{P}) for these three subsets after the 1st, 3rd, 7th, and 14th observation epochs, respectively.

and 14th observations, respectively. For these colored curves, the solid and dashed line-styles indicate systems with $\mathcal{P} > 0.75$ and $\mathcal{P} > 0.99$, highlighting subsets of high-score mock lensed systems.

As expected, systems with smaller Δt_{12} tend to receive high scores earlier, that is, after fewer observations (i.e. at lower $N_{\text{obs},m-b}$). The arrival of the second image generally strengthens the overall signal, both spatially and temporally, while also making these systems more distinct from normal SNe Ia in LRGs. In addition, the gap between the dashed and solid curves narrows as more epochs are observed, indicating that the model becomes increasingly confident with additional data, consistent with Figure 7.

On the other hand, the right panel of Figure D.4 shows the distribution of Δt_{12} for systems with scores below these thresholds, i.e., $\mathcal{P} < 0.75$ and $\mathcal{P} < 0.99$, using solid and dashed line styles. Purple and red curves correspond to results after the 7th and 14th observations, highlighting the distributions at later stages of the time series. The distributions are normalized, with the y-axis showing the fraction (in percentage) of systems in each Δt_{12} bin receiving scores below the thresholds. As expected, systems with larger Δt_{12} initially tend to receive lower scores until the second image is observed, as indicated by their higher representation in the larger Δt_{12} bins.

By construction of our time series, the second image always appears by the final observation ($N_{\text{obs},m-b} = 14$), see Appendix A.2. As a result, the higher Δt_{12} bins no longer show an excess at $N_{\text{obs},m-b} = 14$. In fact, systems with larger Δt_{12} often have wider angular separations and more distinct lensing features, making them easier to recognize once the second image is observed.

A potential complication arises for high- Δt_{12} systems when the second image, being widely separated from the first, is missed during its bright phase due to sparse time sampling, as in our HSC-based simulations. In such cases, while many systems still achieve high scores, some may be misidentified. This issue is expected to be far less significant for LSST, which will provide a much denser cadence than HSC.




Fig. D.4. Distribution of the time delay between the first and second arriving images (Δt_{12}) for mock lensed systems. The left panel shows the cumulative distribution, i.e., the number of systems with Δt_{12} smaller than a threshold Δt_{th} , plotted along the x-axis. The filled cyan histogram represents the full test sample of 5,811 systems. Blue, green, purple, and red curves correspond to the distributions after the 1st, 3rd, 7th, and 14th observations, respectively. For these curves, solid and dashed line-styles indicate the subsets with scores $\mathcal{P} > 0.75$ and $\mathcal{P} > 0.99$. The right panel presents the Δt_{12} distributions for lensed systems with scores below those thresholds, i.e., with $\mathcal{P} < 0.75$ and $\mathcal{P} < 0.99$, using solid and dashed curves, respectively. The y-axis gives the fraction of such systems in each Δt_{12} bin as a percentage. Purple and red curves again correspond to results after the 7th and 14th observations.