PEAKY FINDERS: CHARACTERIZING DOUBLE-PEAKED TYPE IIB SUPERNOVAE IN LARGE-SCALE LIVE-STREAM PHOTOMETRIC SURVEYS

Adrian Crawford [1], Tyler A. Pritchard [2,3,4] and Maryam Modjaz [1]

[1] Department of Astronomy, University of Virginia, Charlottesville, VA, 22903, USA
[2] Department of Astronomy, University of Maryland, College Park, MD 20742
[3] Astrophysical Sciences Division, NASA/Goddard Space Flight Center, 8800 Greenbelt Rd, Greenbelt, MD 20771, USA
[4] Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771

1. ABSTRACT

We present the first photometric population study of double-peaked Type IIb supernovae (SNe IIb). SNe IIb are produced from the core-collapse of massive stars whose outermost hydrogen layer has been partially stripped prior to explosion. These double-peaked light curves, consisting of a shock-cooling emission peak (SCE) followed by the main nickel-powered peak, contain more crucial information about the progenitor system than the typical single-peaked light curves. We analyzed a sample of 14 spectroscopically confirmed SNe IIb with publicly available observations, discovered between 2018–2022, from the ZTF and ATLAS surveys. We developed and fit a linear model, referred to as the "lightning bolt model," to describe the early-time behavior and measure their population statistics. We find the SCE peak lasts, on average, > 5 days above half-maximum light. These SCE rise rates are over 10x faster than the rise to the nickel-powered peak and are comparable to those of fast blue optical transient (FBOT) events. Finally, we present a proof-of-concept alert filter to demonstrate how to translate these population statistics into simple and effective filters to find potential double-peaked SNe IIb in large-scale survey alert streams, like the imminent Vera C. Rubin Observatory Legacy Survey of Space and Time.

2. INTRODUCTION

Some of the most energetic explosions in the universe occur when a massive star ($M > 8\,M_\odot$) ends its life as a core-collapse supernova (CCSN). Observed CCSNe sub-types depend on their progenitor stellar systems, explosion mechanisms, and surrounding environments. Within these CCSNe, some stars will have been stripped of their outermost hydrogen, and sometimes helium, shells and belong to the class called stripped-envelope supernovae (SESNe; Filippenko 1997; Modjaz et al. 2019). The mechanisms that cause this stripping/mass loss are an outstanding question.

Within the core-collapse and stripped-envelope classes lie the Type IIb supernovae (SNe IIb), which are characterized by their spectra which show initially strong hydrogen lines that eventually fade and give way to strong helium lines—suggesting that the progenitor's outer hydrogen layer was partially stripped prior to explosion (e.g., Woosley et al. (1994); Modjaz et al. (2019)). We have directly observed the progenitors of a few of SNe IIb—e.g., 1993J (Aldering et al. 1994), 2011dh (Arcavi et al. 2011), and 2013df (Van Dyk et al. 2014)—many of which imply a yellow supergiant progenitor. However, there are still many proposed and viable progenitor channels that can create an empirical Type IIb SN, specifically binary systems and interactions (Smith 2014; Dessart et al. 2020). Thus, SNe IIb make for interesting and unique probes of the mass-loss mechanisms that strip stars as well as a probe into the lives and deaths of massive stars.

The main peak in the optical light curve of most, if not all, SESNe is driven by the radioactive decay of ^{56}Ni and the shape of this peak is directly related to the ejecta mass (which itself is indirectly related to the progenitor mass) and ejecta velocity (Arnett 1980), given certain opacity and geometric assumptions. However, some SESNe, including SNe IIb, can show an additional earlier peak. This preceding peak is powered by shock-cooling emission (SCE; e.g., Woosley et al. 1987; Modjaz et al. 2019) which occurs as the stellar envelope cools after it has been heated by the shock wave which exploded the massive star. The SCE peak is a tracer of the progenitor stellar envelope radius and shock wave geometries (Rabinak & Waxman 2011; Nakar & Piro 2014) and can, in some rarer cases, also trace the circumstellar material (CSM) Pellegrino et al. (2023); Jacobson-Galán et al. (2022)).

The coupling of the SCE peak with the progenitor information from the nickel-powered peak makes these double-peaked light curves especially advantageous for stellar forensics. The challenge lies in obtaining photometric observations quickly enough to catch the majority of the SCE peak, particularly in the UV as past work have shown is vital for constraining progenitor models (Pellegrino et al. 2023). It is known that the SCE phase of the light curve evolves much quicker than the nickel-powered peak which evolves on a scale of weeks to months (e.g., Tartaglia et al. (2017)). Ho et al. (2023) showed that amongst the fast-evolving transients they found in the ZTF data, the included SNe IIb population showed a range of time above half-maximum flux, ranging from 3-12 days.

Corresponding author: Adrian Crawford
adrian.crawford@virginia.edu

Prior single-object papers characterizing the shock-cooling emission in individual SNe IIb have used specific, detailed models (each with their own physical assumptions) to describe the object's behavior and possible progenitor channels (e.g. Arcavi et al. (2017)). However, vital light curve parameters, such as the observed rise and decline times as well as duration of SCE, have not been quantified in a statistically large sample; this is the main aim of this work. With a more data-driven characterization of the evolution of the SCE peak, and of double-peaked SNe IIb light curves in general, we can better study these information-rich objects through statistical samples and more easily identify these objects amongst the transient alert streams.

While current surveys like Zwicky Transient Facility (ZTF; Bellm et al. 2019), Asteroid Terrestrial-impact Last Alert System (ATLAS; Tonry et al. 2018), All-Sky Automated Survey for Supernovae (ASAS-SN; Kochanek et al. 2017), and Distance Less Than 40 Mpc (DLT40; Yang et al. 2017) have already revolutionized the field, time domain astronomy is about to undergo another transformation with the commencement of the Vera C. Rubin Legacy Survey of Space and Time (Rubin LSST; Ivezić et al. (2019)). Rubin LSST is currently predicted to have over a million transient alerts each night, with $\sim$ 1000 of those alerts being new supernovae. While brokers will largely handle the brunt of the data influx and management (i.e. by generating alerts) developing quick filters/algorithms/classifiers to find objects of particular interest will a be vital non-broker task. Many of the classifiers being developed focus on early classification so that we may best allocate our scarce follow-up resources, and be able to trigger and coordinate multi-band and multi-wavelength observations across multiple facilities for our most promising and interesting objects. Yet, many make use of simulated data, to obtain a large enough training and test set, meaning that any flaws in the theoretical models or simulated dataset are perpetuated through to the classification schema and machine learning features.

Data-driven methods, which utilize real observed data, represent an alternative approach that is not subject to these model-flaws. In this work we adopt a data-driven approach using 14 spectroscopically confirmed SNe IIb to limit the numbers of physical assumptions we make while ensuring a representative sample of double-peaked SNe IIb (keeping in mind that our sample size of 14 objects).

In this work we present the first ever early-time population statistics of double-peaked SNe IIb light curves. Our dataset is described in §3. We create a data-driven, alert-stream-influenced "lightning bolt" model to characterize the SCE peak and nickel-powered peak rise, described in §4. From the lightning bolt model we generate population statistics that constrain the SCE rise time, duration, magnitude variation, among many other interesting parameters outlined in §4.1. We use these population statistics as features in a proof-of-concept alert stream filter, engineered on the ANTARES broker in §5, and discuss future improvements and avenues in §5.2. Finally, we summarize our methods, findings, and implications of the work in §6.

3. DESCRIPTION OF DATASET

Through a combined literature and Transient Name Server search, we compiled a set of SNe IIb whose data were publicly available survey data, were taken between 2018–2022, consisted of 2 photometric bands, and showed evidence of a double-peaked light curve. This resulted in 14 objects which are listed in Table 1.

Because forced photometry light curves are produced at a specific RA, Dec at every time that there is an observation, these light curves can include additional measurements that may not be included in the alert-stream light curves. For this reason, we choose to utilize forced photometry in our analysis so that we may have the best-sampled, composite-survey light curves possible.

For this work, we query forced photometry data from the ZTF forced photometry server (Masci et al. 2023) and the ATLAS forced photometry server (Smith et al. 2020; Shingles et al. 2021) for each object in our sample. We clean the forced photometry light curves by removing bad observations. For ZTF data we apply the quality cuts recommended in Masci et al. (2023): `infobitssci`$<$ 33554432, `scisigpix`$<=$ 25, and `sciinpseeing`$\leq$ 4. Similarly, we apply a reduced chi-squared fit of the PSF to the ATLAS data by selecting observations with `chi/N`$<$ 4. Note, the ZTF forced photometry data products are in flux units while the ATLAS forced photometry data products are in magnitude units. We chose to convert all light curves to AB magnitude space to better match the current ZTF alert-stream.

4. MODELING THE SCE PEAK AND RISE TO NICKEL-POWERED PEAK

In order to ascertain whether a new object belongs to a certain population, one must first know what the population looks like on the whole. While the main nickel-powered peak has been well studied and described, the preceding shock-cooling emission peak has proven to be more elusive. Previously published SCE studies of SNe IIb are limited to a single/handful of objects at a time and rely on theoretical models to interpret the SCE peak. As a complimentary approach, in this work we aim to provide the first population statistics of double-peaked SNe IIb as a class, specifically focused on characterizing the SCE peak's photometric behavior from publicly available alert stream data. Additionally, we provide the largest sample of SNe IIb with SCE to date.

We aim to use these new population statistics to inform and build a filter that can identify these particular objects from an alert stream. This filter development is discussed more in §5.

Table 1. Description of Dataset

IAU Name	Discovery Name	RA	Dec	z	Paper Reference
–	ZTF18aalrxas	15:49:11.64	+32:17:16.68	0.0582	Fremling et al. (2019)
SN 2018dfi	ZTF18abffyqp	16:50:50.084	+45:23:52.44	0.031302	Bruch et al. (2021, 2023)
SN 2019rwd	ZTF19acctwpz	00:10:45.898	+21:08:20.73	0.017	Bruch et al. (2023)
SN 2020ano	ZTF20aahfqpm	13:06:25.176	+53:28:45.53	0.03113	Ho et al. (2023); Khakpash et al. (2024)
SN 2020ikq	ATLAS20lfu	13:36:05.016	+28:59:00.11	0.037	Ho et al. (2023); Khakpash et al. (2024)
SN 2020rsc	ZTF20aburywx	01:19:56.503	+38:11:09.66	0.0313	Ho et al. (2023); Khakpash et al. (2024)
SN 2020sbw[a]	ZTF20abwzqzo	02:46:03.318	+03:19:47.66	0.023033	Bruch et al. (2023)
SN 2021gno[a]	ZTF21aaqhhfu	12:12:10.290	+13:14:57.05	0.006211	Jacobson-Galán et al. (2022); Ertini et al. (2023)
SN 2021heh	ZTF21aaqvsvw	07:59:47.290	+25:21:20.99	0.026648	Soraisam et al. (2022)
SN 2021pb[a]	ZTF21aabxjqr	09:44:46.80	+51:41:14.6	0.033	Das et al. (2023)
SN 2021vgn	ZTF21abrgbex	16:21:10.510	+36:03:40.36	0.032341	This work
SN 2022hnt	ZTF22aafrjnw	11:36:59.754	+55:09:50.25	0.0192	Farah et al. (2025)
SN 2022jpx	ZTF22aajkpen	10:10:10.000	-11:04:50.05	0.015	This work
SN 2022qzr	ATLAS22zpf	00:09:55.001	-05:01:16.09	0.018705	This work

[a]Typed as Calcium-Rich Transients of Type IIb

Our lightning bolt model was created with three objectives in mind: one, be able to characterize the early time photometric behavior, specifically the SCE peak; two, generate model-agnostic, data-driven statistics; and three, create easily translatable population statics for use in alert-stream filtering. In regards to the third goal, when a new transient is identified in the alert stream, its light curve consists of very few observations—in the earliest cases, just two detections. This means that any filtering done on an alert stream must involve simple cuts, especially for quickly evolving objects like SCE peaks. The most popular/common filtering criteria are slope (i.e. is this object evolving at a rate that is expected?) and duration (i.e. is this a persisting transient or extremely short-lived transient). To this end, when developing the lightning bolt model, we chose to use three piecewise lines to characterize and model the early-time behavior as their slopes are the most readily translatable into filter cuts.

With these criteria in mind, we developed the "lightning bolt" model, as shown in Figure 1, which derives its name from the three-line zig-zag shape characteristic of cartoon lightning bolts. Note that the lightning bolt model is only concerned with the earliest photometric behavior and that we do not use or fit observations taken after the nickel-powered peak.

There are 7 parameters that make up the lightning bolt model: m_1 is the slope of the first rise (to SCE peak); m_2 is the slope of the first decline (from SCE peak); m_3 is the slope of the second rise (to nickel peak); b_2 is the magnitude-axis offset of the model; a_1 is the time of the SCE peak; a_2 is the time of the trough between the SCE and nickel peaks; $\log(f)$ is the estimation of the underestimation of the errors.

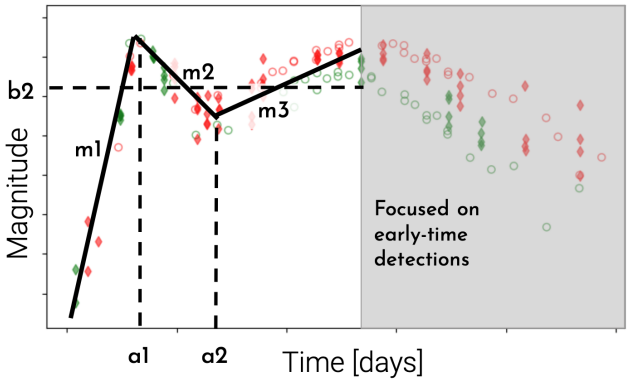

Figure 1. Schematic of the lightning bolt model which is used to fit the early time photometric light curves (stopping before the nickel-powered peak maximum.) The model consists of 7 parameters: m_1 the first rise, m_2 the first decline, m_3 the second rise, a_1 the time of the SCE peak, a_2 the time of trough between the two peaks, b_2 the overall magnitude offset, and $\log(f)$ as estimate of the errors on the errors (not pictured as it is a statistical, not physical, parameter).

The lightning bolt model takes the parametric form below:

$$f(x) = \begin{cases} m_1 x + a_1(m_2 - m_1) + b_2 & x \leq a_1 \\ m_2 x + b_2 & a_1 \leq x \leq a_2 \\ (m_3 x) + a_2(m_2 - m_3) + b_2 & a_2 \leq x \end{cases}$$

However, not all objects have observations comprising the rise to the SCE peak. In these cases, we use a simpler, "two-line" version of the lightning bolt that only fits the decline from the SCE peak and the re-rise to the nickel-powered peak, wher $f(x) = NaN$ for $x < a_1$.

We refer to the original lighting bolt model as the "full" model or the "three-line model" while the latter is exclusively referred to as the "two-line model."

We apply the lightning bolt model to the forced photometry data of the whole 14 object sample from ZTF and ATLAS in the following section §4.1.

4.1. *Fitting to Forced Photometry Light Curves*

As mentioned in §3, we queried and cleaned the ZTF and ATLAS forced photometry light curves for each of the 14 objects. In order to increase the number of observations in each light curve and maintain the two-filter approach, we chose to treat g-band and c-band as one filter and r-band and o-band as another filter. For the remainder of this section, g-band refers to both g and c, and r-band refers to both r and o.

To capture the full rise of the SCE we try to include one "marginal detection" in each band just before SCE when possible. In this work, marginal detections were identified as observations that were significantly dimmer ($\lesssim 4$ mags) than the main SN light curve and had larger errors. We chose to only include marginal detections that fell within ~ 3 days to the first significant detection in a given band (where significant detections are observations that rise significantly above the baseline/background, generally have smaller errors, and are not isolated from other similar points.) In general, our measured slopes on the rise to the SCE peak are lower limits to the possible steepness of the rise.

As mentioned previously, we are only interested in characterizing the early photometric behavior, thus, we exclude any observations that are taken 10 days after a_2, i.e. the trough between the peaks.

We normalized the light curves by shifting and aligning each light curve, in g and r separately, to the trough between the SCE and nickel peak, a.k.a. a_2 in the lightning bolt model, such that $t(a_2) = 15$ days (a fiducial number that allowed nearly all data points to have $t > 0$ days) and $\mathrm{mag}(a_2) = 0$.

After aligning the light curves, we sub-selected the early observations of the light curves, cutting out any observations that fall more than 10 days after the trough. We fit the two-line model to the g-band light curves of SN 2020ano, SN 2021heh, SN 2021pb, SN 2022hnt, SN 2022jpx, and SN 2022qzr. All other g-band light curves and every r-band light curve was fit using the three-line model.

We perform the MCMC fitting using the package `emcee` (Foreman-Mackey et al. 2013), minimizing the log-likelihood, and use the auto-correlation time as a metric for the convergence of the MCMC fits. We found that 128 walkers, $2e6$ iterations, and a 500-step burn-in phase achieved convergence for each of the 14 objects.

Table 2. Best-fit population statistics from lightning bolt model

Parameter	Units	Min	Mean	Max	StdDev	Filter
m_1†	mags/day	-4.1	-2.1	-0.7	1.0	g
m_2	mags/day	0.09	0.26	0.50	0.12	g
m_3	mags/day	-0.22	-0.08	0.01	0.06	g
b_2	mags	-6.6	-3.4	-1.0	1.7	g
a_1†	days	0.1	7.2	11.6	3.3	g
a_2	days	11.3	14.0	17.1	1.5	g
$\log(f)$	...	-4.8	-2.3	0.3	1.3	g
m_1	mags/day	-2.95	-1.10	-0.17	0.76	r
m_2	mags/day	-0.02	0.19	0.68	0.16	r
m_3	mags/day	-0.23	-0.09	-0.02	0.05	r
b_2	mags	-9.5	-2.8	0.2	2.2	r
a_1	days	0.0	7.7	12.3	3.1	r
a_2	days	12.3	14.5	17.5	1.3	r
$\log(f)$	...	-4.6	-2.4	-0.7	1.1	r

† Values calculated using the 8 (out of 14) objects with observations along the rise to the SCE peak.

In Figure 2 we now show the mean of the 14 best-fit models in each band along with the standard error of the mean (SEM). We choose to calculate the mean because of the small-number statistics that we are working with. We calculate the mean model fit by taking the mean of each of the parameters and plugging those values into our lightning bolt model, Eq. 1. The SEM is shown to quantify the reliability of our mean model rather than the spread in our population.

We describe the population statistics from the mean MCMC best-fit parameters in Table 2. By using the lightning bolt model we were able to obtain rise/fall rates for the three early-time regimes of double-peaked SNe IIb lightcurves. Notably, the rise to the SCE peak (m_1 in our model) evolves at a rate of ≈ 2 mags/day for the g-band and ≈ 1 mags/day in the r-band. In the g-band, this rise to the SCE peak evolves an explosive 25x faster than the rise to the nickel-powered peak. For the r-band, the rise to the SCE peak evolves 12x faster than the nickel-powered rise. Even the decline from the SCE peak (which evolves at around an eighth to a sixth the rate of the SCE rise) still evolves $\approx$ 3x faster than the nickel-rise in the g-band and $\approx$ 2x faster in the r-band. Furthermore, the fall from the SCE peak lasts an average of just seven days. This is a much shorter timescale than the weeks to months that the nickel-powered peak evolves over. We explore this quick SCE evolution timescale further in §4.2.

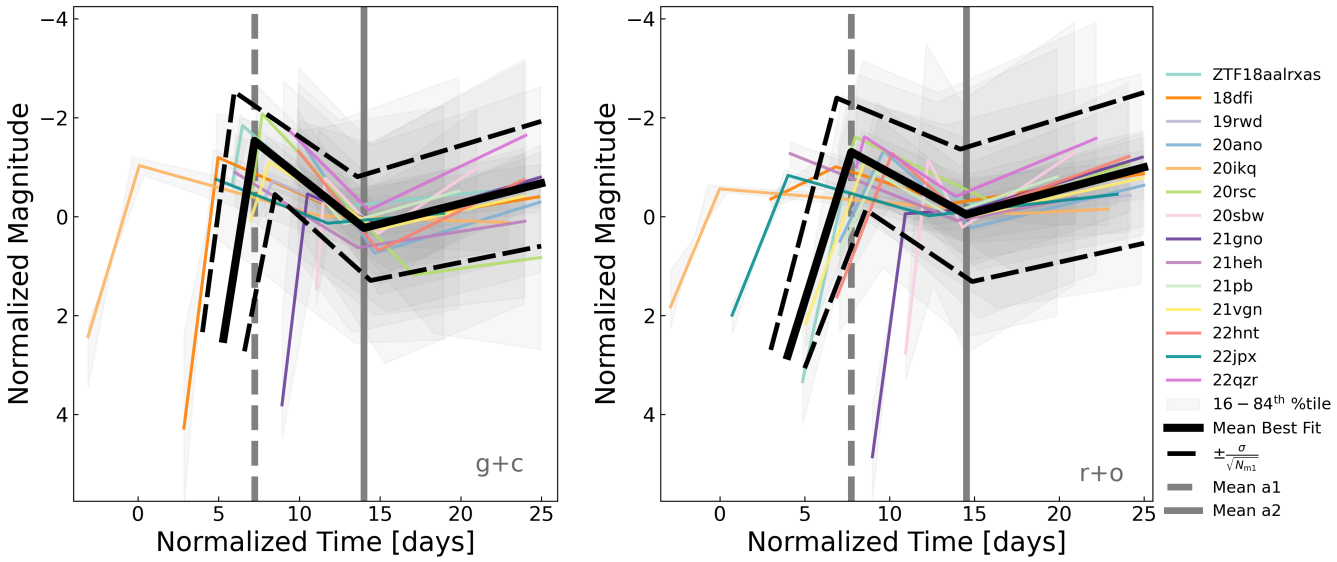

Figure 2. Analysis of the sample's best-fit results using ZTF and ATLAS forced photometry observations. Solid-colored lines are the median fit from the MCMC results with $16-84^{\rm th}$ percentile error shown in the light gray shaded region. The dashed gray line is a_1 and solid gray line is a_2. The solid black line is the mean of the 14 best-fit models and we show the standard error of the mean (SEM) in the dashed black lines. Note that while all forced photometry light curves are normalized by eye such that the trough falls at $t = 15$ days, the trough from the best-fit MCMC results may fall before, at, or after this fiducial time. The predominant source of scatter between the various light curves is the difference in time between a_1 and a_2. We see that the slopes of the rise to the SCE peak (m_1), while occurring at different relative times, are still quite similar in steepness; the same is true for the slope of the rise to the nickel-powered peak (m_3). *Left panel*: g+c bands, *Right panel*: r+o bands.

Table 3. Additional population statistics and features as derived from best-fit lightning bolt model parameters.

Feature[a]	Min	Mean	Max	StdDev	N	Units
$(a_2 - a_1)^g$	3.3	7.0	11.2	2.5	8	days
$(a_2 - a_1)^r$	2.0	6.8	17.6	3.6	14	days
$a_1^g - a_1^r$	-1.86	-0.64	0.10	0.58	8	days
$a_2^g - a_2^r$	-6.3	-0.5	0.8	1.8	14	days
$\mathrm{mag}(a_1)^g$	-2.07	-1.15	-0.45	0.52	8	mags
$\mathrm{mag}(a_1)^r$	-1.62	-1.04	-0.06	0.41	14	mags
$\Delta\mathrm{mag}(a_2 - a_1)^g$	0.63	1.49	3.26	0.68	8	mags
$\Delta\mathrm{mag}(a_2 - a_1)^r$	-0.09	0.96	1.55	0.44	14	mags
$\Delta\mathrm{mag}(a_1^g - a_1^r)$	-0.66	-0.16	0.38	0.29	8	mags
$\Delta\mathrm{mag}(a_2^g - a_2^r)$	0.04	0.37	1.37	0.33	14	mags

[a]Superscripts refer to the band the fitting was performed on. $\Delta\mathrm{mag}(t_1^{\rm band} - t_2^{\rm band})$ refers to the difference in magnitude between the two times.

4.2. *Connection to Fast Blue Optical Transients: Peak Luminosity and Timescale of SCE*

From the results of the previous section, it is clear that the SCE peak evolves on a much shorter timescale than that of the nickel-powered peak. In this section, we further characterize the SCE properties by measuring its absolute magnitude and calculating the time above half-maximum flux ($t_{1/2}$). Further, we compare our results to current FBOT parameter spaces as presented in Ho et al. (2023) and find that all 14 of our objects' SCE peak fall within the typical FBOT timescale of $t_{1/2} < 12$ days.

First, we compute the absolute magnitude of the SCE peak for each object by the following steps. We use a central wavelength of 4830Å for the g-band and 6260Å for the r-band. We use the packages `dustmaps` (Green 2018) and `extinction` (Barbary 2021) to calculate line-of-sight Milky Way extinction for each SN using the Fitzpatrick & Massa (2007) formulation.

To calculate the distance modulus, we first search for published values, which we obtained for ZTF18aalrxas, 21gno, and 21heh. For the remaining object without published distances, we use the redshifts listed in Table 1. We assume cosmological parameters as outlined in WMAP9 (Hinshaw et al. 2013). We use two methods of calculating the absolute magnitude depending on the redshift of the objects. If the object is in the Hubble flow ($z \geq 0.015$ for this work) we use:

$$M = m - 5\log_{10}\left(\frac{D}{10\,\mathrm{pc}}\right) + 2.5\log 10(1+z) \qquad (1)$$

If the object has a redshift $z < 0.015$ we use the following instead, $M = m - 2.5\log_{10}(D/10\,\mathrm{pc})^2$.

The distances, line-of-sight extinction values, and absolute magnitudes of the SCE peak in both bands are available upon request.

In Figure 3, we recreate the Fast Blue Optical Transient (FBOT) parameter space presented in Drout et al. (2014); Ho et al. (2023) of absolute magnitude versus time above half-maximum in order to compare the fast-evolving (but well-known) SCE peak of our objects with the emerging population of fast-evolving transients (with unknown progenitor or powering-sources). The gray points are literature values taken from Ho et al. (2023) whose absolute magnitudes were measured in the g-band. We plot the FBOTs of all spectroscopic subtypes in the Ho23 sample, but highlight the SNe IIb in particular with an additional border around their plotting symbol (square).

We overlay our measurements of the SCE peak for the 8 objects that have data comprising the SCE rise and fall. Our g-band measurements are denoted with the larger colored squares with black edges. For the objects without full SCE data, we instead use the r-band measurements (plotted with as smaller boxes with red outlines). The 16th/84th percentiles are the lower/upper bound errors.

The mean values of our 8 SNe IIb with g-band photometry is over-plotted in black with $\pm1\sigma$ error bars and lies at $\mathrm{Mag}_g = -17.02\pm1.39$ and $t_{1/2} = 4.95\pm1.49$ days. We can draw a number of important conclusions from Figure 3: the mean as well as the individual SN IIb SCE values in our sample are well within the range of the Ho23 FBOTs and certainly within the 12-day classical FBOT cutoff. This means that the well-known SCE phenomenon in SNe IIb can be found in fast/high-cadence transient searches. However, it also means that this well-known SCE phenomenon in SNe IIb can be mistaken as something new if the main second, Ni-driven peak is not observed/followed due to too short of a time baseline in surveys (see Fig. 1 in Ho23 and section 5.8 in Khakpash et al. (2024)). Usually when people refer to FBOTs or rapidly evolving SNe in the literature they refer to truly exotic objects like AT2018cow, which only have one peak and are of high luminosity (see Section 5 in Ho et al. 2023, not the well-known SCE peaks of SNe IIb.

While we know that the early detection of SESNe with SCE is necessary for obtaining progenitor information from the light curve, early detection is also equally important for planning spectroscopic observations. Even when utilizing ground-based Target of Opportunity (ToO) observations, which are typically obtained within days to weeks, they are often too slow to observe SCE, as seen in Fig 3 with the sample mean's incredibly short ~ 5 days above half-maximum flux. The rapid time evolution of the SCE peak limits spectroscopy to observing facilities that can execute spectroscopic observations with minimal turnaround time.

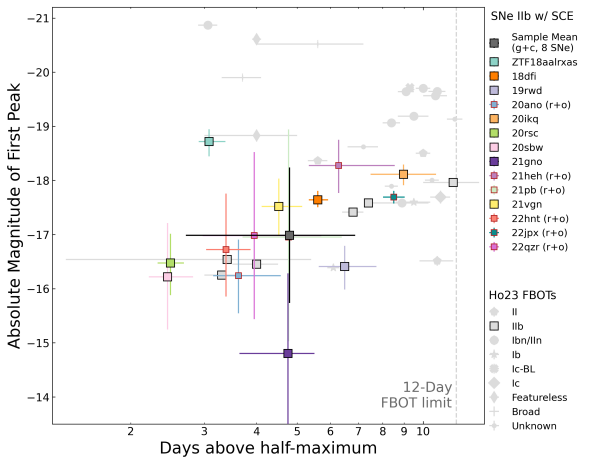

Figure 3. Comparison between the SCE peak absolute magnitude and days above half maximum with literature FBOTs (Ho et al. 2023). Pale gray markers represent objects from Ho et al. (2023). Colored square markers represent objects from this work. The black square represents the mean (of the 8 object sample) for the SCE peak (in the g+c filters). All objects in this work fell below the 12-day FBOT cutoff, with the sample spending an average of 5 days above half-maximum light. Note, values from Ho et al. (2023) are calculated using the g-band. However, not all objects in this work had enough data in the g+c bands to accurately calculate time above half-max. Thus, for the 6 objects missing this information, we use the r-band photometry and they are denoted by the boxes with red outlines in the figure. Also note that 3 objects are shared between this work and Ho et al. (2023)—20ano, 20ikq, and 20rsc—and the differences between our measurements are described in the text.

5. FEATURE ENGINEERING PRELIMINARY ANTARES FILTER

In the upcoming age of the Vera C. Rubin Observatory LSST, there will be an estimated 10 million alerts each night (Ivezić et al. 2019; Graham et al. 2024). There has already been important groundwork laid in the development of transient classifiers, each with their own classification goals as mentioned in the introduction. Most of these classifiers focus on categorizing transients into broad SNe and transient classes, often grouping stripped envelope SNe under a single "Ibc" class. Because of the highly diverse nature and behavior of stripped envelope SNe, the Ibc class tends to have a lower purity score than other transient/SN classes, i.e. objects that are not actually Ibc's are misclassified as such thus diluting the overall class behavior (e.g., Muthukrishna et al. (2019); de Soto et al. (2024)). Part of this low purity can, in some cases, be attributed to the use of flawed

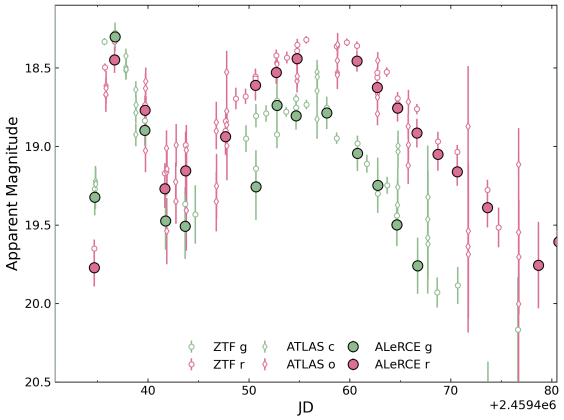

Figure 4. Light curve of SN 2021vgn comparing the coverage between a single alert stream data and multi-survey forced photometry data. The alert stream data is taken from the ALeRCE detections for 21vgn are plotted in the larger circles with black outlines. The forced photometry data is taken from ZTF and ATLAS and are plotted as the smaller empty circles and diamonds respectively. Note that on the rise to the first peak, there are only two detections in each band from the alert data, one close to explosion and one at the SCE peak. In comparison, the combined forced photometry has around 5 observations per band spread from first detection to SCE peak. This light curve sampling disparity will influence how we filter off the alert stream, namely, that broader/more generalized cuts will be prioritized over specific/exact fit measurements. Future time-domain survey designs must prioritize shorter cadences which we approximate here with the multi-survey forced photometry data.

simulated lightcurves, which, when used as a training set for a machine learning algorithm, propagates uncertainty and error into the classification schema itself.

As we saw in the previous section, §4.1, the average double-peaked IIb spends an average of 5 days above half-maximum during the first SCE peak; yet the SCE peak contains important and unique probes of SNe IIb progenitors that are not accessible with only the nickel-powered peak, namely, probes of the radius of the progenitor and the mass/radius of the CSM (Soderberg et al. 2012; Pellegrino et al. 2023). Thus, detailed, high-cadence photometry, specifically during the SCE peak, is necessary for the most detailed and insightful analysis of possible progenitors. On top of the detailed photometry, spectroscopy also holds important information that is not contained within the light curves alone.

Our goal, in this classification instance, is to be able to reliably and efficiently identify potential double-peaked IIb SNe from alert streams, such as ZTF and Rubin LSST. Being able to tag potential targets will allow interested astronomers to quickly call for finer-grain photometric follow-up using lesser-subscribed telescopes. Additionally, in an age of overwhelming photometric detection and sparing spectroscopic resources, being able to efficiently narrow down the deluge of newly discovered objects into a manageable subset of high-significance objects lets us probe the edges of our physical understanding through targeted and efficient spectroscopic follow-up studies.

Earlier in this work we combine forced photometry from two surveys, ZTF and ATLAS, to create an approximation of a high-cadence light curve for a single source. Now, we contrast those forced photometry light curves with the sparser alert-stream light curves, as seen in Figure 4. In this figure, we use the ALeRCE alert stream which runs on the public ZTF survey data and thus has a cadence of 3 days (Förster et al. 2021). Just looking at the alert light curves of this work's sample, we see a wide spread in the difference in time between an object's first few detections. The shortest length of time between the first two detections of an objects were intranight while the longest was around 2-3 days (i.e. on par with public survey cadence). Thus, when creating filters and features that are specifically designed to be applied to infant light curves with few observations, we must keep this spread in fidelity and cadence in mind.

The first step in the creation of this classifier was to measure and describe the shared properties (i.e. features) of this class. We did so by utilizing the best-fit parameters from the MCMC fitting as described in the previous sections (§4.1). We use the mean and standard deviations of the features outlined in Tables 2 and 3 to create a preliminary alert filter on the ANTARES broker (Matheson et al. 2021). We do not include $log(f)$ as a parameter as it is a statistical value borne from our MCMC methodology and is not an inherent physical or photometric feature. We do not include b_2 as a feature as this parameter described the offset of the whole lightning bolt model in normalized-magnitude units and is not a true physical property of the light curves. Finally, we do not include a_1 or a_2 as individual features as this information is more valuably encoded in the Table 3 features. The creation and testing of this "zeroth order"/proof-of-concept filter is further described in the following section.

5.1. *"Low-Tech" Proof-of-Concept ANTARES Filter*

Here the use of "low-tech" refers to the lack of machine learning or AI used in the creation and implementation of this version of the alert filter. For our proof-of-concept testing we choose to use historical alert stream light curves (for a subset of this work's 14 object sample), meaning that the light curves of each object have already been observed and contain only public alert stream detections (and non-detections). Thus, the alert stream light curve data can include partial SCE peak evolution and can extend past the peak nickel-powered peak. Note that this filter is the first step in creating a truly "live-stream" alert filter that reads in new transient detections as they arrive and does not have access

to the full light curve evolution. However, the proof-of-concept filter approach described in this section would easily translate into an archival search tool.

The high level workflow of the filter is to take an object's alert package (called an "alert locus" in the ANTARES vernacular) and pass it through a series of quality checks and feature cuts to "make sure" it is a supernova before tallying up the total number of the features from our target sample that are present in the alert object. The output is a 3-tiered tagging system, based on the number of features met/present, ranging from Gold ("Highly probable double-peaked IIb"), Silver ("Probable double-peaked IIb") to Bronze ("Potential double-peaked IIb"). If not enough features are met the object is simply skipped/ignored.

Taking a closer look at each of these overarching steps in the workflow, in order to increase the likelihood that the light curve we are looking at belongs to a supernova we first pass the alert light curve through a series of three checks that ask if alert object is coincident with moving solar system object [if yes, skip object, if there are at least 30 minutes between the first and latest alert detection [if no, skip object], if the alert object is coincident with known stellar object [if yes, skip object].

These particular cuts were informed by prior work in detecting fast-evolving targets in transient surveys, namely Andreoni et al. (2021); Ho et al. (2020).

If an object's alert light curve passes these initial cuts, we then move onto calculating the 22 possible features describing our sample of double-peaked SNe IIb (as derived from Tables 2 and 3). First we must identify what shape the light curve has in order to determine the phase and compare to the correct model parameters. We split the alert stream light curve observations into g- and r-band and perform the next steps on each band independently.

We split the light curve into rise/fall/re-rise chunks and fit a line to each phase of the light curve. If our fitted slopes fall within $\bar{x} + 2\sigma$ of the measured population statistics from Table 2 then we tag that feature. After calculating the individual slopes for both bands and assessing if they meet the slope features, we then assess whether the object meets any of the time/magnitude/filter offset features from Table 3.

After all the feature criteria have been assessed, we add up the number of features present in that object's alert stream light curve. There are 22 total possible features across the two bands. If fewer than 10 features are present, the object is most likely NOT a potential target and is skipped. If 10-11 features are present, we tag the object as "bronze" meaning it is potentially a double-peaked IIb SNe. If 12-13 features, we tag the object as "silver" meaning this is a probable target object. Finally, if 14 or more features are present we tag the object as "gold" meaning there is a strong likelihood that the alert stream light curve belongs to a double-peaked IIb SNe. These cutoffs were determined by optimizing the recovery rate of the alert light curves that had the fullest evolution/coverage.

We tested the performance of this filter on a set of 10 target objects (i.e. known double-peaked IIb SNe identified in this work: 18dfi, 19rwd, 20sbw, 21gno, 21heh, 21pb, 21vgn, 22hnt, 22jpx, 22qzr) and 33 non-target SNe objects (e.g., SNe II, IIP, Ib, Ic, Ia). We do not include the other 4 objects presented as a part this work in our target sample as their alert light curves did not contain the decline from the SCE peak in at least one band.

Of the 10 target objects, 8 were correctly identified as double-peaked SNe IIb (4 Gold, 1 Silver, and 3 Bronze) while 2 did not meet the required number of features and so were not tagged. Of the 33 non-target objects, 28 were correctly not tagged while 5 were incorrectly tagged as potential double-peaked SNe IIb (1 Silver and 4 Bronze). For the target objects that were (incorrectly) discarded as non-target objects, their alert light curves contained a partial decline from the SCE peak only in a single band. Thus, the total number of features identified, out of the 22 possible sample-derived features, did not meet the minimum of 10 features needed to be tagged.

In summary, implementing our proof-of-concept filter on historical, non-live-stream light curves taken from the ZTF/ANTARES alert stream resulted in 8 true positives, 28 true negatives, 2 false negatives, and 5 false positives.

The overall accuracy of the filter was 0.837. The completeness was 0.800 and the purity was 0.615. From the initial accuracy and completeness, we see that the filter is adept at identifying both true positives (identifying double-peaked SNe IIb as such) and true negatives (ignoring non-double-peaked SNe IIb).

5.2. *Discussion of Proof-of-concept Alert Stream Filter*

A potentially powerful improvement to make to the alert stream/survey filter would be the implementation of simultaneous fits across photometric bands—making use of all photometric information at once—rather than splitting the light curves into separate bands. This would allow single-band filtering (currently skipped in this alert filter) and increase flagging information in both fine- and coarse-grain light curves. This combination of multiple filters' information will be especially necessary once the Rubin LSST transient search begins as the current survey strategies will implement a rolling cadence that cycles through multiple filters.

Another potential improvement would be the development of a larger and more diverse training sample which in turn could be used to create a "smarter" machine-learning-based filter. The population statistics presented in this sample are generated from 14 objects in two bands. If one were to include more objects and more filter information, there would be potentially additional, or at least robuster, features to translate into a filter. We plan to explore this avenue in future work.

When optimizing alert stream filter completeness, it's important to not let the purity suffer, i.e. labeling more non-target objects as targets. Framing this concern in terms of photometric (specifically UV) and spectroscopic follow-up, we cannot afford—both in time and money—to send hundreds of potentially interesting objects to these follow-up facilities; we must be skeptical and restrained. In an era with millions of transient alerts each night, it would be folly to expect to catch and care about each and every alert, so regardless of completeness, astronomers as a whole must be okay with missing some of our target objects.

6. CONCLUSIONS

We present 14 spectroscopically confirmed SNe IIb with double-peaked light curves using publicly available survey data from ZTF and ATLAS forced photometry (observed between 2018–2022). Each object's light curve shows an initial peak powered by shock-cooling emission (SCE) which occurs when the shock passes through and heats up the stellar envelope, and a secondary peak powered by the radioactive decay of nickel-56 (i.e., the classical stripped-envelope SNe light curve powering mechanism).

We develop a "lighting bolt" model to measure and describe the early-time photometric behavior, specifically focusing on quantifying the rise and fall rates of the SCE peak, which we robustly fit using MCMC.

We generate the first ever early-time photometric population statistics from the mean of the 14 fits. We find that the rise to the SCE peak evolves 25x faster than the rise to the nickel-powered peak, at an average rate of ≈ 2 mags/day in the g-band. We find that the decline from the SCE peak lasts an average of one week in both bands.

We measure the time above half-maximum flux ($t_{1/2}$) for the SCE peak for 8 objects in the g-band and 6 in the r-band as well as absolute magnitudes, corrected for line-of-sight MW extinction, for all objects. We find that in the g-band, the average $t_{1/2}$ was ≈ 5 days and that average SCE peaks at ≈ -17 magnitude. We compare the spread of the SCE from SNe IIb to the latest FBOT parameter space (Ho et al. 2023) and find that all SCE light curves lie below the 12-day classical FBOT cutoff, with the shortest $t_{1/2} \approx 2.5$ days and the longest $t_{1/2} \approx 9$ days. We interpret our findings to mean that this well-known SCE phenomenon in SNe IIb could be mistaken as an exotic FBOT/fast-evolving transient, if the main second, Ni-driven peak is not observed/followed due to too short of a time baseline in surveys.

Finally, we present a preliminary framework for a future alert stream filter that implements feature engineering from the lightning bolt model using the ANTARES broker. We test the filter performance on archival light curves, comprising only alert stream data, for 10 target objects (i.e. known double-peaked SNe IIb) and 33 non-target objects. Initial testing resulted in an overall accuracy of 83.7% and a completeness of 80.0%.

SNe IIb are important probes into the mechanisms that lead massive stars to lose their outermost envelopes. The double-peaked light curves of SNe IIb are uniquely rich in progenitor tracers as the additional SCE peak holds direct tracers of the progenitor stellar envelope as well as tracers of the shock wave geometries and CSM environment. Past works have shown the importance of observing this SCE peak in the UV, where the emission peaks, and the importance of observations in the UV and optical not just on the decline of the SCE peak but along the rise as well (Pellegrino et al. 2023). Thus, observing these objects, especially at early-time during the SCE peak, is very useful for improving our stellar models and better understanding supernova physics. Future work for these objects lies in computing the ratios between the SCE and Ni-powered peaks on a statistically large sample as well as a meta-analysis of all published SNe IIb with double peaks across all bands.

ACKNOWLEDGMENTS

This work is a condensed version of a paper that was submitted to the Astrophysical Journal and is currently in review. Read the full length paper at the following link: http://arxiv.org/abs/2503.03735. AC acknowledges support from the Virginia Space Grant Consortium (VSGC) in part from the Graduate Research Fellowship. M.M. and the METAL group at UVA acknowledge support in part from ADAP program grant 80NSSC22K0486, from NSF grant AST-2206657, and from *HST* program GO-16656. TP acknowledges that the material is based upon work supported by NASA under award number 80GSFC24M0006. The ZTF forced-photometry service was funded under the Heising-Simons Foundation grant #12540303 (PI: Graham).

REFERENCES

Aldering, G., Humphreys, R. M., & Richmond, M. 1994, AJ, 107, 662, doi: 10.1086/116886

Andreoni, I., Coughlin, M. W., Kool, E. C., et al. 2021, The Astrophysical Journal, 918, 63, doi: 10.3847/1538-4357/ac0bc7

Arcavi, I., Gal-Yam, A., Yaron, O., et al. 2011, The Astrophysical Journal, 742, L18, doi: 10.1088/2041-8205/742/2/l18

Arcavi, I., Hosseinzadeh, G., Brown, P. J., et al. 2017, ApJL, 837, L2, doi: 10.3847/2041-8213/aa5be1

Arnett, W. D. 1980, ApJ, 237, 541, doi: 10.1086/157898
Barbary, K. 2021, extinction: Dust extinction laws, Astrophysics Source Code Library, record ascl:2102.026
Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002, doi: 10.1088/1538-3873/aaecbe
Bruch, R. J., Gal-Yam, A., Schulze, S., et al. 2021, ApJ, 912, 46, doi: 10.3847/1538-4357/abef05
Bruch, R. J., Gal-Yam, A., Yaron, O., et al. 2023, ApJ, 952, 119, doi: 10.3847/1538-4357/acd8be
Das, K. K., Kasliwal, M. M., Fremling, C., et al. 2023, ApJ, 959, 12, doi: 10.3847/1538-4357/acfeeb
de Soto, K. M., Villar, V. A., Berger, E., et al. 2024, ApJ, 974, 169, doi: 10.3847/1538-4357/ad6a4f
Dessart, L., Yoon, S.-C., Aguilera-Dena, D. R., & Langer, N. 2020, A&A, 642, A106, doi: 10.1051/0004-6361/202038763
Drout, M. R., Chornock, R., Soderberg, A. M., et al. 2014, The Astrophysical Journal, 794, 23, doi: 10.1088/0004-637x/794/1/23
Ertini, K., Folatelli, G., Martinez, L., et al. 2023, MNRAS, 526, 279, doi: 10.1093/mnras/stad2705
Farah, J. R., Howell, D. A., Terreran, G., et al. 2025, arXiv e-prints, arXiv:2501.17221, doi: 10.48550/arXiv.2501.17221
Filippenko, A. V. 1997, ARA&A, 35, 309, doi: 10.1146/annurev.astro.35.1.309
Fitzpatrick, E. L., & Massa, D. 2007, ApJ, 663, 320, doi: 10.1086/518158
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, Publications of the Astronomical Society of the Pacific, 125, 306–312, doi: 10.1086/670067
Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, AJ, 161, 242, doi: 10.3847/1538-3881/abe9bc
Fremling, C., Ko, H., Dugas, A., et al. 2019, The Astrophysical Journal, 878, L5, doi: 10.3847/2041-8213/ab218f
Graham, M. L., Bellm, E., Guy, L., et al. 2024, LSST Alerts: Key Numbers, doi: 10.5281/zenodo.11062350
Green, G. 2018, The Journal of Open Source Software, 3, 695, doi: 10.21105/joss.00695
Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, The Astrophysical Journal Supplement Series, 208, 19, doi: 10.1088/0067-0049/208/2/19
Ho, A. Y. Q., Perley, D. A., Beniamini, P., et al. 2020, The Astrophysical Journal, 905, 98, doi: 10.3847/1538-4357/abc34d
Ho, A. Y. Q., Perley, D. A., Gal-Yam, A., et al. 2023, ApJ, 949, 120, doi: 10.3847/1538-4357/acc533
Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111, doi: 10.3847/1538-4357/ab042c
Jacobson-Galán, W. V., Venkatraman, P., Margutti, R., et al. 2022, ApJ, 932, 58, doi: 10.3847/1538-4357/ac67dc
Khakpash, S., Bianco, F. B., Modjaz, M., et al. 2024, arXiv e-prints, arXiv:2405.01672, doi: 10.48550/arXiv.2405.01672
Kochanek, C. S., Shappee, B. J., Stanek, K. Z., et al. 2017, PASP, 129, 104502, doi: 10.1088/1538-3873/aa80d9
Masci, F. J., Laher, R. R., Rusholme, B., et al. 2023, arXiv e-prints, arXiv:2305.16279, doi: 10.48550/arXiv.2305.16279
Matheson, T., Stubens, C., Wolf, N., et al. 2021, The Astronomical Journal, 161, 107, doi: 10.3847/1538-3881/abd703
Modjaz, M., Gutiérrez, C., & Arcavi, I. 2019, Nature Astronomy, 3, 717, doi: 10.1038/s41550-019-0856-2
Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., & Hložek, R. 2019, PASP, 131, 118002, doi: 10.1088/1538-3873/ab1609
Nakar, E., & Piro, A. L. 2014, The Astrophysical Journal, 788, 193, doi: 10.1088/0004-637x/788/2/193
Pellegrino, C., Hiramatsu, D., Arcavi, I., et al. 2023, The Astrophysical Journal, 954, 35, doi: 10.3847/1538-4357/ace595
Rabinak, I., & Waxman, E. 2011, ApJ, 728, 63, doi: 10.1088/0004-637X/728/1/63
Shingles, L., Smith, K. W., Young, D. R., et al. 2021, Transient Name Server AstroNote, 7, 1
Smith, K. W., Smartt, S. J., Young, D. R., et al. 2020, PASP, 132, 085002, doi: 10.1088/1538-3873/ab936e
Smith, N. 2014, ARA&A, 52, 487, doi: 10.1146/annurev-astro-081913-040025
Soderberg, A. M., Margutti, R., Zauderer, B. A., et al. 2012, ApJ, 752, 78, doi: 10.1088/0004-637X/752/2/78
Soraisam, M., Matheson, T., Lee, C.-H., et al. 2022, ApJL, 926, L11, doi: 10.3847/2041-8213/ac4e99
Tartaglia, L., Fraser, M., Sand, D. J., et al. 2017, The Astrophysical Journal, 836, L12, doi: 10.3847/2041-8213/aa5c7f
Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, PASP, 130, 064505, doi: 10.1088/1538-3873/aabadf
Van Dyk, S. D., Zheng, W., Fox, O. D., et al. 2014, The Astronomical Journal, 147, 37, doi: 10.1088/0004-6256/147/2/37
Woosley, S. E., Eastman, R. G., Weaver, T. A., & Pinto, P. A. 1994, ApJ, 429, 300, doi: 10.1086/174319
Woosley, S. E., Pinto, P. A., Martin, P. G., & Weaver, T. A. 1987, ApJ, 318, 664, doi: 10.1086/165402
Yang, S., Valenti, S., Cappellaro, E., et al. 2017, ApJL, 851, L48, doi: 10.3847/2041-8213/aaa07d