

Searching for late-time interaction between Type Ia supernovae and distant circumstellar material using the Zwicky Transient Facility

A dissertation submitted to the University of Dublin
for the degree of Doctor of Philosophy

Jacobus (Jacco) Hendrikus Terwel

*School of Physics
Trinity College Dublin
The University of Dublin*

Supervisor:
Prof. Kate Maguire

2025



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Name: Jacobus Hendrikus Terwel

Signature: **Date:** 30/09/2024

Summary

The nature of the progenitor systems and explosion mechanisms that give rise to Type Ia supernovae (SNe Ia) is still debated. In rare cases the ejecta interact with circumstellar material (CSM) that was ejected from the progenitor system before the explosion. As it is being swept up by the expanding ejecta, the CSM starts to glow and becomes an additional light source that can be bright enough to alter the SN evolution drastically. By studying the interaction signature, the properties and composition of the CSM can be revealed, which can constrain the type of progenitor system it was ejected from. However, most previous studies have focused on finding CSM that was ejected shortly before the SN Ia explosion and still resides close to the explosion site, resulting in short delay times until the interaction starts.

In this thesis, I search data from the Zwicky Transient Facility (ZTF) for signs of late-time interaction between SN Ia ejecta and distant CSM, where interaction starts > 100 days after the explosion. I use SN 2015cp, where interaction was discovered 664 days after peak brightness, as a prototype event and develop a pipeline to search for late-time rebrightening events in ZTF forced photometry light curves. By binning the late-time light curve data, I push the detection limit as deep as possible and recover faint signals below the detection limit of individual data points.

I search through three samples of transients: 1) The ZTF second data release (ZTF SN Ia DR2), which contains all spectroscopically classified SNe Ia discovered by ZTF from March 2018 to October 2020, 2) the pre-ZTF sample, which consists of all transients discovered between 2008 and 2018, 3) all SNe Ia discovered by ZTF before 8 July 2023, whose late-time evolution I monitored from 13 November 2023 to 9 July 2024 to find active late-time rebrightening events and follow these up with secondary telescopes. In total, I analyse 11 973 unique spectroscopically confirmed Type Ia SNe and 2 728 other transients. Out of these, I identify seven SNe Ia that show signs of possible late-time CSM interaction starting between 250 and 2 200 days after the explosion and lasting between 100 and 500 days. I also discovered a short (~ 50 d) signal in SN 2020qxx, which I spectroscopically confirmed to be a period of rebrightening around 1 150 days after peak light in the SN rest frame. Besides these, also recover several other types of objects, including sibling transients, declining tails of bright SNe, and a group of (faint) nuclear transients. This shows my pipeline's ability to find various types of faint signals that are only recoverable through binning.

Using the well-defined nature of the ZTF SN Ia DR2 and simulations of the survey I estimate that $< 0.5\%$ of normal SNe Ia display late-time (> 100 days after the peak) strong $H\alpha$ -dominated CSM interaction. This is equivalent to an absolute rate of 8_{-4}^{+20} to 54_{-26}^{+91} $\text{Gpc}^{-3} \text{ yr}^{-1}$, assuming a constant SN Ia rate of $2.4 \times 10^{-5} \text{ Mpc}^{-3} \text{ yr}^{-1}$ for $z \leq 0.1$. Weaker interaction signatures of $H\alpha$ emission, more similar to the strength seen in SN 2015cp, could be

more common but are difficult to constrain with the depth of ZTF. By using a simple CSM model with bremsstrahlung as the main emission mechanism during the interaction, I show that interaction with a thin, distant shell of CSM containing $\lesssim 5 M_{\odot}$ of material can produce a signal that is detectable with ZTF, even if the interaction starts up to six years after the SN.

Most of the identified candidates reside close to the nuclei of their host galaxies, which suggests that environment or specific progenitor characteristics play a role in the production of the potential late-time CSM signatures in these SNe Ia. For candidates close to the host nucleus, distinguishing between nuclear activity and late-time CSM can also be difficult. The exception to this is SN 2020qxx, a SN Ia with an earlier period of CSM interaction as well as a short late-time signal. Through spectroscopic follow-up I confirm four transient emission lines, which I associate to H β , Ca II, N I, and K I, and are blueshifted by $5150 - 5920 \text{ km s}^{-1}$. This late-time signal can be interpreted as swept-up nearby CSM of which a part interacts with a cloud of more distant CSM at 1150 days after peak brightness.

Strong late-time CSM interaction is very rare, and the only way to study them is by systematically searching for these events by monitoring a large group of known SNe using a survey such as ZTF. Discovering late-time interaction signals while they are still active is crucial for deep photometric and spectroscopic follow-up to constrain their properties. CSM trace the history of the progenitor system, and late-time interaction with distant CSM contains clues on the progenitor system as it was long before exploding. By measuring CSM properties such as its location, mass, geometry, and composition, it is possible to learn about this history and put constraints on the type of progenitor system.

*To my family, who have always stood by me and encouraged me
to follow my passion and my dreams.*

Now the J.H. Terwel degeneracy has finally been broken.

Acknowledgements

First of all, I want to thank Prof. Kate Maguire, who has supported me and provided guidance for me to get where I am now. You gave me the freedom to do things my way, and fully supported me every time I decided I wanted to be somewhere else than Dublin. I am extremely grateful for everything you have done to make my multinational journey possible. I like to think that I was still able to do a satisfactory amount of work, given that I now have three more years of working with you to look forward to. I promise I'll be in Dublin a bit more this time around.

I also want to thank the postdocs in the TCD transients group. Georgios Dimitriadis, thank you for your support in getting me into the positions that I did, and thank you for reducing all those IDS spectra for me. Umut Burgaz, thank you for the enlightening conversations we've had, I look forward to working together with you some more. Simon Prentice and Mark Magee, though we never got to work together in person due to (cough) external reasons, I still learned a lot from you and want to thank you for helping me at the start of my PhD.

Then there are the other PhD students in the group, Maxime Deckers and Luke Harvey. Though our paths were destined to diverge at some point, I guess that happened a bit sooner than expected. Still, doing this PhD together with you has truly been delightful. Maxime, I wish you all the best with your family, and make sure to teach them to speak Dutch. Luke, you are proof that the smallest of chance encounters can ignite a runaway process to become the most important thing in life, I wish the two of you all the best in the south of France. At the very least the climate there beats that of Dublin.

Of course I cannot forget the other astronomers at Trinity, my collaborators at the Zwicky Transient Facility, and the people from the PESSTO collaboration. Thank you all for the many enlightening discussions and the amazing opportunities that helped me improve my work in many ways.

I want to thank my friends in La Palma (or wherever you are now), the staff and students from the ING, NOT, and other telescopes of the ORM, for the amazing opportunities these studentships have given me. My time with all of you on this island has changed me forever. You are the greatest group of friends I've ever had, and I'm not just talking about the size. The main reason why I decided to spend half my PhD here is because of all of you. And a special thanks to all of you who generously observed my targets for me: 77 classifications and multiple papers that make use of the data, this would not have been possible without your help.

To my friends back in the Netherlands, Koen, Jesse, Marieke, Yann, we all have chosen very different paths in life, following our own passions and interests. Yet our time together started a friendship that will last a lifetime, even if some of us choose to live on a faraway island (or La Palma).

I never would have started this journey without the everlasting support of my family. Mom, I know you would have preferred me to live closer, but even though the town I grew up in is over the mountain, that mountain does not house the telescopes I need for my work. Thanks for standing by me nonetheless. Dad, your levelheadedness and calm outlook on life is a state of mind I can only strive to achieve. And when I inevitably end up breaking something, you are always there to help fix it. Twan, my big little brother, you are one of my closest friends, and your dry humour can make (or break) my day. Let's play another game sometime. Grandma, it is rare to see someone with such a love for life as you, you are the reason why I'm always up for a stupid idea. And Grandpa, your inquisitive mind shines through in both our conversations and our games of Shogi, and your interest in my work makes you my biggest supporter. Knight to B8.

Last but not least, 赵允超, though we only met close to the end of this endeavour of mine, you have shown me more love, support, and enthusiasm than I ever could have dreamt of. I look forward to the day when we finally meet in person, and every day after that.

List of Publications

Publications

1. Transitional events in the spectrophotometric regime between stripped envelope and superluminous supernovae ([Prentice et al., 2021](#))
2. SN 2016dsg: A Thermonuclear Explosion Involving a Thick Helium Shell ([Dong et al., 2022](#))
3. An elliptical accretion disk following the tidal disruption event AT 2020zso ([Wevers et al., 2022](#))
4. Photometric study of the late-time near-infrared plateau in Type Ia supernovae ([Deckers et al., 2023](#))
5. A Systematic Study of Ia-CSM Supernovae from the ZTF Bright Transient Survey ([Sharma et al., 2023](#))
6. Early-time spectroscopic modelling of the transitional Type Ia Supernova 2021rhu with TARDIS ([Harvey et al., 2023](#))
7. SN 2022joj: A Peculiar Type Ia Supernova Possibly Driven by an Asymmetric Helium-shell Double Detonation ([Liu et al., 2023a](#))
8. Ground-based and JWST Observations of SN 2022pul. I. Unusual Signatures of Carbon, Oxygen, and Circumstellar Interaction in a Peculiar Type Ia Supernova ([Siebert et al., 2024](#))
9. Ground-based and JWST Observations of SN 2022pul. II. Evidence from Nebular Spectroscopy for a Violent Merger in a Peculiar Type Ia Supernova ([Kwok et al., 2024](#))
10. A precursor plateau and pre-maximum [O II] emission in the superluminous SN2019szu: a pulsational pair-instability candidate ([Aamer et al., 2024](#))
11. ZTF SN Ia DR2: Searching for late-time interaction signatures in Type Ia supernovae from the Zwicky Transient Facility ([Terwel et al., 2025](#))
12. The complex circumstellar environment of supernova 2023ixf ([Zimmerman et al., 2024](#))
13. ZTF SN Ia DR2: Overview ([Rigault et al., 2025a](#))
14. ZTF SN Ia DR2: Peculiar velocities impact on the Hubble diagram ([Carreres et al., 2025](#))

-
15. **ZTF SN Ia DR2: Colour standardisation of Type Ia Supernovae and its dependence on environment** ([Ginolin et al., 2025a](#))
 16. **ZTF SN Ia DR2: Environmental dependencies of stretch and luminosity of a volume limited sample of 1,000 Type Ia Supernovae** ([Ginolin et al., 2025b](#))
 17. **ZTF SN Ia DR2: Impact of the galaxy cluster environment on the stretch distribution of Type Ia supernovae** ([Ruppin et al., 2025](#))
 18. **ZTF SN Ia DR2: Study of Type Ia Supernova lightcurve fits** ([Rigault et al., 2025b](#))
 19. **ZTF SN Ia DR2: Evidence of Changing Dust Distributions With Redshift Using Type Ia Supernovae** ([Popovic et al., 2025](#))
 20. **ZTF SN Ia DR2: Exploring SN Ia properties in the vicinity of underdense environments** ([Aubert et al., 2025](#))
 21. **ZTF SN Ia DR2: The secondary maximum in Type Ia supernovae** ([Deckers et al., 2025](#))
 22. **ZTF SN Ia DR2: Cosmology-independent constraints on Type Ia supernova standardisation from supernova siblings** ([Dhawan et al., 2024](#))
 23. **ZTF SN Ia DR2: The spectral diversity of Type Ia supernovae in a volume-limited sample** ([Burgaz et al., 2025a](#))
 24. **ZTF SN Ia DR2: Simulations and volume limited sample** ([Amenouche et al., 2025](#))
 25. **ZTF SN Ia DR2: The diversity and relative rates of the thermonuclear SN population** ([Dimitriadis et al., 2025](#))
 26. **The Einstein Probe transient EP240414a: Linking Fast X-ray Transients, Gamma-ray Bursts and Luminous Fast Blue Optical Transients** ([van Dalen et al., 2024](#))
 27. **ZTF SN Ia DR2: An environmental study of Type Ia supernovae using host galaxy image decomposition** ([Senzel et al., 2025](#))
 28. **ZTF SN Ia DR2: Properties of the low-mass host galaxies of Type Ia supernovae in a volume-limited sample** ([Burgaz et al., 2025b](#))
 29. **ZTF SN Ia DR2: High-velocity components in the Si II λ 6355** ([Harvey et al., 2025](#))

Contents

List of Publications	vii
List of Figures	xiii
List of Tables	xv
1 Introduction	1
1.1 The final stages of stars	1
1.1.1 Stars below $8 M_{\odot}$	1
1.1.2 Stars above $8 M_{\odot}$	2
1.2 The transient universe	4
1.2.1 Core-collapse SNe	6
1.2.2 Thermonuclear SNe	7
1.2.2.1 Standardizable candles	9
1.2.3 Other types of transients	10
1.2.3.1 SLSN - Superluminous Supernova	10
1.2.3.2 AGN - Active Galactic Nucleus	10
1.2.3.3 TDE - Tidal Disruption Event	11
1.2.3.4 ANT - Ambiguous Nuclear Transient	12
1.2.3.5 Variable stars	12
1.3 SN Ia subclasses	12
1.3.1 Photometric differences	14
1.3.2 Spectroscopic differences	14
1.3.3 The SN Ia-91T subclass	15
1.4 SN Ia progenitor systems	16
1.4.1 Single degenerate	16
1.4.2 Double degenerate	18
1.5 SNe with circumstellar material	19
1.5.1 Interaction between SN ejecta and CSM	21
1.5.2 Other CSM signatures	23
1.5.3 The SN Ia-CSM subclass	23
1.5.4 CSM in CCSNe	24
1.5.5 Distant CSM	25
1.6 In this thesis:...	27
2 Observing in the optical regime	29
2.1 Types of observations	29
2.1.1 Imaging	31
2.1.2 Spectroscopy	31
2.2 Telescopes	33

2.2.1	Zwicky Transient Facility	33
2.2.2	Nordic Optical Telescope	34
2.2.3	Gran Telescopio CANARIAS	35
2.2.4	Other observations	36
2.3	Calibration images	37
2.3.1	Bias	37
2.3.2	Dark	37
2.3.3	Flatfield	38
2.3.4	Arc	38
2.4	Reduction	39
2.4.1	Bias, dark, and flat corrections	40
2.4.2	Cosmic-ray removal and image stacking	40
2.4.3	Source extraction and standard star calibration	41
2.4.3.1	Photometry	41
2.4.3.2	Spectroscopy	46
2.4.4	Difference imaging	47
2.4.5	Forced photometry	49
2.5	General considerations for observing	50
2.5.1	Location	50
2.5.2	Telescope, instrument, observation type, and setup	51
2.5.3	Night plan	52

3 ZTF SN Ia DR2: Searching for late-time interaction signatures in Type Ia supernovae from the Zwicky Transient Facility 55

3.1	Data	56
3.1.1	ZTF light curve data	56
3.1.2	Baseline correction	57
3.2	Analysis	59
3.2.1	Binning & filtering program	59
3.2.2	Removal of SN Ia radioactive tail detections	60
3.2.3	SuperNova Animation Program (SNAP)	62
3.2.4	Simulated interaction recovery fractions	64
3.2.4.1	The interacting SN model	64
3.2.4.2	Simulating the observing campaign	67
3.2.4.3	Simulated interaction recovery	69
3.2.4.4	Impact of reference image depth	73
3.3	Results	73
3.3.1	Initial summary of detected events	74
3.3.1.1	Known Ia-CSM	77
3.3.1.2	Siblings	78
3.3.1.3	Kinky tails	81
3.3.2	Additional tests of promising events	82
3.3.2.1	Scene modelling photometry	84
3.3.2.2	AGN contamination	85
3.3.2.3	Presence of a sibling close to the SN location	85
3.3.2.4	Late-time interaction candidates	93
3.3.3	Late-time CSM interaction rates based on my candidate objects	99
3.4	Discussion	103

3.4.1	Late-time interaction is rare in SNe Ia	103
3.4.2	Properties of SNe Ia with late-time interaction	105
3.4.3	Limitations of the analysis	107
3.4.4	By-product pipeline detections	108
3.5	Conclusions	109
4	ZTF-observed late-time signals of pre-ZTF transients	113
4.1	Data	114
4.2	Analysis	117
4.2.1	Detection pipeline modifications	117
4.2.2	Removing false positives	119
4.3	Results	120
4.3.1	Pre-ZTF transients still active in ZTF	121
4.3.1.1	AT 2017gpv - a 14hls-like event	124
4.3.2	Siblings	125
4.3.3	Non-SN sources of flux excesses in ZTF	128
4.3.4	Final shortlist of late-time interaction	131
4.3.4.1	Estimating CSM masses	134
4.4	Discussion	147
4.4.1	Rarity of late-time signals from SNe	147
4.4.1.1	Siblings	147
4.4.1.2	Serendipitous nuclear transient detections	148
4.4.1.3	CSM interaction	149
4.4.2	Limitations	151
4.5	Conclusions	153
5	A real-time search for Type Ia SNe with late-time CSM interaction in ZTF	155
5.1	Data	156
5.2	Analysis	157
5.2.1	Follow-up observations and reduction	158
5.3	Results	159
5.3.1	Faded objects	162
5.3.2	Follow-up results	165
5.3.3	SN 2020qxx	171
5.3.3.1	Late-time photometry	172
5.3.3.2	Host spectroscopy	172
5.3.3.3	Late-time spectroscopy at the SN location	175
5.4	Discussion	179
5.4.1	Detection challenges	179
5.4.2	Follow-up campaigns	180
5.4.3	The late-time signal in SN 2020qxx	181
5.5	Conclusions	184
6	Conclusions and future work	187
6.1	Conclusions	187
6.2	Future work	191
6.2.1	Improvements to the pipeline	191

6.2.2	Alternative uses for the pipeline	192
6.2.3	Other future work	193
References		199

List of Figures

1.1	Schematic view of the onion structure of a highly evolved massive star.	3
1.2	<i>gri</i> composite image of NGC 4216 using observations taken by ZTF.	5
1.3	Rest frame, absolute mangitude ZTF <i>gri</i> light curve of a normal SN Ia.	8
1.4	SN Ia subclasses.	13
1.5	Artistic impressions of single and double degenerate systems.	17
1.6	Schematic view of a SN Ia-CSM event.	20
1.7	Schematic view of the CSM interaction region.	22
2.1	Image and partial spectrum of SN 2024nqr and SN 2024pgd.	30
2.2	Throughput of different filters and grisms used in this thesis.	34
2.3	Schematic view of a point source on a CCD	42
2.4	Example of a reference, science, and the resulting difference image.	48
2.5	Example of a night plan.	53
3.1	Example of a binned light curve.	61
3.2	Example frame from the SuperNova Animation Program (SNAP).	63
3.3	Spectral comparision between SN 2011fe + $H\alpha$ model and SN 2015cp.	67
3.4	<i>r</i> and <i>i</i> model light curves.	68
3.5	Simulated late-time interaction recovery as a function of redshift.	70
3.6	Recovered known SN Ia-CSM objects.	76
3.7	Recovered sibling transients.	78
3.8	Light curve of SN 2018big and its sibling SN 2019nvm.	79
3.9	Recovered SNe Ia with kinks in their tails.	82
3.10	Objects whose late-time signal was due to photometry issues.	85
3.11	Objects whose late-time signal can be explained without CSM.	86
3.12	Spectrum of the late-time signal in SN 2020alm.	91
3.13	Light curves of the three candidate objects.	94
3.14	Colour curves of the three candidate objects.	95
3.15	Binned SMP light curves of two of the three candidate objects.	96
3.16	Chance of one detected event vs. late-time CSM interaction fraction	103
4.1	Classes in the pre-ZTF sample.	115
4.2	Sample size as a function of redshift for all classes.	116
4.3	Examples of pre-ZTF SNe with tails detectable in ZTF.	122
4.4	pre-ZTF and ZTF light curve of AT 2017gpv.	124
4.5	Binned flux of the recoverd AGN.	129
4.6	Binned flux of the recovered variable stars.	130
4.7	Binned light curves of the objects that require more investigation.	133
4.8	M_{CSM} curves for different conditions.	138
4.9	Comparison of detected late-time nuclear transients with an ANT.	140

5.1	Binned light curves of the three objects that could not be followed up on.	163
5.2	Binned light cuves of the objects that were followed up on.	165
5.3	Location and host galaxy of SN 2019zbq, imaged with ZTF and GTC. .	169
5.4	GTC spectrum of the late-time transient in the SN 2019zbq host galaxy.	170
5.5	Pan-STARRS DR1 colour image of the SN 2020qzx location and host. .	171
5.6	Spectra of SN 2020qzx.	173
5.7	Regions with transient emission lines in SN 2020qzx spectra.	174
5.8	Individual spectra of SN 2020qzx.	176
5.9	Individual spectra of SN 2020qzx zoomed in on transient line regions. .	177

List of Tables

3.1	DR2 sample size and its reduction in each step of the analysis process.	58
3.2	List of spectra used in the SN 2011fe model.	65
3.3	z_{50} recovery values for all simulations.	72
3.4	List of objects that passed the initial visual inspections.	75
3.5	Objects with a detected sibling transient.	79
3.6	Additional tests performed on the tem most promising objects.	83
3.7	Details of the used comparison transients.	87
3.8	Parameters used for the rate estimation simulations for each object. . .	101
4.1	Different baseline regions used in the pre-ZTF sample pre-processing. .	118
4.2	List of pre-ZTF SNe with tails detectable in ZTF.	122
4.3	pre-ZTF transients with a detected sibling in ZTF.	127
4.4	Pre-ZTF transients that were detected due to stellar or AGN variability.	131
4.5	Objects with ZTF detections that warrant further investigation.	132
4.6	Number of recovered signals for each SN type.	148
5.1	Objects with a detected late-time signal.	160
5.2	Follow-up observations of transients with late-time signals.	161
5.3	Fits of the transient emission lines in the late-time SN 2020qxx spectra.	178

Introduction

1.1 The final stages of stars

A star is a big ball of plasma that is gravitationally bound together. Consisting primarily of hydrogen and helium upon first formation, the extreme heat and pressure in stellar cores allows them to fuse hydrogen into helium, which generates the energy needed to maintain hydrostatic equilibrium and prevent further gravitational collapse. [Kippenhahn et al. \(2013\)](#) gives an excellent and detailed review of the evolution of stars in different mass ranges, and I will only quickly go over it in very broad terms.

When the hydrogen supply is depleted the inward pressure from gravity causes the core to contract until something is able to stop it once more. This can be matter in the core turning degenerate and providing an electron degeneracy pressure, or the start of helium fusion. What happens next, and the ultimate fate of a star, depends on its mass, which can range from $\sim 0.08 M_{\odot}$ to well over $100 M_{\odot}$.

1.1.1 Stars below $8 M_{\odot}$

The lightest stars have a theoretical lifespan larger than the current age of our universe, and as such the final stages of their evolution cannot be observed yet. For stars with a mass between approximately $0.6 M_{\odot}$ and $8 M_{\odot}$, the outer layers increase dramatically in size when the core is depleted of hydrogen, and they become red giants as they fuse helium into carbon and oxygen in their cores and hydrogen in a shell around the core. Eventually, the outer layers are blown away while the core is left behind to cool down and become a white dwarf (WD).

WDs are degenerate objects that are primarily made up of carbon and oxygen. As electrons are fermions, the Pauli exclusion principle ([Pauli, 1925](#)) forces each one to occupy a different quantum state. Due to their high density in the WD, the gas becomes degenerate as particles are forced into higher energy states. The degenerate

electron gas provides a pressure that is independent of temperature, allowing the WD to theoretically cool to absolute zero without any significant change in the stellar structure (although some models predict WDs to crystallise at a certain temperature, [Mochkovitch 1983](#); [Isern et al. 1997](#)). Assuming that the electron gas is non-relativistic, it can be shown that the radius of a WD is related to its mass as

$$\frac{R_{\text{WD}}}{R_{\odot}} = 0.010 \left(\frac{M_{\odot}}{M_{\text{WD}}} \right)^{1/3}, \quad (1.1)$$

with R_{WD} the radius of the WD and M_{WD} its mass. This shows that WDs are roughly planet-sized, and that more massive WDs are smaller. However, as the mass increases, the electrons in the WD become relativistic and Equation 1.1 is no longer valid. In this case the radius decreases much more drastically with increasing mass, up to the limit of

$$M_{\text{Ch}} \approx 1.4M_{\odot}, \quad (1.2)$$

which is the so-called Chandrasekhar limit ([Chandrasekhar, 1931](#)). A WD at the Chandrasekhar mass M_{Ch} has a radius of 0, and the electron degeneracy pressure cannot stop a WD from fully collapsing in on itself.

1.1.2 Stars above $8 M_{\odot}$

More massive stars are able to burn the carbon that is left after burning helium. New elements are created and subsequently burned under increasingly more extreme conditions, resulting in a core that is structured like an onion. Each ring consists of a different fusion product, with heavier elements towards the core. Figure 1.1 shows a schematic view of the layered structure of a highly evolved massive star. Each stage of fusion is less efficient, and once iron is reached, further fusion consumes energy rather than releasing it. The core becomes unable to support itself and collapses into a neutron star or black hole, depending on the mass of the star. The outer layers fall inwards but bounce back, gaining enough energy to become unbound from the star. $\sim 99\%$ of




Figure 1.1: Schematic view of the onion structure of a highly evolved massive star just before core collapse, with light elements on the outside and heavier elements towards the core. Typical values are given along the horizontal and vertical axes. This figure originally comes from chapter 35 of [Kippenhahn et al. \(2013\)](#).

the total energy that is released in this process is carried away by neutrinos, while only $\sim 1\%$ of the energy ($\sim 10^{51}$ ergs) is deposited in the ejecta (Mirizzi et al., 2016). The high densities allow for elements heavier than iron to be fused in these brief moments. The stellar demise is known as a core-collapse supernova (CCSN) (Kippenhahn et al., 2013). Section 1.2.1 contains more details on these events.

1.2 The transient universe

Astronomy is filled with large scales, and when studying processes as they evolve, our limited life span often prevents us from studying a single object through its entire evolution. Most processes occur on such long timescales that we only have a snapshot of the state of each object. It is only by combining a large ensemble of objects together with theory and simulations that we are able to understand the bigger picture.

Transients are, by definition, fleeting things, and in extrasolar astronomy these are objects or events that occur on short timescales of mere hours to a few years. This means we are able to observe the entire evolution of a single object, as long as it is discovered early. The most famous examples of transients are supernovae (SNe), exploding stars at the end of their lives releasing enough light that they can outshine their entire host galaxy for a few weeks to months before fading away again. Besides these there are other types of transients as well, a selection of which are given in section 1.2.3.

Figure 1.2 is a *gri* composite image of NGC 4216, taken with the Zwicky Transient Facility (ZTF, see section 2.2.1). The left image is made using data before the start of 2024. The right image is made using data from January 2024 and clearly shows an extra source in the southern part of the galaxy. This is SN 2024gy, a Type Ia SN (see section 1.2.2) first discovered on 4 January 2024 (Itagaki, 2024). Despite the galaxy being at a distance of 16.194 ± 0.491 Mpc (de Vaucouleurs et al., 1991), the SN still outshines the Milky Way stars that are in the same image.



Figure 1.2: *gri* composite image of NGC 4216 using observations taken by ZTF.
Left: composite image of observations taken before 1 January 2024.
Right: composite image of observations taken between 5 and 19 January 2024, the first two weeks after the first detection of the Type Ia SN 2024gy. (Credit: Benjamin Nobre Hauptmann)

1.2.1 Core-collapse SNe

As was mentioned at the end of section 1.1.2, stars above $8 M_{\odot}$ end their lives in a CCSN when the core implodes into a compact object due to the effects of gravity. The outer layers fall in as well but bounce back outwards as the enormous amounts of released gravitational energy from the collapsing core is imparted on the outer layers, heating them up and expelling them outwards at high velocities. The light coming from the cooling of this material and the radioactive decay of newly formed unstable nuclei into stable ones cause the SN to rise rapidly in brightness. Within hours after the explosion it is bright enough to be observed in other galaxies, peaking after several days to weeks at an absolute B-band magnitude in the range of -16.5 mag to -18.5 mag depending on the type of CCSN (Richardson et al., 2014). This is what we observe as the SN event, and by analysing their spectra around peak light they can be classified based on the visible elemental emission and absorption lines.

The composition of the outer layers of the progenitor depends heavily on the life of the star. Most CCSNe have hydrogen in their envelope, which will show up in the SN spectra as strong Balmer P Cygni profiles. These are known as Type II SNe. Photometric features of the light curve (the SN brightness as it changes over time) are also used in SN identification, leading to subtypes such as SN IIL or SN IIP for Type II SNe that show a linear decline or a plateau, respectively. High mass-loss rates, due to a high stellar wind (Woosley et al., 1995) or the presence of a binary companion (Podsiadlowski et al., 1992), can however strip the hydrogen layer, leaving the helium shell as the outermost part of the progenitor at the moment of explosion. When there is no hydrogen visible in the spectrum, the SN is of Type I. CCSNe that show no H lines but do have He lines in their spectra are called Type Ib SNe. In some extreme cases even the helium layer has been stripped away, resulting in SNe without H or He lines in their spectra, known as Type Ic SNe. Transitional objects, where a layer is mostly but not completely stripped, are also known (Alsabti & Murdin, 2017).

Photometric classification is cheaper as photometry is easier to obtain, but it cannot differentiate between subtypes that are defined on the presence, strength, or absence of specific elemental lines, leading to broader classifications with less certainty. Besides this, many of the differences that make certain SNe interesting are visible at early times, while a SN needs to evolve up to a certain point to have a light curve that can be used for classification. For these reasons, a common strategy is to use dedicated large-scale surveys such as the ZTF to find new SNe within days after explosion using photometry, and using different telescopes to follow these up with spectroscopy for classification.

1.2.2 Thermonuclear SNe

Besides CCSNe are the more commonly observed thermonuclear SNe (Type Ia SNe, SNe Ia), which are exploding WDs after fusion in their core has been ignited through some mechanism (see section 1.4). The more massive the WD, the easier it is to ignite as the nuclei are forced closer together. Since the WD is degenerate, the energy released by a single fusion event can completely be used to heat up the surrounding material and let it fuse as well, leading to a runaway process where the entire WD tries to fuse to iron-peak elements on a timescale of seconds. This process disrupts the entire star, expelling the material outward at high velocities (Hoyle & Fowler, 1960). While not all material in the WD is fused, most of the mass ends up in ^{56}Ni , with typical SNe Ia producing between $0.3 - 0.8 M_{\odot}$ of ^{56}Ni (Stritzinger et al., 2006). This is an unstable nucleus and decays as

$$^{56}\text{Ni} \xrightarrow{6.1 \text{ d}} ^{56}\text{Co} \xrightarrow{77.3 \text{ d}} ^{56}\text{Fe}, \quad (1.3)$$

with ^{56}Fe being the stable end product of the decay chain. Both decays are β^+ decays and release a positron, an electron neutrino, and radiation. The positron quickly annihilates with an electron in the ejecta and creates more radiation. As the ejecta are dense, this radiation cannot escape and heats the ejecta instead. This provides the energy that powers the broadband light of the SN.




Figure 1.3: ZTF *gri* light curve of a SN Ia in the rest frame and absolute magnitude, corrected for Milky Way extinction but not for host extinction. The *g*-band clearly shows a change in the decline rate as ^{56}Co decay becomes dominant over ^{56}Ni decay (as shown by the broken black line). The *r*- and *i*-band show the second peak.

The key spectral signatures of a SN Ia are the absence of hydrogen lines and broad Si II absorption features in their spectra around peak brightness. At first, the ejecta are dense and opaque, and only photons from the outer layers can escape. As the ejecta expand, they become more transparent, and the slower moving inner layers become visible as the photosphere recedes inwards through the layers. Therefore, by taking spectra at different phases, a map of the elemental abundances in velocity space can be built. As different models of the progenitor system and explosion mechanisms predict different compositions, this is a popular way to compare models to observations (e.g. [Dilday et al., 2012](#); [De et al., 2019](#); [Dutta et al., 2022](#); [Harvey et al., 2023](#)).

The SN reaches peak brightness around three weeks after the explosion and starts to fade again. The main power source that governs the rate at which the SN fades is the decaying ^{56}Ni , resulting in a linear decline of the light curve in magnitude space. In the near-infrared there is a second peak a few weeks after the first, which is suggested to be the result of a decrease in opacity due to the changing ionization state of iron group elements ([Kasen, 2006](#)). Figure 1.3 shows an example of a well-observed normal SN Ia.

1.2.2.1 Standardizable candles

Most SNe Ia peak around an absolute B-band magnitude of -19.3 mag, making them brighter than CCSNe. Some SNe Ia become brighter than others, but they also evolve slower. This is called the Phillips relation (Phillips, 1993), and it can be used to infer the absolute peak brightness based on how fast it evolves. Tripp (1998) improved the standardization by adding a second term: the so-called colour luminosity term. This accounts for the intrinsic SNe Ia colour, as bluer SNe Ia are brighter. These days a third term is included, which accounts for environmental dependencies and is implemented as a step function (Kelly et al., 2010; Sullivan et al., 2010). The resulting standardization formula is given by

$$\mu_{\text{obs}} = m_B - M_0 + \alpha x_1 - \beta c - \gamma p, \quad (1.4)$$

where μ_{obs} is the observed distance modulus. m_B and M_0 are the apparent and absolute B-band magnitude, respectively. x_1 is the SN stretch, c its colour, and α and β are the coefficients used to correct for these effects. γ is the size of the step function, and p determines when this step needs to be applied.

Using the standardisable nature of SNe Ia allows their distance to be measured, as well as the distance between Earth and the host galaxies they reside in. At the same time, the redshift of the host galaxy can be measured from the spectra. With these values together, the Hubble constant H_0 , which measures the expansion rate of the universe, can be measured. Di Valentino et al. (2021) discussed the various measurement of H_0 that have been done, both using SNe Ia and through other methods. Currently, one of the best measurements using SNe Ia was performed by the SH0ES team (Riess et al., 2022), who used the Pantheon+ sample (Scolnic et al., 2022) to measure the expansion rate of the universe and found it to be $H_0 = 73.30 \pm 1.04 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

1.2.3 Other types of transients

There are other types of transients besides the SNe mentioned above. Some SNe become much brighter than normal events and are therefore called superluminous SNe (SLSNe). Supermassive black holes (SMBHs) in the centre of galaxies can cause a variety of transient events such as tidal disruption events (TDEs) or ambiguous nuclear transients (ANTs), or a galaxy could have an active galactic nucleus (AGN). Some types of stars also pulsate, and vary in brightness over time. All of these can show up as transients in surveys looking for changes in the night sky on a day-to-day basis.

1.2.3.1 SLSN - Superluminous Supernova

SLSNe are SNe that have a higher peak luminosity than normal SNe, in some cases reaching over 5 magnitudes brighter than normal events at their peak ([Gal-Yam, 2019](#)). Like normal SNe, they can be classified into Type I and Type II based on the absence or presence of hydrogen emission lines in their spectra. One of the main puzzles regarding these types of events is the source that powers their superluminous nature. Different types of models have been put forward, using magnetars ([Maeda et al., 2007](#)), black holes ([Dexter & Kasen, 2013](#)), radioactivity ([Kasen et al., 2011](#)), or interaction with circumstellar material (CSM) ([Yan et al., 2017](#)) to push the brightness significantly above that of regular SNe. Each of these has different strengths and weaknesses.

1.2.3.2 AGN - Active Galactic Nucleus

An AGN is the central region of a galaxy in which a SMBH accretes matter at a high rate, releasing a part of the potential energy as kinetic energy and radiation. As a result of this, the AGN has a very high luminosity compared to the rest of the galaxy, with different orientations of the system resulting in different types of AGN ([Antonucci, 1993](#); [Urry & Padovani, 1995](#)). As the stream of matter being accreted is inhomogeneous, the accretion rate and luminosity change over time. In some cases AGN variability is so extreme that it changes the entire spectrum of the AGN, with

(dis)appearing broad emission lines and continuum flux. These are so-called changing-look AGN (CL-AGN) (see [Ricci & Trakhtenbrot, 2023](#), for a review). AGN variability can be used to estimate the SMBH mass through reverberation mapping ([Wandel et al., 1999](#); [Peterson & Horne, 2004](#)). Fast moving clouds close to the AGN change the strength of line emission based on how strongly they are illuminated by the AGN. The delay between a luminosity change in the AGN continuum and the broad emission line features coming from these clouds gives a measure of their distance, while the width of the emission lines gives a measure of their velocity. From these measurements the SMBH mass can be estimated using Kepler's third law.

1.2.3.3 TDE - Tidal Disruption Event

In some cases a star can get too close to the SMBH and get tidally disrupted ([Rees, 1988](#); [Strubbe & Quataert, 2009](#)). In a TDE the star gets ripped apart due to the strong tidal force caused by the gravity of the SMBH. The debris forms an accretion disk of hot gas that gets accreted onto the SMBH. The distance at which the TDE occurs depends largely on the SMBH mass, and for SMBHs more massive than $\sim 10^8 M_\odot$ the tidal disruption radius lies within the event horizon, meaning that TDEs cannot be observed around the most massive SMBHs ([Hills, 1975](#)). The tidal radius also depends on the compactness of the star. Compact objects such as WDs are very difficult to disrupt and thus have a smaller tidal radius. A red giant, on the other hand, has outer layers that are very loosely bound to the star, and the star is therefore easier to disrupt at a larger radius, though this might also result in a partial tidal disruption where only the outer layers are stripped while the stellar core survives the encounter. TDEs can become brighter than SNe, with the brightest event ever observed being AT 2021lwx (ZTF20abrbeie, nicknamed "Scary Barbie"), which is thought to be the disruption and accretion of a giant molecular cloud ([Subrayan et al., 2023](#); [Wiseman et al., 2023](#)).

1.2.3.4 ANT - Ambiguous Nuclear Transient

Some nuclear variability does not quite fit the known classes of AGN or tidal TDEs, showing characteristics of both types of events. These events have been named ambiguous nuclear transients (ANTS; [Kankare et al., 2017](#); [Hinkle et al., 2022](#); [Hinkle, 2024](#); [Hinkle et al., 2024](#); [Wiseman et al., 2025](#)) and rise up quickly, reaching peak brightness within a few weeks before declining very slowly over hundreds of days, although their decline rates are variable.

1.2.3.5 Variable stars

Some stars show variability as well, with the brightness of certain types of stars oscillating by a magnitude or more. The most famous class of variables are Cepheids, whose regular oscillation period of the order of days is related to their absolute magnitude. They are bright enough to be visible in nearby galaxies, which makes them useful to calibrate SN Ia distance measurements when building the distance ladder ([Gibson et al., 2000](#); [Saha et al., 2006](#)). Other types of stars, such as Mira variables, are much more extreme in their change of brightness and have periods > 100 days ([Bedding & Zijlstra, 1998](#)).

1.3 SN Ia subclasses

As has been stated in section [1.2.2.1](#), most SNe Ia explosions evolve in a very similar way with a tight relation between their peak, colour, and rate at which they evolve. These are called normal SNe Ia, or SNe Ia-norm. This name suggests the existence of abnormal SNe Ia besides the normal ones. This is indeed the case, in fact there is an entire zoo of subclasses that have photometric and/or spectroscopic differences from SNe Ia-norm. Figure [1.4](#) (taken from [Taubenberger, 2017](#)) shows the different SN Ia subclasses together by plotting their peak absolute B-band brightness as a function of $\Delta m_{15}(B)$, which is the difference in brightness in the B-band between the SN at peak brightness and 15 days after the peak in the SN rest frame expressed in magnitudes.




Figure 1.4: Subclasses of SNe Ia. The peak absolute B-band magnitude is plotted as a function of Δm_{15} in the same band. Normal SNe Ia lie along the Phillips relation, which is shown with a black line. Most of the non-normal subclasses can be separated from normal SNe Ia based on these two values alone. The main exception are the 91T-like SNe, which form a subclass based on their spectral differences from normal SNe Ia. This figure was taken from [Taubenberger \(2017\)](#).

Recently, [Dimitriadis et al. \(2025\)](#) made a similar plot but in the g -band, using the carefully curated sample of SNe Ia from the Zwicky Transient Facility’s second data release (ZTF SN Ia DR2, [Rigault et al., 2025a](#), [Smith et al., in prep.](#)). They show that $\sim 75\%$ of all SNe Ia are normal events that can be standardized and used for distance measurements. After that, the two largest subclasses are the SNe Ia-91T ($\sim 12\%$) and SNe Ia-91bg ($\sim 6\%$), leaving $\sim 7\%$ as the combined total of various other subclasses.

Different subclasses have a preference for different environments. By studying these, the characteristic properties of their light, and the connection to their environments, we can learn more about the properties of their progenitor systems and the exact mechanisms that are involved in SN Ia explosions.

1.3.1 Photometric differences

The Phillips relation is easily visible in Fig. 1.4, and SNe Ia-norm lie in a narrow strip around it. Most subclasses can be distinguished from SNe Ia-norm by photometry alone. The (aptly named) fast decliners fade unusually quickly after reaching peak brightness, while SNe Ia-CSM barely fade at all in the first few weeks after reaching their peak brightness. SNe Iax never reach the peak brightness of a normal SN Ia, which in some models is explained through the explosion failing to fully disrupt the WD ([Jordan et al., 2012](#); [Kromer et al., 2013](#)). Super-Chandrasekhar SNe Ia on the other hand are brighter than what would be expected from a M_{Ch} WD explosion. Many subclasses are named after a prototypical SN that defines the subclass.

1.3.2 Spectroscopic differences

Besides photometric differences, most subclasses are also spectroscopically different from SNe Ia-norm. Some have stronger emission lines of particular elements, others have weaker ones. Ca-rich transients (e.g. [Perets et al., 2010](#); [Kasliwal et al., 2012](#); [Frohmaier et al., 2018](#)) show Ca II emission lines in their nebular spectra that are much stronger than for other SN Ia types. The spectra of super-Chandrasekhar SNe

show that their ejecta move considerably slower compared to other subclasses, which, together with their high luminosity, is consistent with a double degenerate merger scenario (Howell et al., 2006; Hicken et al., 2007; Yamanaka et al., 2009).

1.3.3 The SN Ia-91T subclass

This subclass lies at the bright end of the Phillips relation, and thus cannot be identified photometrically but can therefore be normalized and used for distance measurements. However, spectroscopy reveals that this subclass evolves differently as it rises towards peak brightness, showing a blue pseudo-continuum with two strong Fe III absorption multiplets instead of the intermediate-mass element lines which are seen in normal SNe Ia. Si II can be detected around peak, but remains rather weak compared to normal SNe Ia. After the SN has reached its peak, the spectra look the same as normal SNe Ia, making it impossible to distinguish SN Ia-91T from normal event when using post-peak spectra.

SNe Ia-91T have been shown to have a preference for younger stellar populations, as they predominantly explode in late-type galaxies (Filippenko et al. 1992; Ruiz-Lapuente et al. 1992, although this preference is not seen in Dimitriadis et al. 2025). It is the largest SN Ia subclass after the normal SNe Ia (which account for $\sim 75\%$ of events), containing $\sim 12\%$ of all SNe Ia (Dimitriadis et al., 2025). Note that $\sim 40\%$ of the objects included in the SN Ia-91T subclass are in fact SNe Ia-99aa (Dimitriadis et al., 2025), a transitional group of objects that bridge the gap between normal SNe Ia and SNe Ia-91T (Taubenberger, 2017).

The existence of a transitional subclass favours an explanation with a single explosion mechanism with variable parameters. Calculations agree that explosion models need a significant amount (a few percent) of ^{56}Ni in the outermost layers to reproduce the spectral peculiarities of SNe Ia-91T (see e.g. Sasdelli et al., 2014). This means it has to be created in the outer layers or mixed outward. One proposed model is a gravitationally confined detonation (Fisher & Jumper, 2015), where a strongly off-center deflagration rises up as a buoyancy-driven plume to deposit some ashes in the

outer layers. The WD is not able to expand much before the deflagration turns into a detonation, resulting in a large amount of ^{56}Ni that will power a bright event.

1.4 SN Ia progenitor systems

WDs are stable objects on their own, but they can be made to explode under certain circumstances. Not all stars are single, and WDs in a binary system can interact with the other star and start accreting matter coming from the so-called donor star. As the radius of a WD is inversely proportional to its mass (Equation 1.1), it will shrink as it becomes more massive until the point is reached where fusion is ignited in the core of the WD, triggering its explosion.

Assuming that the stars have been in the binary system since formation, and since more massive stars have a shorter lifespan, the more massive star has turned into the WD. Therefore, the donor has to be $< 8 M_{\odot}$ (though due to stellar and especially binary evolution the mass of the donor has likely changed since formation). The donor could still be burning hydrogen or have evolved to the later stages in its life, up to and including becoming a WD as well. The many possibilities for the donor star is called the progenitor problem, which is currently still being debated. Different progenitor scenarios can broadly be put into one of two categories based on whether the donor star is a compact object (double degenerate) or not (single degenerate). Artistic impressions of these two scenarios are shown in Fig. 1.5.

1.4.1 Single degenerate

The most straightforward and classical way to explode a WD is by adding mass to a near- M_{Ch} WD until carbon is ignited in the core due to compressional heating. The companion star fills its Roche lobe, and material from the outer layers is syphoned onto the WD. The donor can fill its Roche lobe while it is still a main sequence star, when it has evolved into a red giant, or it may even be a helium star whose hydrogen layer has been stripped due to binary interactions (Whelan & Iben, 1973; Nomoto, 1982).




Figure 1.5: Left: Artistic impression of a single degenerate system, where the white dwarf accretes material coming from the surface of the donor star. Image credit: STFC / David Hardy. **Right:** Artistic impression of a double degenerate system where two white dwarfs spiral in to each other while releasing gravitational waves. Image credit: NASA / Tod Strohmayer (GSFC) / Dana Berry (Chandra X-Ray Observatory).

In delayed-detonation models the WD first expands during a deflagration phase, which then transitions into a detonation after the WD becomes unbound ([Khokhlov, 1991](#); [Mazzali et al., 2007](#)).

The homogeneous nature of normal SNe Ia could be explained by this scenario as the WD explodes in a narrow mass range, but the required accretion rate is in a narrow range of $\mathcal{O}(10^{-8}) \text{ M}_{\odot} \text{ yr}^{-1}$ as well (e.g. [Piersanti et al., 2014](#); [Wang, 2018](#)). This makes it difficult to explain the observed SN Ia rate. Other challenges for this scenario include the absence of a surviving companion in SN remnants (e.g [Schaefer & Pagnotta, 2012](#); [Ruiz-Lapuente et al., 2018](#)) and the absence of swept up H or He in most objects ([Leonard, 2007](#); [Tucker et al., 2020](#)).

It is also possible that the WD and red giant companion enter a common envelope phase. The WD merges with the hot core of the red giant, forming a rapidly rotating, WD-like degenerate core with $M_{\text{core}} > M_{\text{Ch}}$ inside the red giant. The angular momentum prevents the core from collapsing in on itself but it is gradually lost until it can no longer support itself, triggering the SN Ia explosion in the so-called core degenerate scenario ([Kashi & Soker, 2011](#)).

1.4.2 Double degenerate

If both objects are WDs, the primary (more massive) WD can still accrete material from the secondary WD in a stable Roche lobe overflow scenario (Piersanti et al., 2003a,b). In non-classical models, a sub- M_{Ch} WD can explode by first slowly accreting and building up a layer of helium on its surface, which explodes after a sufficient amount of material is gathered. As this explosion wraps around the outside of the WD, it sends a shockwave into the interior of the star, which culminates near the centre and temporarily increases the local pressure. If this is enough to ignite carbon fusion, the second detonation is triggered, which disrupts the star and results in the main SN (Taam, 1980; Livne & Arnett, 1995; Shen & Bildsten, 2009; Fink et al., 2010).

In other scenarios the WDs fully merge, either dynamically or violently. In the dynamical scenario the WDs lose angular momentum as it is radiated away through gravitational waves (Iben & Tutukov, 1984; Webbink, 1984). An artist's impression of this scenario is given on the right side of Fig. 1.5. As the WDs come closer, the less massive companion fills its Roche lobe, and material starts flowing onto the primary. As the companion loses mass, its radius increases, leading to more mass loss until the entire star is disrupted. Around half of the material forms a disk around the surviving WD while the rest falls directly onto its surface. Very little material is expected to be flung out of the system. As the system evolves, it eventually explodes in a SN Ia.

These scenarios have an explanation for the absence of a surviving companion in normal SNe Ia, and also explain the observed lack of hydrogen-rich companion material in nebular spectra. However, these models allow for a wide range of explosion masses and configurations, making it difficult to explain why normal SNe Ia behave so similarly (Liu et al., 2023b).

In collisions or violent mergers of two WDs a detonation can occur during the merger at the location of the accretion stream due to its high density and temperature (Rosswog et al., 2009; Pakmor et al., 2010, 2012). This might either directly start carbon fusion at the ignition site or cause a surface explosion that again wraps around

the WD and compresses its interior causing a second ignition, depending on the system and WD masses. The asymmetry of this system when it explodes is expected to cause significant asymmetry in the ejecta and its composition in different directions. This is expected to lead to different amounts of polarization of the SN light, depending on the angle between the plane of rotation and the line of sight (Wang et al., 2007; Bulla et al., 2016).

1.5 SNe with circumstellar material

The progenitor system of a SN is not completely isolated, and a significant amount of material can end up close to the explosion site before the explosion occurs. This is called circumstellar material (CSM), and since it usually originated from the progenitor system, CSM can give new insights into the history of the progenitor system and the sequence of events that led it to explode.

CSM can be created by different mechanisms depending on the progenitor system, and for SNe Ia its composition depends heavily on the type of donor star present. In the single degenerate scenario, CSM rich in hydrogen can be created by a WD generating a fast wind, blowing away a part of the material it received from the mass transfer (Nomoto et al., 2005). Alternatively, accretion in the single degenerate scenario can lead to unstable burning of the accreted material, resulting in large amplitude flashes called novae. Such events could expel a part of the outer layer which ends up as a detached CSM shell around the progenitor system (e.g. Dilday et al., 2012; Wang, 2018). In the double-degenerate scenario, part of the tidally disrupted secondary WD becomes unbound from the system, creating (H-poor) CSM in a disk around the primary WD, which may be able to produce detectable signatures depending on the time between the tidal disruption and the SN Ia explosion (Raskin & Kasen, 2013). The core degenerate scenario could provide a massive amount of hydrogen-rich material near the explosion site (e.g. for SN 2014J, Soker, 2015), though the signatures of this scenario could be ambiguous with interacting SNe IIn, leading to confusion (Inserra et al., 2016).




Figure 1.6: Schematic view of a SN Ia-CSM event. **Top:** The progenitor system containing a WD feeding on material from the donor star. Through various mechanisms a fraction of this material can be lost from the system and end up as CSM in a disk or shell at a distance d around the progenitor system. The size and extent of the CSM can vary depending on how it is created, this figure is not to scale. **Middle:** The WD explodes as a SN Ia. This can be a SN Ia-normal or a peculiar object, depending on the progenitor system. The ejecta start travelling outwards from the explosion site. **Bottom:** The ejecta catch up to the CSM and start interacting with it, causing the CSM to start emitting light as well as it is swept up. This results in a slowing down of the photometric light curve decline and the emergence of narrow emission lines in the spectra.

Tsebrenko & Soker (2015) proposed that the core degenerate and double degenerate scenarios could lead to SNe Ia exploding inside planetary nebulae, possibly resulting in observable time-varying Na ID absorption lines or CSM interaction.

Figure 1.6 shows a schematic view of the creation of a CSM shell (or ring depending on how mass is lost from the SN Ia progenitor system) and the subsequent interaction between the CSM and SN ejecta. When the SN explodes the ejecta will initially expand within the gap between the progenitor system and the CSM, and its spectrum will look like that of a non-interacting class of SN Ia that may or may not display absorption features caused by the CSM if it is in the line of sight. Compared to the ejecta, which are travelling at v_{ej} , the CSM is practically stationary but at a distance d from the explosion site. This means that the ejecta need a time $t = d/v_{ej}$ to reach the CSM and start to interact with it. Only then does the event transform into a SN Ia-CSM with new narrow emission lines. This later transformation explains why in some cases objects are reclassified as Ia-CSM after first being classified as something else (e.g. SN 2020aekp, which was first classified as a SN Ia by Dahiwalé & Fremling 2021 but reclassified as a SN Ia-CSM by Perley et al. 2021).

1.5.1 Interaction between SN ejecta and CSM

Light from the initial explosion is able to (partially) ionize the CSM. Depending on the distance of the CSM, it may still have a considerable ionization fraction by the time the ejecta reach the CSM. Additionally, some radiation from radioactive decays in the ejecta can re-ionize a part of the CSM, even when the CSM is optically thin.

As both the CSM and ejecta have a very low density, the interaction is mostly collisionless. However, as the ejecta reach the CSM, the complex magnetic field in the ejecta interacts with the ionized part of the CSM and accelerate it, creating a forward shockwave. The required energy comes from the ejecta which is slowed down, creating a reverse shock. Charge exchange interaction and the different velocities between the neutral and ionized parts of the CSM causes the neutral material to be swept up as well. A part of the accelerated neutral material can cross the shock boundary again




Figure 1.7: Schematic view of the interaction between the SN ejecta and CSM. The explosion site is below the picture and the ejecta are traveling up when it encounters CSM. At the top of the picture is the unshocked CSM, which has not yet felt the impact of the ejecta (other than the shock precursor). Behind the forward shock is the shocked CSM, followed by the shocked ejecta. A cold dense shell forms at their boundary where the material piles up, mixes, and radiates through line emission. Finally, behind the reverse shock is the freely expanding unshocked ejecta which has yet to feel the impact with the CSM. This figure is adapted from [Smith et al. \(2008\)](#).

and have charge exchange interactions with the ions there and create a shock precursor in the unshocked CSM ([Marcowith et al., 2016](#)).

The CSM interaction creates four different regions. Starting from the outside these are: unshocked CSM, shocked CSM, shocked ejecta, and unshocked ejecta (see figure 1.7, which is adapted from [Smith et al., 2008](#)). The unshocked and shocked CSM and ejecta are separated by the forward and reverse shock, respectively. At the boundary between the shocked CSM and ejecta material piles up and cools, forming a cold dense shell which radiates the narrow line emission that is characteristic of CSM interaction ([Smith, 2017](#)). Rayleigh-Taylor instabilities form at the boundary between the shocked CSM and ejecta as they generally have different densities, leading to mixing of the material. The interaction at the shock fronts produces X-ray and UV radiation which reheats the cold dense shell, allowing it to sustain strong line emission for a long time ([Chevalier & Fransson, 2017](#)). CSM interaction is an extra power source that brightens the SN, and can be strong enough to become the dominant light source and significantly slow the evolution of the SN light curve or even stall it completely for a period of time.

1.5.2 Other CSM signatures

Another way in which the presence of CSM has been inferred is through the presence of narrow, time-varying absorption lines in the SN spectra, indicating that they evolve alongside the SN (e.g., [Patat et al., 2007](#); [Sternberg et al., 2011](#)). However, these absorption lines are not always seen, and their behaviour varies from object to object, making it difficult to interpret them consistently. Light echoes can also provide an opportunity to study circumstellar and interstellar material, though the difficulty to geometrically separate them from the SN location and their weak nature makes this only possible for SNe in the Milky Way or nearby galaxies ([Patat, 2005](#); [Krause et al., 2008](#); [Gaur et al., 2016](#)).

1.5.3 The SN Ia-CSM subclass

SNe Ia-02ic, also known as SNe Ia-CSM, are the subclass of thermonuclear explosions that show signs of CSM interaction. This often shows up as a peculiar long-lived light curve along with narrow Balmer emission lines, which is a clear indicator of this signal coming from CSM as SNe Ia are expected to be hydrogen-poor. The prototypical event, SN 2002ic, showed narrow $H\alpha$ and $H\beta$ lines, and the slowly declining light curve was found to be consistent with $\sim 1.3 M_{\odot}$ of CSM ([Hamuy et al., 2003a,b](#); [Wood-Vasey et al., 2004](#); [Nomoto et al., 2005](#)).

Currently, there are several dozen known SNe Ia-CSM ([Aldering et al., 2006](#); [Silverman et al., 2013](#); [Sharma et al., 2023](#)). In all cases the interaction started within two months after the explosion, leading to their (re-)classification after spectroscopic confirmation. The amount and duration of interaction varies quite significantly from one object to another. While some members such as SN 2018gkx and SN 2020xtq mainly feature an unusually slow decline rate, other objects such as SN 2020aekp have a plateau for several 100 days before starting to fade away slowly ([Sharma et al., 2023](#)).

SN 2011km (PTF11kx) is a well-observed SN Ia-CSM event that shows signs of a complex CSM consisting of multiple shells with which the SN ejecta interact. [Dilday](#)

[et al. \(2012\)](#) show that its photometry is similar to 91T-like objects before the CSM interaction begins, and explain SN 2011km using a symbiotic nova progenitor system. [Leloudas et al. \(2015\)](#) also suggested a link between SNe Ia-91T and SNe Ia-CSM as they show similarities in their peak brightness and spectroscopy before the start of the interaction. In some cases, a 91T-like SN could start interacting with CSM hundreds of years after the explosion, as has been suggested for Kepler’s SN ([Patnaude et al., 2012](#); [Katsuda et al., 2015](#)).

Another interesting SN Ia-CSM is SN 2020eyj, which [Kool et al. \(2023\)](#) present as the first detection of a SN Ia-CSM interacting with He-rich material. Up until then all members of the subclass had shown strong $H\alpha$ emission and weak He signatures. SN 2020eyj, however, showed little to no H present in the CSM. This suggests that the progenitor system contained a He star. This was also the first time a SN Ia was detected in the radio. Non-detections in normal SNe Ia suggest a clean environment for the ejecta to expand in, while in this case there is a lot of material present.

1.5.4 CSM in CCSNe

Some CCSNe show interaction with CSM as well. The events are subclassified as SNe IIn, SNe Ibn, or SNe Icn, with the ‘n’ for narrow emission lines. As stated in section [1.2.1](#), the difference between these three main types of SNe is the amount of material that has been stripped prior to explosion. Mass can be lost through binary interaction stripping away the outer layers of the star prior to explosion, which can end up in a disk in the orbital plane of the system ([Solar et al., 2024](#)). Massive stars have also been known to have strong stellar outflows and winds that could form more spherical structures. The existence of Wolf-Rayet stars shows that very massive stars can lose their entire hydrogen envelope, and they are thought to be one of the very last stages in the life of the star before it explodes in a SN Ib or SN Ic ([Groh et al., 2013](#); [Gal-Yam et al., 2022](#)). Current models prefer the binary scenario due to the low amount of confirmed disappearances of progenitor stars ([Van Dyk et al., 2023](#)).

The CSM often resides close to the progenitor, causing signs of interaction to be present almost immediately after the explosion. Examples of the interaction starting much later have however been found as well, suggesting that in these cases the CSM has been ejected a long time ago (Miller et al., 2010; Fox et al., 2013). The CSM can also be thick enough to obscure features of the underlying SN, depending on the geometry and orientation of the system (Chugai et al., 2004; Smith et al., 2015). The short distance between the SN and CSM suggests that the material was ejected recently, and in some cases precursor events have been observed prior to the final SN explosion (e.g. SN 2011ht, Fraser et al. 2013b and SN 2020pzb, Elias-Rosa et al. 2024).

Recently, SN 2021yfr was discovered to be a new type of SN and given the classification of SN Icn (Gal-Yam et al., 2024). This object showed highly ionized silicon, sulphur, and argon lines but barely any carbon, oxygen, helium, or hydrogen. This suggests the progenitor was a massive star stripped all the way down to its final layer before the iron core. The classification leaves space for the yet-to-be-discovered explosion of a slightly less stripped star in a SN Type Icn, which would have no carbon but strong oxygen, neon, and magnesium emission lines, as well as their non-interacting counterparts and transitional objects. (Schulze et al., 2024).

1.5.5 Distant CSM

As stated in section 1.5 and shown in figure 1.6, there is a delay between the SN explosion and the start of CSM interaction due to the distance between the progenitor system and CSM. All objects that are currently classified as SNe Ia-CSM on the Transient Name Server (TNS)¹ have shown signs of interaction at or around peak brightness. Assuming an ejecta velocity of $v_{\text{ej}} \sim 20\,000 \text{ km s}^{-1}$ and that the interaction starts on average around the SN peak ~ 3 weeks after the explosion, the CSM is located at a distance $d \sim 3.5 \times 10^{15} \text{ cm}$ of the progenitor system (Sharma et al., 2023). If however the CSM is at a distance of $d \sim 10^{17} \text{ cm}$, ejecta with the same v_{ej} would need ~ 1.5 years to catch up and start interacting.

¹<https://www.wis-tns.org/>

In the effort to systematically search for such late-time CSM interaction, [Graham et al. \(2019b\)](#) looked at old (≥ 1 year) SNe using the *Hubble Space Telescope* (HST). They focused their search on subclasses that are associated with CSM interaction, such as SNe Ia-91T ([Leloudas et al., 2015](#)). Out of 72 targets, only ASASSN-15og and SN 2015cp were found to show late-time CSM interaction. ASASSN-15og is a SN IIn with detected CSM interaction around its peak, and was used as a control object. SN 2015cp has been classified as a SN Ia-91T without signs of CSM interaction around its peak. This showed that late-time CSM interaction may be systematically missed due to SNe Ia not being actively followed at these phases. From a progenitor point of view, this means that material can be ejected from the system prior to the explosion, giving it time to travel further before being caught up by the SN ejecta.

[Dubay et al. \(2022\)](#) used archival UV-band data from the *Galactic Evolution Explorer* (GALEX) to look for late-time CSM interaction in SNe Ia. Out of a sample of 1080 SNe Ia, four were detected in the UV near peak, but none showed signs of late-time CSM interaction. They injected model light curves into their data to recover and constrain a range of CSM interactions differing in start time and maximum luminosity. Using this and the objects in [Graham et al. \(2019b\)](#), they show that this type of CSM interaction is rare, constraining the occurrence rate of late-time CSM interaction between 500 to 1000 d after the initial discovery of the SN in $< 5\%$ of the SNe Ia at a strength similar to SN 2015cp, and a decreasing percentage as the interaction gets stronger.

In theory, a SN showing late-time interaction could be interacting with a nearby object rather than material that used to belong to the donor star in the progenitor system, such as a nearby molecular cloud or if the explosion occurred inside a planetary nebula ([Tsebrenko & Soker, 2015](#)). The later the interaction starts, the further from the progenitor system it occurs, and the larger the chance of the ejecta encountering a nearby structure. WD progenitors have long lifespans, and there is an additional delay between the formation of the WD and the SN Ia explosion. This means that the system is likely far from the molecular cloud that birthed it and any nearby cloud

would be more likely a chance encounter, suggesting a busy environment such as a galactic nucleus. However, SN 2015cp is at a distance of ~ 11 kpc from its host, making it unlikely for a cloud to be within 10^{17} cm of the explosion site (Graham et al., 2019b). This shows that late-time interaction with distant CSM that originated from the progenitor system can be observed in certain cases.

1.6 In this thesis:...

I will present my work on the search for signs of late-time interaction in SNe Ia using observations taken by ZTF. CSM is connected to the mass loss history of the progenitor system, and more distant CSM has been ejected a longer time ago. By studying late-time CSM it is possible to learn more about the mass loss history of the progenitor system, and it can hold clues to the type of progenitor scenario that the system belonged to.

In chapter 2 I will introduce ZTF as well as the other facilities that have been used to gather data that is presented in this thesis. I will also give an overview of the basics of observing, some of the types of data that can be obtained, and the basic steps that are required to reduce them to something that can be used in further analysis.

Chapter 3 is an adaptation of Terwel et al. (2025), and presents my search for late-time signals in the ZTF SN Ia DR2, which consists of 3 628 spectroscopically classified SN Ia that were first discovered between March 2018 and October 2020. It also details the pipeline and analysis tools I developed to search in a consistent manner (Section 3.2).

It might be possible that a late-time signal is observed by ZTF while the main transient occurred before the survey started. In chapter 4 I search for this type of event using a slightly modified version of the pipeline used for the DR2 search. As the analysis is exactly the same whether the original transient was a SN Ia, SN II, or any other type of transient, I search through a sample of 7 718 transients that were first discovered between 1 January 2008 and 1 January 2018.

Up to this point my search has been through archival data only, with the obvious drawback that it is very likely that any recovered late-time signal cannot be followed up on. By further modifying the pipeline, I search in chapter 5 for late-time signals in SNe Ia in real-time using the latest ZTF observations with the goal to follow up photometrically and/or spectroscopically using other telescopes.

Finally, in chapter 6 I conclude on the results of searching through these three samples and finish with possible continuations of this search, modifications, and possible alternative usages for this or a similar pipeline.

To convert between apparent and absolute magnitude I assume a flat Λ CDM cosmology with $H_0 = 67.7 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.310$ (Planck Collaboration et al., 2020) throughout this thesis.

Observing in the optical regime

Astrophysicists face the challenge of not being able to set up and control their experiments. The universe is our laboratory, but all we can do is see or detect the results while often not knowing the exact setup of the experiment. Models are made to explain and predict the behaviour of planets, stars, and galaxies, but ultimately, observations are needed to compare against and test our models. My work relies heavily on observational data, and in this chapter I will introduce the different types of observations that are used throughout this thesis (section 2.1), telescopes and instruments that are used to obtain these observations (section 2.2), and a quick overview of how to calibrate the raw images and extract useful data (sections 2.3 and 2.4). I will also give a general overview of what to consider when planning observations in section 2.5.

2.1 Types of observations

All optical observations are, in essence, images taken by a camera. Light falls onto a pixel on the detector (a charge-coupled device, CCD) and frees some amount of electrons. If more light hits the pixel, more electrons are freed. At readout these electrons are counted per pixel, or group of pixels if binning is applied, and turned into a digital number called a count or analog-to-digital unit (ADU). The gain (e^- / ADU conversion rate) is a unique property of each readout channel, and sets the conversion rate between the measured number of electrons and the resulting value that is stored. One of the main advantages of CCDs is their low read noise levels (mean e^- added to a pixel upon readout), with modern CCDs having a read noise level of only a couple of electrons. Observations have contributions from different noise sources, but as long as the total count rate is in the linear regime of the CCD, there is a linear relation between the received flux and the final count. By using calibration images, it is then possible to calculate the observed flux (Howell, 2006). Section 2.3 explains the different types of calibration images, and section 2.4 explains their usage in image reduction.




Figure 2.1: Left: Image and partial spectrum of SN 2024nqr (left) and SN 2024pgd (right), two SNe Ia simultaneously active in the same galaxy. The image was taken without a filter and used to align the $1.0''$ slit (horizontal dashed lines) over both SNe. The resulting spectrum, taken with grism #4, shows three traces as white vertical stripes. The outer two line up with the two SNe, while the middle trace is from the edge of the host galaxy, which is in the slit (vertical dotted lines for guidance). The horizontal lines in the spectrum are sky lines coming from atmospheric emission, and the white spots in the spectrum are due to cosmic rays. This data was taken with NOT+ALFOSC on the night of 28 July 2024 while testing an experimental rapid response mode (RRM, credit: Samuel Grund Sørensen). **Right:** Compass showing the orientation of the left image. Note that North is not exactly straight up.

2.1.1 Imaging

Imaging is one of the simplest observing modes, as it is just taking a photo of a part of the sky. The top of figure 2.1 shows a raw photometric image taken with the *Alhambra Faint Object Spectrograph and Camera* (ALFOSC) on the Nordic Optical Telescope (NOT) without using a filter (See section 2.2.2 for more information about the NOT and ALFOSC). The images are monochromatic, meaning they only have a value for the light intensity received per pixel. To obtain colourful images multiple observations have to be made using different filters, which are then combined to represent different colours in the final image.

Faint objects can be observed by increasing the exposure time in a single image, or by stacking multiple images together to increase the effective exposure time. Stacking images can be useful for, e.g., reducing cosmic ray interference, avoiding overexposure of a bright source close to a fainter target, or constructing time series. When stacking images, it is common practice to dither the telescope: applying a small offset between exposures to ensure that the target hits a different part of the CCD to avoid issues with bad pixels ruining otherwise good observations. While this decreases the effective size of the fully stacked image, as long as the edges are not needed there is no issue (Zackay & Ofek, 2017a,b).

2.1.2 Spectroscopy

Spectroscopy goes one step beyond just taking a photo. In slit spectroscopy, instead of a filter to select a wavelength range to observe, now a slit restricts the observable region of the sky to a narrow band along one axis of the detector (e.g. horizontal). After the slit, the light hits a grating or grism (a grating and prism combined), which diffracts the light based on wavelength across the second axis of the detector (vertical). The rule density on the grating or grism dictates the wavelength spread of the light: the more rules per unit distance, the bigger the dispersion, and the higher the spectral resolution of the resulting image. The trade-off is that a smaller part of the spectrum

2. OBSERVING IN THE OPTICAL REGIME

can be observed at a time, and there is less light being received per pixel which reduces the SNR unless the exposure time is increased to account for this. Any point-like source that is observed becomes a line in the spectral direction, called a trace. Of course, atmospheric turbulence and a finite pixel size slightly smear point sources, resulting in the trace usually being a couple pixels wide in the spatial direction, depending on the weather conditions. Extended sources such as galaxies create extended traces.

There is some freedom in choosing the orientation of the slit. This is called the position angle of the slit. As rotating a single piece of the setup might cause alignment issues, often the entire instrument is rotated instead. This means that the orientation of the slit and grism with respect to the detector remains constant. Instead, the location of North changes on the resulting image (e.g. in figure 2.1 North is slightly to the right to allow both SNe to be in the slit simultaneously). Slits and grisms are then called "horizontal" or "vertical" based on the direction they work in with respect to the detector. For instance, a vertical slit allows light to fall along a couple of columns on the detector while blocking the rest. If a horizontal grism is placed after this slit, it will disperse the different wavelengths from left to right on the detector. Light from a point source going through this setup will end up as a horizontal trace on the resulting image. If there are multiple targets near each other, and they can be in the slit at the same time, the required position angle can be calculated from the two target positions. If there is a single target to be observed, the position angle can be anything, but usually the parallactic angle is chosen. In this orientation the slit is perpendicular to the horizon, and prevents losses from differential diffraction (different colours diffracting differently when entering the atmosphere at an angle, [Filippenko 1982](#)).

The trace of a target observed with the slit at the parallactic angle will be slightly diagonal on the detector. This is because differential diffraction occurs vertically, so if the slit is oriented in the same way to prevent losing flux at the edge of the spectrum the spectral direction is along the horizon. The target appears to be at slightly different altitudes when comparing its position in red and blue light, and these are dispersed to opposite sides along the spectral direction (which is orthogonal). If this difference

is, for example, equivalent to 2 pixels on the CCD, the resulting trace (assuming it is horizontal in this example), if centered on row n in the middle, will go from row $n - 1$ on one side to row $n + 1$ on the other side.

The bottom panel of figure 2.1 shows a section of the spectrum taken of the SNe in the top panel image. The two SNe are drawn out into vertical traces, and a third trace belonging to the edge of the host galaxy can be seen in the middle. The horizontal lines are sky emission lines, and while these can technically be used to estimate the conversion from pixel position to wavelength, standardized arc frames will result in a much better wavelength calibration (see section 2.4).

2.2 Telescopes

Most of the data used in this thesis comes from the Zwicky Transient Facility (ZTF), and follow-up observations have been made using the Nordic Optical Telescope (NOT), and the Gran Telescopio CANARIAS (GTC), which will be introduced below. Some additional data comes from other sources, which are listed for completeness. The same filter names (*gri*) are used for filters at different telescopes, which have slight differences. In the rest of this thesis I will use *gri* to refer to the ZTF filters, unless specified otherwise.

2.2.1 Zwicky Transient Facility

The Zwicky Transient Facility (ZTF, [Bellm et al. 2019a,b](#); [Graham et al. 2019a](#); [Masci et al. 2019](#); [Dekany et al. 2020](#)) is an optical large-sky survey observing the entire northern night sky above $\text{Dec} \approx -30^\circ$ every 2 to 3 nights in three broadband optical filters g ($\lambda_{eff} = 4746.48 \text{ \AA}$), r ($\lambda_{eff} = 6366.38 \text{ \AA}$), and i ($\lambda_{eff} = 7829.03 \text{ \AA}$), which are similar to the well-known SDSS *gri* filters. The efficiency of these filters is plotted as a function of wavelength in figure 2.2. The survey saw its first light in October 2017 and it formally began scientific operation in March 2018, and has been running continuously until the time of writing this thesis.




Figure 2.2: Throughput as a function of wavelength of the different filters used to gather the bulk of the data in this thesis. g filters are shown in green, r in orange, i in red, and the different telescopes are shown with different line styles (Continuous for ZTF, dashed for NOT, dot-dashed for GTC). The SDSS filter transmissions (dotted lines) are shown for comparison. The black dashed line is the WG345 order-blocking filter that was used for some spectroscopic observations taken with the NOT. For the grisms only the wavelength ranges are shown, not their efficiency at each wavelength.

The observations are made using the 48'' aperture Schmidt-type design Samuel Oschin Telescope, which is based at the Palomar Observatory in Southern California. Each exposure is 30 s long, can go a limiting magnitude of ~ 20.5 mag, and covers an area of ~ 47 deg² with a plate scale of $1.01''$ pixel⁻¹. The camera is divided into a 4×4 grid of CCDs, each of which has four readout channels called quadrants. This results in each observation producing 64 separate images, each with its own readout channel identifier (rcid). Similarly, the observed region of the sky is divided into different telescope pointings called fields to ensure that the same region of the sky is observed in the same way each time, aiding with the reduction of the data. This results in each combination of filter, field, and rcid being a set of observations of a particular part of the sky using a specific setup.

2.2.2 Nordic Optical Telescope

The Nordic Optical Telescope (NOT)¹ is a 2.56 m telescope located at Observatorio Roque de Los Muchacos in La Palma, Spain, at an elevation of 2382 m above sea level. It hosts several instruments for observing in the optical and near-infrared, both for imaging and spectroscopy. The *Alhambra Faint Object Spectrograph and Camera*

¹<https://not.iac.es>

(ALFOSC) was used to obtain the data used in this thesis. I will only discuss the parts relevant to this thesis, further details on this instrument and details on the other instruments can be found on the NOT website.

ALFOSC is a versatile instrument mounted in Cassegrain and can be used for imaging, spectroscopy, and (spectro)polarimetry. As there are several wheels equipped to hold a variety of optical elements, the instrument can switch quickly between different setups between observations. The images can cover up to $6.4' \times 6.4'$ per exposure, and the CCD has a plate scale of $0.2138'' \text{ pixel}^{-1}$. In this thesis g ($\lambda_{cen} = 4800 \text{ \AA}$), r ($\lambda_{cen} = 6180 \text{ \AA}$), and i ($\lambda_{cen} = 7710 \text{ \AA}$) are used for photometry. For spectroscopy grism #4 is used to split the light vertically, together with a horizontal $1.0''$ slit if the seeing was $\leq 1.3''$ or a horizontal $1.3''$ slit if the seeing was $\geq 1.3''$. Grism #4 has a dispersion of $3.3 \text{ \AA pixel}^{-1}$ and a wavelength range from $3200 \text{ \AA}$ to $9600 \text{ \AA}$, but as the response at short wavelengths is poor, the spectra used in this thesis are cut at $4000 \text{ \AA}$. For some spectra an order-blocking filter (WG345) is used to avoid second-order diffracted blue light overlapping with first-order diffracted red light on the detector. Figure 2.2 shows the filter transmission curves and wavelength range of the grism.

2.2.3 Gran Telescopio CANARIAS

The Gran Telescopio CANARIAS (GTC)¹ is a 10.4 m telescope at Observatorio Roque de los Muchachos in La Palma, Spain, and is the largest optical/near-infrared telescope on the island. Its primary mirror is made up of 36 hexagonal pieces, creating an effective collection area of 73 m^2 , ideal for observing very faint targets. The GTC can host up to six instruments at a time in various focal positions, allowing for a large variety of observations to be made. One of the most commonly used instruments is OSIRIS+, the upgraded version of OSIRIS: the *Optical System for Imaging and low-Intermediate-Resolution Integrated Spectroscopy*.

OSIRIS+ has an unvignetted field-of-view of $7.8' \times 7.8'$ and a standard plate scale of $0.254'' \text{ pixel}^{-1}$. Since the standard readout has 2×2 binning, the plate scale can be

¹<https://www.gtc.iac.es>

increased to $0.127'' \text{ pixel}^{-1}$ if so desired. Like ALFOSC, this instrument is also built to easily switch between different setups between observations. For photometry the r ($\lambda_{\text{cen}} = 6410 \text{ \AA}$) filter is used in this thesis, and for spectroscopy the R1000R grism with a $1.0''$ vertical slit is used. R1000R splits the light horizontally over the detector with a range of $5100 \text{ \AA}$ to $10000 \text{ \AA}$ with a dispersion of $2.62 \text{ \AA pixel}^{-1}$. These filter transmission curve and grism wavelength range are shown in figure 2.2.

2.2.4 Other observations

Small amounts of data coming from other telescopes and surveys are presented in this thesis as well. This includes a follow-up observation of SN 2019ldf in section 3.3.2.4 in g and r using the *ESO Faint Object Spectrograph and Camera version 2* (EFOSC2, Buzzoni et al. 1984) imaging spectrograph on the ESO New Technology Telescope (NTT) in La Silla, Chile as part of the extended Public ESO Spectroscopic Survey of Transient Objects+ (EPESSTO+, Smartt et al. 2015).

To complement ZTF data of several SNe, in chapter 4 optical photometry from the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS, Chambers et al. 2016), (intermediate) Palomar Transient Factory (PTF, Law et al. 2009; Rau et al. 2009, iPTF, Kulkarni 2013), All Sky Automated Survey for SuperNovae (ASASSN, Shappee et al. 2014; Jayasinghe et al. 2019), Asteroid Terrestrial-impact Last Alert System (ATLAS, Tonry et al. 2018), and *Global Astrometric Interferometer for Astrophysics* (Gaia, Gaia Collaboration et al. 2016) are used. In chapters 3, 4, and 5 near-infrared photometry from the *Wide-Field Infrared Survey Explorer* (WISE, Wright et al. 2010) is used as well.

For difference imaging, data from the DESI Legacy Imaging Surveys Data Release 9 (Dey et al., 2019) were used as references for the SN 2019ldf follow-up data, and data from Pan-STARRS and the Sloan Digital Sky Survey (SDSS, York et al. 2000; Adelman-McCarthy et al. 2006; Gunn et al. 2006; Smee et al. 2013) were used as reference images for follow-up observations taken with the NOT and GTC.

2.3 Calibration images

Before observations can be used for science, the images need to be calibrated. This is done using different types of calibration images, each of which measures and corrects for different effects of the telescope and detector. Table 4.1 from [Howell \(2006\)](#) has an overview of the different types of calibration images used for photometry. Usually calibrations are taken during the day or twilight in order to avoid losing valuable observing time during the night. It is standard practice to take multiple calibration images and use an odd number in the reduction to find median values and remove interference from e.g. cosmic rays.

2.3.1 Bias

The first type of calibration image is the bias, which is made by reading out the CCD without exposing. The resulting image contains the number of counts that will be in every exposure regardless of what has been observed or with what exposure time. In other words, measuring the bias can be thought of as measuring the offset to correct for in every other image, and the underlying noise level in each frame.

2.3.2 Dark

Any detector that is not at a temperature of 0 K will have some amount of noise due to thermal effects. This can free electrons in pixels over time, creating a dark current and increasing the noise over time. The effect can be measured by exposing for the same amount of time as the science images that were taken but without letting any light hit the CCD. This is called a dark frame.

As this is a thermal effect, it can be reduced to negligible amounts by cooling the instrument. This saves precious observing time, as otherwise dark frames would ideally have to be taken at the same temperature as the target was observed, which is easiest to do directly after the science exposure. By cooling the detector with e.g. liquid

nitrogen, this noise source can be avoided instead of having to correct for it, saving time and the amount of images that need to be taken in the process.

2.3.3 Flatfield

The amount of light that the telescope receives is converted into a digital number, but there is no guarantee that this conversion rate is the same for each pixel. This can be due to intrinsic differences between the pixels, or outside effects such as dust reducing the amount of light received on a part of the detector. To correct for this, an evenly illuminated field has to be observed, resulting in an image called a flat or flatfield. By ensuring that each pixel receives the same amount of light, the different counts will reflect the varying responses per pixel.

Any evenly illuminated object can be used for this, such as the inside of the telescope dome to create dome flats. A more perfect evenly lit source, however, is the sky, and by using this, sky flats can be taken. While it is usually too bright during the day causing the CCD to saturate even with the narrowest filter and shortest exposure time, there is a window during twilight where the sky is darker but not dark enough to observe stars yet, perfect for taking flats. As a general rule, narrowband filters need a brighter sky, and in the evening these need to be done before the broadband filters. After that, assuming similar efficiencies between filters, blue filters need brighter skies than red filters, forcing a specific order in which the sky flats need to be taken during the short window when this is possible. Of course, if flats are taken in the morning, the order has to be reversed.

2.3.4 Arc

In spectroscopy one of the axes has the lowest wavelength at one end and the highest wavelength at the other end of the image. To know where each wavelength falls on the detector, arc frames are needed. These are taken by observing a lamp filled with a known set of elements (e.g. He, Ne, or Th and Ar). The wavelengths of the emission

lines are known very precisely, and by matching these with the observed lines in the arc image a pixel-to-wavelength conversion can be found, called the wavelength solution.

Usually, arcs can be taken during the day when the telescope is idle. However, in some cases the mechanical flexure of the telescope, caused by it being in a different position during observing, can introduce an uncertainty in the wavelength calibration unless an arc is taken with the telescope in the same position as for the target. In these cases an arc is usually taken directly before or after the target is observed, or between exposures of the target (Howell, 2006).

2.4 Reduction

After all observations have been taken, it is time to analyse them. A thorough overview of this can be found in Berry & Burnell (2000). The first step is to reduce the raw data into the required format to work with. After that, additional analysis techniques can manipulate the reduced images directly or the data that has been extracted from them. The response function of a detector can be written as

$$R_{ij}(f, t, \lambda) = B_{ij} + D_{ij}(t) + F_{ij}(\lambda) \times f \times t, \quad (2.1)$$

where R_{ij} is the CCD response of pixel i, j as a function of the integrated flux of the target $f \times t$ during the exposure which lasted a time t . The goal is to measure the flux f , which requires knowing and correcting for the bias level B_{ij} , dark current $D_{ij}(t)$, and pixel response $F_{ij}(\lambda)$. Each type of calibration image is used to measure one of these values. Note that it is assumed that there are no cross or higher order terms in Eq. 2.1. In other words, the CCD is in its linear regime. When a pixel receives too much light and gets close to saturating, it is no longer in its linear regime and more terms appear in Eq. 2.1, making it much more difficult or even impossible to measure the observed flux.

2.4.1 Bias, dark, and flat corrections

Using the calibration images from section 2.3, the raw science images can be reduced to something a flux level can be measured from.

First, the master bias frame is created and subtracted from every other image. As both f and t are 0, the bias measures B_{ij} directly and can then immediately be removed.

Secondly, the dark frames measure $D_{ij}(t)$ for a specific t , but the master dark frame can only be used on science observations with the same exposure time. Alternatively, it is possible to subtract $B_{ij} + D_{ij}(t)$ in a single step by not separating out the bias term using bias images first.

Finally, every science image is divided through the normalized master flat to equalize the pixel responses. There is still a factor $F(\lambda)$ present as the detector efficiency is wavelength-dependent, but that value is now independent of the pixel position, allowing values from across the CCD to be compared.

2.4.2 Cosmic-ray removal and image stacking

At this point it is often good practice to run a cosmic-ray removal algorithm to find these and mask them by replacing affected pixels with the median of neighbouring pixels that are not affected by the cosmic ray. This can be done using, for instance, L.A.Cosmic (van Dokkum, 2001), which I used through Astro-SCRAPPY (McCully et al., 2018) when reducing follow-up photometry of several objects in this thesis.

L.A.Cosmic uses Laplacian edge detection to detect the sharp edges that are characteristic for cosmic rays as they tend to only affect the one or few pixels they hit and not their neighbours. Point sources and extended objects on the other hand vary much smoother from pixel to pixel. The algorithm convolves the image using a Laplacian kernel to find large changes between neighbouring pixels, and checks for symmetry in the identified pixels (A critically sampled or undersampled point source is notoriously difficult to distinguish from a cosmic ray, but should show some symmetry around its

centre while a cosmic ray would not). Identified cosmic rays are then replaced by the median value of its unaffected neighbours. Cosmic rays that affect a larger area can be removed iteratively as each iteration removes the outer edges of the affected area until the entire cosmic ray is removed. For more details about this algorithm I refer the reader to [van Dokkum \(2001\)](#) where it is introduced and thoroughly explained.

If multiple images are taken of the same field or object they can be stacked to reduce background noise and increase the SNR of the observed objects. Sometimes the observations have been taken with dithering to avoid the same objects being on the same pixels in every exposure, which means the images have to be shifted to align all objects between the images before stacking them. As cosmic rays are only expected to strike a particular position in a few images, they can be removed through sigma clipping, (detecting and removing outliers from the computation).

2.4.3 Source extraction and standard star calibration

Filters are never 100% transparent at any wavelength, and the CCD responds differently to different wavelengths as well. To correct for this, one last type of calibration image is used: the standard star. This was not mentioned in section [2.1](#) as observing a standard star is exactly the same as observing the actual science target. The only difference is that the expected result of the observation is known and can be used to correct for the wavelength-dependent efficiency of the instrument. Standard star calibration is usually done around the same time as the final extraction of the science target to ensure that both the standard star and science target have been treated exactly the same throughout the reduction process.

2.4.3.1 Photometry

The goal of photometry is to measure the received flux of a source (assumed here to be a point source such as a star for simplicity). In theory this is simply a counting exercise, where all that needs to be done is count the signal at the position of the source and subtract the contribution from the background. In practice this immediately




Figure 2.3: Schematic view of a point source on a CCD, with the pixel grid shown. The black circle is the source, the innermost circle shows the source aperture, and the grey region shows the annulus used to estimate the background.

gives an issue. Digital images are pixelated, with each pixel having a single number of counts. Assuming that a point source results in a circular area of the CCD being illuminated, this begs the question of how to deal with pixels that only partially cover the source (See figure 2.3). Generally there are three ways to deal with this: Ignore partially covered pixels, include them fully, or apply some weighting dependent on the estimated fraction of coverage. When this choice is made, the total signal can be estimated through PSF (point spread function) photometry or aperture photometry.

Source and background extraction

In PSF photometry a 2D profile function is fitted to the observed PSF of the source, which is then integrated over to estimate the total signal. Commonly used functions to model the PSF include the Gaussian $G(r)$, modified Lorentzian $L(r)$, and Moffat $M(r)$, which can be expressed as

$$G(r) \propto e^{\left(\frac{r^2}{2a^2}\right)}, \quad L(r) \propto \frac{1}{1 + (r^2/a^2)^b}, \quad M(r) \propto \frac{1}{(1 + r^2/a^2)^b}, \quad (2.2)$$

with r the distance from the centre of the point source and a and b fitting parameters. Assuming that the used function is a good representation of the data, it can be used to estimate the source signal at any point. If it is also assumed that all point sources should have the same PSF shape, then it is possible to fit the function on the brightest stars in the image and use the found PSF on fainter objects. One of the main advantages of PSF photometry lies in the fact that the model of a point source can be used to extract a single source in crowded fields where light from different sources are blended, such as a SN in a galaxy background (Stetson, 1992).

Aperture photometry is simply summing the pixel values in a defined area (which does not have to be circular). While computationally simple, this can lead to errors in regions with severe blending, as there is no way to separate the contributions from different sources. The size of the chosen aperture is also important. A large aperture, such as the one shown in figure 2.3, contains more of the source but also more of the background. As the point source is brightest at its centre, at some point more noise than signal is added, reducing the final S/N of the measurement. Therefore, while an aperture with radius $r = 3 \times \text{FWHM}$ includes $\sim 100\%$ of the source light (Merline & Howell, 1995), $r \approx 1 \times \text{FWHM}$ gives the highest S/N (Howell, 1989). Aperture corrections can then be applied to account for the missing light (Howell, 1989; Stetson, 1992).

Up until here the photometry of the source is the total amount of counts in the source area. This includes signal coming from the background, which must be estimated

and subtracted to get a good estimate of the source signal. The background level is commonly estimated by placing an annulus around the source (e.g. the grey area in figure 2.3) where no sources are present. By estimating the mean pixel value the mean background level can be estimated, with a good estimate generally requiring three times the amount of pixels as are contained in the source aperture (Merline & Howell, 1995). To do this robustly, outliers should be removed when calculating the mean to remove things like bad pixels and contamination from close by sources.

Flux calibration

After removing the background, the remaining counts f are that of the source and have the same units as flux. This means that they can be converted to an instrumental magnitude m_{inst} through

$$m_{\text{inst}} = -2.5 \log(f), \quad (2.3)$$

which only differs by a constant from the Vega or AB system magnitude m . This constant is generally referred to as the zero point or ZP, and can be estimated by using one or more standard stars. If the filter in which the observation is taken is very commonly used, such as the SDSS filters gri , any image is likely to contain some stars whose magnitude in these filters are known. These can then be used to calibrate against instead of using a standard star in a different observation. Generally speaking the ZP can be written as

$$\text{ZP} = m_{\text{inst}} - m + \alpha_c \times C + \alpha_a \times AM, \quad (2.4)$$

with C the colour of the source, AM the airmass of the observation, and α_c and α_a coefficients for these terms. To find the ZP, Eq. 2.4 can be fitted using the standard star(s). For more precise calculations higher order terms can be included, as well as a term accounting for possible time variability of ZP (as has for instance been done in Ofek et al., 2012).

Uncertainty estimation

No measurement is meaningful without an estimation of the associated uncertainty. In photometry, the uncertainty comes from a few different sources. The uncertainty on a signal $N_\star$ is its square root. The background adds uncertainty N_S for each pixel included in the calculation, as does the uncertainty coming from the dark current in the CCD N_D . As the amount of used pixels needs to be taken into account, a factor $(1 + n_{\text{pix}}/n_B)$ has to be included with n_{pix} and n_B the amount of pixels in the source and background, respectively. Read noise adds more uncertainty N_R , but as this is not a Poisson process it scales differently with the signal strength. Finally, as the image uses counts as units, which was converted from the amount of e^- in a CCD pixel, the uncertainty in this conversion needs to be taken into account as well through a factor $G\sigma_f$. Here, G is the gain, and $\sigma_f \approx 0.289$ is the 1σ error in an A/D converter (Merline & Howell, 1995). This leads to the CCD equation for the signal-to-noise ratio S/N:

$$\text{S/N} = \frac{N_\star}{\sqrt{N_\star + n_{\text{pix}} \left(1 + \frac{n_{\text{pix}}}{n_B}\right) (N_S + N_D + N_R^2 + G^2\sigma_f^2)}}. \quad (2.5)$$

As $\text{S/N} = 1/\sigma$, where σ is the standard deviation of the measurement, this gives a measure of σ on the flux and can be transformed to magnitudes by

$$\sigma_m = \frac{1.0857}{\text{S/N}}, \quad (2.6)$$

where the factor 1.0857 comes from error propagation.

Limiting magnitude

Finally, with the above measure of the uncertainty on the photometry of a given source, the limiting magnitude can be estimated. An easy definition of the limiting magnitude is the faintest source that can be measured in the image up to some confidence level (usually 5σ). For a $n\sigma$ limiting magnitude the $\text{S/N} = n$. Solving the CCD equation (Eq. 2.1) for $N_\star$ gives

$$N_{\star} = \frac{n}{2} \left(n \pm \sqrt{n^2 + 4 \frac{n_{\text{pix}}}{n_B} (N_S + N_D + N_R^2 + G^2 \sigma_f^2)} \right). \quad (2.7)$$

As $N_{\star} > 0$, only the positive solution makes sense. Using Eq. 2.3 this can then be converted into a limiting magnitude.

2.4.3.2 Spectroscopy

In spectroscopy, at this point the data is still a 2D image from which a 1D spectrum needs to be extracted. An observed object is a line along the spectral direction of the image. Arcs and tellurics (spectral absorption and emission lines from the Earth’s atmosphere) run along the spatial axis. When looking along the spatial axis (assuming this is the horizontal axis), a single row looks approximately like a flat background with a gaussian on top, which is the object. By finding the approximate location of the object in each row and fitting a model to it, the object trace is found.

The extraction itself is done by (weighted) averaging the flux in some region around the trace (e.g. rows $trace - 5$ to $trace + 5$). To remove the background flux, the background is traced in a similar way and subtracted (e.g. by using rows $trace - 20$ to $trace - 15$ and rows $trace + 15$ to $trace + 20$ to estimate the background on both sides of the target spectrum). This gives a 1D spectrum of the object.

The wavelength calibration can be done by extracting a 1D spectrum of the arc along the same trace as the object, mapping the peak locations to the known emission line wavelengths of the emitting material, and fitting a model to interpolate the wavelength to the entire spectrum to get a wavelength solution for the spectrum.

If the extracted object was a standard star, its spectrum can be divided by the known spectrum of that star to obtain the sensitivity function $F(\lambda)$ of the detector. To get a good measure of the sensitivity function it is best to get a high S/N spectrum of a star that has few emission or absorption lines, such as very metal-deficient subdwarfs. Many different catalogues of good spectrophotometric standard stars can be found online, with different stars being good standards in different wavelength ranges.

By dividing the extracted trace of a science target through the sensitivity function,

the spectrum can be flux-calibrated. Up to a certain extent the sensitivity function also contains a measure of the tellurics, which will be removed from the science target spectrum. This works best when the standard and science target have been observed close together in time and airmass, so that it can be assumed that the sky conditions are more or less the same for the two observations. This generally leaves areas with increased uncertainties however. Several other methods exist, each having their own advantages and disadvantages (see e.g. [Rudolf et al., 2016](#), and references therein for more details). In some cases, only a relative sensitivity function is known, resulting in a calibrated spectrum in an unknown flux unit. In these cases, properly calibrated photometry of the object can be used to flux-calibrate the spectrum by integrating the spectrum over the filter transmission curve.

2.4.4 Difference imaging

Large surveys such as ZTF observe the night sky to find new transients and monitor known ones. A powerful method to easily detect changes in the sky is through difference imaging. This method uses a previously constructed reference image and subtracts this from the new science image in such a way that any constant sources such as (non-varying) stars and galaxies are removed. The result is a so-called difference image, which only contains sources whose brightness has changed or have moved (e.g. planets, asteroids, high proper motion stars) between the observations. Figure 2.4 shows an example of a reference image, a science image, and the resulting difference image. A historically popular algorithm for difference imaging is High Order Transform of PSF ANd Template Subtraction (HOTPANTS, [Becker 2015](#))¹, which is used for image subtraction in section 3.3.2.4. Besides this, there are other algorithms as well. For instance, ZTF uses ZOGY ([Zackay et al., 2016](#))². In chapter 5 AUTOPHOT ([Brennan & Fraser, 2022](#))³, which internally calls HOTPANTS or ZOGY, is used to subtract reference images from follow-up observations obtained with the NOT and GTC.

¹<https://github.com/acbecker/hotpants>

²<https://github.com/cenko/ZOGY>

³<https://github.com/Astro-Sean/autophot>




Figure 2.4: Example of an r -band reference image of NGC 7814, the science image of the same region, and the resulting difference image with the host galaxy removed and only the transient (SN 2021rhu, SN Ia-86G) remaining. (Credit: Luke Harvey)

A detailed explanation of difference imaging is beyond the scope of this thesis, and I will only give a very broad overview of the main idea of the technique. For more details I refer the reader to [Alard & Lupton \(1998\)](#) for a description of the algorithm used in HOTPANTS and [Zackay et al. \(2016\)](#) for a detailed discussion of different image subtraction algorithms and the description of the ZOGY algorithm.

Assuming that we have a reference image (e.g. by stacking several previous observations that were made in good observing conditions), the first step is to align the reference and science images to ensure that they show the exact same part of the sky. This includes corrections for rotation, zoom, and scaling (making sure that an n th mag source has the same amount of counts in both images), among others. Next is the actual image subtraction process. [Zackay et al. \(2016\)](#) show that the reference image R can be seen as the true sky T convolved with the PSF of the reference image P_r , while accounting for the atmospheric transparency, telescope, and detector transmission and integration time through a factor F_r ([Zackay & Ofek, 2017a,b](#)). The resulting expression for R can then be written as

$$R = F_r T \otimes P_r + \epsilon_r, \quad (2.8)$$

where $\otimes$ denotes the convolution and ϵ_r is a noise component in R . Assuming that there is a new point source at position q with flux α in the science image N , its

expression can be written as

$$N = F_n T \otimes P_n + \alpha F_n \delta(q) \otimes P_n + \epsilon_n, \quad (2.9)$$

where F_n , P_n , and ϵ_n denote the same quantities for the science image as their reference counterparts, and $\delta(q)$ is a two-dimensional image which is zero everywhere except at position q where it is one. If there is no new source in the science image $\alpha = 0$ and the second term disappears, resulting in the expressions for R and N looking the same. Through some algebra an expression for the optimal transient detection statistic can be found, with the associated difference image having no subtraction residuals or deconvolution artifacts, and only transient sources surrounded by white noise.

2.4.5 Forced photometry

Through PSF photometry the location and brightness of each source in the difference image is determined. These are then compared to the locations of known sources to separate new from known ones. Each location has however been observed for the entire duration of the survey, which means that it is also possible to measure the flux at the location of a known transient before and after it was visible in the images and create a light curve for the full duration of the survey.

This is called forced photometry, because the PSF function is forced to center on a specific location instead of finding the best-fitting position for the centroid. All ZTF light curves that are used in this thesis have been obtained through FPBOT (Reusch, 2023)¹, which uses forced photometry. When there is nothing but noise at that location, the measured flux will be 0 within the error. When there is a source at the target location, it will be measured, but if the source is not at the centre of the PSF, the fit will have trouble converging, resulting in a large uncertainty. Artefacts such as cosmic rays, imperfectly subtracted difference images, or light bleeding effects from saturated bright nearby stars can also affect the accuracy of the photometry measurement.

¹<https://github.com/simeonreusch/fpbot>

2.5 General considerations for observing

During my PhD I have spent two years doing observational studentships in La Palma. The first year was with the Isaac Newton Group of Telescopes (ING) and the second year with the NOT. I gained first-hand experience with the specifics of observing in the optical regime and the considerations that come with it. I will briefly go over these in this section. These studentships also gave me the unique opportunity to quickly follow up on interesting transients, which was very valuable for the objects I will discuss in chapter 5.

2.5.1 Location

Although this is usually already done before constructing a telescope, the first thing to consider is the observing location. When purely aiming for the best observations possible, there are three main things to consider when choosing where to observe from:

- Weather: Clear, stable sky conditions for most nights throughout the year, and low atmospheric distortion (e.g. seeing) are vital to ensure good quality data on a regular basis. Low-hanging clouds can be avoided by being high above sea level, while simultaneously decreasing the amount of air light has to travel through to reach the detector, decreasing atmospheric influence.
- Light pollution: Darker skies allow for observations of fainter objects. Even the presence of a (partially) illuminated moon significantly changes the depth that can be reached with the same exposure time. Many observatories have (inter)national laws to control the amount of light pollution and ensure that good quality data can be obtained.
- Target observability: The target location needs to be reachable by the telescope to be observable. For instance, the NOT can observe down to an elevation of 7° . If a target is above 7° elevation for only 10 minutes during the night, then it can only be observed for a few minutes after acquiring the target. Low elevation

observations are also generally not desired, as the closer to zenith an observation is made, the less atmosphere there is between the target and the telescope. The atmosphere reduces the data quality through turbulence (seeing), broadband absorption (clouds, dust), narrowband interference (tellurics, sky lines), and differential diffraction, among others.

Ideally, observatories should be located on top of high mountains in areas with stable and clear weather with as small a nearby population as possible, while still being accessible enough for transporting materials and observing staff. One of the best locations in the world that meets these requirements is Roque de los Muchachos on La Palma, a small Spanish island in the Atlantic Ocean off the coast of Morocco. At around 2300 m above sea level, the telescopes are built on the highest peak of the mountain, far from most communities on the island that are much closer to sea level, and the temperate climate ensures good sky conditions for most nights around the year. On top of that, in 1988 the government put the so-called "Law of the Sky" in place to minimize light pollution, e.g. by limiting the use of street lights and restricting flight paths over the island ([Gobierno de España, 1988](#)).

2.5.2 Telescope, instrument, observation type, and setup

Depending on the type of observations and the brightness of the target, there is a choice of what hardware to use. Telescope, instrument, observation type, and desired setup(s) have to be considered together, as some choices will affect other ones.

Bigger telescopes can observe fainter targets, but it is also more difficult to obtain observing time. On the other hand, smaller telescopes are less oversubscribed (a measure of requested versus available observing time) but are more limited in observation depth, even with long exposure times.

Secondly, different instruments, which are often telescope-specific, have different observing capabilities. Photometry and spectroscopy are very standard observing modes, and most telescopes have at least one instrument that can offer this. Even though ALFOSC and OSIRIS+ can both observe in these modes, there are still differences in data

quality and resolution, even if the same object is observed at the same time. However, for polarimetric observations, OSIRIS+ cannot be used while ALFOSC can, limiting the options for this observation mode. The sensitivity when wanting to observe faint targets can also matter when choosing which instrument to use, as does the plate scale or dispersion. In imaging, if two objects have a separation similar to or smaller than the plate scale, they cannot be resolved. The same is true in spectroscopy for narrow emission lines that are close to or narrower than the dispersion of the used setup.

Lastly, the specific setup has to be considered as well. For photometry, which filters are desired? If a very specific or rare filter is required, this may limit the options of telescopes and instruments. For spectroscopy there are other choices, such as fiber or slit spectroscopy, different gratings or grisms depending on the desired resolution and wavelength range, neutral density filters to observe targets that are otherwise too bright for the instrument, and order-blocking filters to remove contamination from second-order blue light in red parts of the spectrum.

2.5.3 Night plan

Lastly, it is good to have a plan of what to observe at each point during the night in order to avoid losing observing time during the night. Most proposals already have a list of targets and standard stars to observe with known exposure times when they are submitted to request observation time, but the detailed plan is usually made mere hours before the night starts as it depends on, e.g., the current weather, target priority, and specific time constraints (e.g. for transits). Calibration images might need to be taken during the night as well. All of these things need to be considered when trying to maximize the time used to expose and observe targets, and minimize the overheads from, e.g., positioning, target acquisition, and readout.

Time spent repositioning the telescope can be reduced by finding the path between targets that minimizes telescope and dome movement throughout the night. The target acquisition time depends on the type of observation, but also on the experience and tiredness of the observer. With photometry a field is observed, meaning that a small




Figure 2.5: Night plan for the NOT on the night of 10 August 2024. Targets are plotted with their altitude as a function of universal standard time. Local sidereal time is shown on top. The target priority has been colour-coded, with the coloured bars showing the amount of time each observation is expected to take. Green targets have already been completed, and the red vertical line shows the current time. Several unscheduled backup targets are shown in case the plan has to be updated during the night.

offset is often not disastrous for the science. With spectroscopy the target needs to be identified and placed in the slit or fiber before the exposure can start, costing extra time. Readout times are detector-specific but can be sped up by windowing and binning if only a part of the CCD is needed and a worse resolution is acceptable. Considering readout times can be especially important when multiple shorter exposures are taken instead of a single long one.

Nothing is certain during the night. Weather conditions can change suddenly, technical problems can occur, a high-priority target can be discovered during the night and override planned observations, or observations might go so smoothly that they are completed faster than expected. A flexible schedule with a priority list and backup targets helps to adapt to these situations quickly. After all, an idling telescope in (half-)decent observing conditions is a waste of resources. Figure 2.5 shows an example night plan for the NOT, made using VISPLOT¹, with space for adaptability built in.

¹<https://www.not.iac.es/observing/forms/visplot/index.php>

ZTF SN Ia DR2: Searching for late-time interaction signatures in Type Ia supernovae from the Zwicky Transient Facility

This chapter details my search for late-time interaction signals in the ZTF second data release (ZTF SN Ia DR2, [Rigault et al., 2025a](#), Smith et al., in prep.) and is based on [Terwel et al. \(2025\)](#), which was accepted for publication in A&A. This chapter contains the same details as [Terwel et al. \(2025\)](#) but expands on section 3.2.3 with an extra figure. The data reduction, analysis and interpretation were performed by me. Where relevant, I have noted the specific parts that were carried out by collaborators.

As stated in section 2.2.1, the magnitude limit for single ZTF exposures is ~ 20.5 mag. By binning post-SN observations together for each object in the ZTF SN Ia DR2 sample, I attempt to push beyond this limit to find any faint signals that have gone previously unnoticed due to being close to the detection limit. This requires building and testing an automated pipeline, the explanation of which is given in section 3.2. In chapters 4 and 5 I will slightly modify this pipeline to use it on different samples with the same goal of finding faint late-time excesses.

In section 3.1, I introduce the sample I use in my search for optical signals for late-time CSM interaction in the ZTF data stream. In section 3.2, I present my custom pipeline for the identification of late-time flux excesses, and set up a simulation to test it and estimate the detection efficiency of my pipeline. Section 3.3 shows the result of running my sample through my pipeline and provides further investigation on some interesting objects. Section 3.4 discusses the results, and I conclude in section 3.5.

3.1 Data

My aim is to look for late-time (> 100 days after peak brightness) signatures of CSM interaction in the largest sample of SNe Ia to date. This has been obtained by the ZTF. I am particularly interested in events that appear to be normal SNe Ia from their spectra and light curves around the peak but may display signs of late-time interaction, as seen in SN 2015cp (Graham et al., 2019b). My starting sample is 3 628 events that were discovered by ZTF from March 2018 to October 2020 (hereafter the ZTF data release 2, ZTF DR2). Each event is spectroscopically classified as a SN Ia or one of its sub-classes. An overview of the ZTF DR2 is presented in Rigault et al. (2025a), including the sample definition, properties, and use for cosmology. In this study, since I am searching for likely rare signatures of interaction in the ZTF light curves, I am as inclusive as possible in my sample definition and include all SNe Ia in the DR2 covering a redshift out to $z = 0.288$.

3.1.1 ZTF light curve data

ZTF observes in three optical bands *gri* on a 2 – 3 day cadence. Reference images, mainly made using observations at the start of ZTF, are subtracted from the science images using the ZOGY image subtraction algorithm (Zackay et al., 2016) to produce difference images. Forced photometry is used at the transient location on the difference images using ZTFFPS (Reusch, 2020) to get a measure of the observed flux at each epoch. This includes non-detections before each SN was first detected and after each SN has faded below detection limits.

Light curve quality cuts on specific light curve points are applied as in Smith et al. (in prep.). I do not correct the light curves for Milky Way extinction in my initial analysis but do consider it when focusing on specific objects of interest in section 3.3.

Another approach for extracting photometry at the transient location is by using Scene Modelling Photometry (SMP; Holtzman et al., 2008). SMP avoids PSF degradation and spatial resampling that is present in difference imaging. Instead, all frames

are fit simultaneously with a model M of the stars, galaxies, SNe, and background in the region simultaneously with only the SNe being allowed to vary over time. All frames have their own PSF such that for each frame I the relation $I = \text{PSF} \otimes M$ holds. Isolated stars are used to determine the astrometric and photometric transformations between the frame, and stars and SNe are required to have the same relative position in each frame.

SMP is extracted for a few selected objects of interest in section 3.3 to test if the identified late-time detections are independent of the approach. When using SMP, one has to define an ‘off’ time and an ‘on’ time. The observations taken during the off time are used to create a model of the region, or scene, where the SN occurs. This is then used as a template during the on time to calculate and remove host contributions to the photometry, leaving just the transient itself (Lacroix et al., in prep.). The advantage of this method is the significantly lower uncertainty in the model compared to the difference imaging technique, allowing one to find fainter detections. The more images that are included in the off time, the better the construction of the model will be. Since I assume that a signal from late-time interaction could occur at any point after the SN, The off time is defined as everything up to shortly before the SN explodes, and the on time as everything after this moment. As SMP is used only for further investigation, I can use the found SN peak times in the initial analysis to set the off and on times.

3.1.2 Baseline correction

Issues in the construction of the reference images, such as inaccurate flat-fielding or artefacts in the images that are subsequently co-added can result in a systematic offset in the forced photometry light curve made using different images. The baseline correction technique is used to correct for this (Yao et al., 2019; Miller et al., 2020).

To estimate the necessary baseline correction, I calculated the weighted mean of the flux of all data points up to 40 days before the estimated SN peak (which is assumed to be the highest flux detection deemed real), and subtracted it from the light curve. Baseline corrections are done separately for each combination of band (gri), field

Table 3.1: Initial sample size and its reduction in each step of the analysis process.

Criterion	Removed	Objects left
Initial DR2 sample	-	3 628
No photometry at 100+ days	109	3 519
No robust late-time detections ^(a)	2 953	566
Presence of SN Ia tail ^(b)	432	134
Removed on visual inspection ^(c)	101	33
No late-time CSM interaction ^(d)	30	3

^(a)For a ‘robust detection’, at least four positive detections (two or more adjacent bins with $\geq 5\sigma$) are required out of the 16 possible combinations of bin size (25, 50, 75, 100 days), and the starting position of the bin varied by 25, 50 or 75% of the bin size.

^(b)I tested for the presence of a radioactive tail of the SN Ia as described in section 3.2.2 and removed those where this was the most plausible explanation.

^(c)Each remaining light curve was inspected using SNAP (Section 3.2.3) for possible issues causing late-time detections. See section 3.3 for a discussion of the reasons events were removed.

^(d)For each remaining light curve I checked in detail if the late-time detections could be explained without CSM interaction starting at late times.

(telescope pointing), and rcid (part of the camera, which is arranged in 4×4 charge-coupled devices (CCDs) with four readout channels each, giving a total of 64 readout channels), as each of these combinations uses different, unique reference images. To be able to apply a baseline correction, at least two observations are needed. If this is not possible all observations with that band, field, and rcid combination are removed.

Since I am interested in post-SN detections, my baseline correction method using only pre-SN data differs from the one used in Smith et al. (in prep.) with both pre- and post-SN data. A comparison between the methods found that for most objects my corrections agree with the ones used in Smith et al. (in prep.) within the uncertainties. This is as expected as objects with late-time flux excesses are expected to be rare, meaning that the two baseline correction methods should give the same result for most objects.

If there is insufficient data to perform a baseline correction, the relevant data (based on field, filter and rcid) is removed from the light curve. If this includes data around the peak position, the peak position in the light curve may change and therefore, the position of the peak is recalculated.

3.2 Analysis

To systematically search for objects with late-time flux excesses, I developed a custom pipeline. In section 3.2.1, I describe how the late-time photometry for each object is binned to reach deeper magnitude limits, as well as the scheme used to select objects with robust, significant detections. In section 3.2.2, I identify and remove bright nearby SNe Ia whose late-time detections are due to the SN radioactive decay tail. In section 3.2.3, I describe my method of visually inspecting images of potentially interesting sources, and in section 3.2.4, I detail the use of SIMSURVEY to simulate SNe Ia with late-time interaction signatures. Table 3.1 shows my sample size after each step of analysis that is discussed in the subsequent sections.

3.2.1 Binning & filtering program

After pre-processing, the late-time observations are binned in phase (days since the observed peak in the observer frame) to push the detection limit beyond the limit of the individual observations. I define late-time observations as being at least 100 days after the estimated date of SN peak brightness in the observer frame. I remove all SNe Ia that have no data in any band beyond this phase. The exact choice of 100 days is arbitrary but was chosen as a balance between minimising spurious detections due to the light curves being dominated by SN light at earlier times and maximising the phase range over which the interaction can be searched for.

Binning of the light curves is performed in each band separately. Larger bins are better for pushing the magnitude limit as deep as possible, but they sacrifice temporal sensitivity. To balance the time sensitivity and magnitude limit, I use bins with widths of 100, 75, 50, and 25 days. To make sure that the placement of the bin edges does not affect my results, I repeat the binning four times for each bin size while decreasing the size of the first bin by 25% in each iteration. This results in a total of 16 trials for each band.

Each bin starts at the position of a data point to avoid empty bins. A gap in the data that is larger than the bins being used can cause the bins after the gap to always be placed in the same location despite the size modifications of the first bin. To avoid this I re-apply this modification of adjusting the size of the first bin after the data gap to trial different start positions. Lastly, if a bin would only contain one or two points, and the phases of these points occur no later than 10% of bin size of the previous bin (e.g., if the bin size is 25 days, they would occur within 2.5 days of the end of the bin), the previous bin is increased in size to include these points. An example of the bin placement is shown in figure 3.1.

For each bin the weighted mean and uncertainty of the observed flux are calculated. When binning flux measurements taken from difference images, the uncertainty of the reference images used in the difference imaging procedure has to be considered, as this will limit the depth of the binned observations (Strotjohann et al., 2021) and this is added in quadrature to the weighted uncertainty of the binned flux.

After the binning procedure each light curve has undergone 16 trials per band across four bin sizes and four bin placements. An attempt in a specific trial is considered significant if it has two or more adjacent bins with $\geq 5\sigma$ detections. A late-time detection is considered ‘robust’ if at least four out of 16 attempts have significant detections, suggesting that the detections are insensitive to bin placement and/or size. I make this choice of robust detection in at least four attempts to ensure that I am not dominated by spurious detections but that I can still detect a long but faint signal that can only be picked up in the four trials involving the largest (100 day) bins.

3.2.2 Removal of SN Ia radioactive tail detections

For some nearby SNe Ia, the normal light curve tail powered by the radioactive decay of $^{56}\text{Co} \rightarrow ^{56}\text{Fe}$ is still visible at the phases investigated here (>100 days after the peak), possibly triggering false positives in my pipeline. To test if the fading tail is the reason for detections after 100 days, I checked if the detections follow a declining power law in flux space consistent with that of a normal SN Ia tail. I take all bins, normalise to




Figure 3.1: First 250 days of the *gri*-band light curves of SN 2019hbb in flux scaled to the peak flux (top panel) and magnitude (bottom panel) space. Binning starts 100 days after the estimated peak date (vertical black dashed and solid lines), using 25 d bins. The *g*-, *r*-, and *i*-bands are shown in green, orange, and red, respectively. Before 100 days, I show the unbinned detections with their uncertainties (coloured circles) and non-detections (inverted triangles). After 100 days I show the bins as horizontal lines to show their size, a circle to show their mean value, and the shaded region showing the 1σ uncertainty (dashed regions are non-detections). A bin is deemed a non-detection if the flux $f < 5\sigma_f$. The 5σ magnitude limit is calculated and shown as a downward arrow. In both the *g*- and *r*-bands, the first bin is a detection and there are multiple adjacent bins with detections, triggering the tail-fitting procedure which is done in flux space (see section 3.2.2). The resulting tail fits are shown in the green and red lines, respectively, with their 1σ uncertainties as hashed regions. The half-life times are $t_{1/2,g} = 70 \pm 6$ d ($\chi^2_{\text{red}} = 0.6$) and $t_{1/2,r} = 27 \pm 4$ d ($\chi^2_{\text{red}} = 1.4$). This tail is therefore deemed to be a normal SN Ia tail.

the brightest data point, and fit a declining power law. To ensure that the tail matches the earlier data points, I include the unbinned observations between 60 and 100 d after the peak, making sure that if there are N bins, only the latest $N/2$ unbinned detections are used to ensure that the fit focusses on the bins and not the unbinned points. All points are fitted using $1/\sigma_f$ weights, where σ_f is the flux uncertainty of point f .

A successful fit of a declining normal SN Ia tail has a reduced chi-square of $\chi^2_{\text{red, fit}} < 5$ and fitted half-life of $t_{1/2}$ with uncertainty $\sigma_{t_{1/2}}$ satisfying $t_{1/2} - 5\sigma_{t_{1/2}} \leq 50$ d. I chose a threshold of 50 d as [Dimitriadis et al. \(2017\)](#) showed that this is the approximate decay time scale for a normal SN Ia at these phases. A fit with a high $\chi^2_{\text{red, fit}}$ value could have failed due to bad or uncertain data, or due to the late-time detections not following a power law decay. Fits with a $t_{1/2}$ significantly larger than that of a normal SN Ia tail suggest an additional luminosity source contributing to the light curve at these phases. Figure 3.1 shows an example where this tail fitting procedure determines the late-time detections in an object to be a normal declining SN Ia tail. 432 SNe Ia that were flagged as having late-time detection are discounted from further discussion because their light curves can be explained by a normal fading SN Ia tail.

3.2.3 SuperNova Animation Program (SNAP)

After performing the binning and filtering, and removing SNe Ia with contamination from the radioactive tail, I am left with 134 SNe Ia with robust late-time detections (see Table 3.1). Since the binning and filtering program is designed to handle a large number of light curves and cannot be tailored specifically to suit the peculiarities of a single object, it is possible that there are objects remaining with issues in the data or data processing (e.g. cosmic rays, bad subtractions), resulting in false positive detections. Therefore, I manually check the difference imaging to search for potential issues. To do this efficiently, I use the SuperNova Animation Program (SNAP)¹.

Using ZTFQUERY ([Rigault, 2018](#)), SNAP takes all difference images of the requested sky position during the requested time period(s) in the requested band(s) and shows

¹https://github.com/JTerwel/SuperNova_Animation_Program




Figure 3.2: Frame of a movie made using SNAP showing the r -band difference images of ZTF18aasdted (SN 2018big). The left side shows a cutout of the difference image, with the red + at the location of ZTF18aasdted. The magenta + is at the location of ZTF19abqhobb (SN 2019nvm). The plots above and right of the image show the mean value of the middle five rows and columns, respectively, with the green square indicating which rows and columns are used for this. The compass shows the orientation of the image, and the wireframe gives a different way of viewing the pixel values. General information of the image and object is shown in the top right corner, as well as some statistics of the shown cutout. The plot in the bottom right shows the light curve at the red +. The black dashed line shows the mjd of the currently shown image, and the black circle highlights the data point associated with it. The bottom right light curve can be replaced by a plot showing the latest spectrum at the time of the shown difference image.

them in chronological order in an animation. At the start of each animation the reference images in all bands are shown. The program can show the image in grey-scale, a three-dimensional wire-frame representation of the intensities measured per pixel, the averaged values along both axes of the image, the observation date and duration, the peak and mean pixel values of the shown region, the last spectrum taken before the currently shown image, and highlight the resulting forced photometry point in the light curve corresponding to the plotted images. Figure 3.2 shows a frame of such a SNAP movie.

Using SNAP, issues in the difference images can be identified, including SN ghosts (the SN is visible in the reference image, leaving a negative imprint in difference images after it has faded), cosmic rays, and bad pixels (NaN, or a large negative number). Variability of a separate source can also be seen, which, when close by, can contaminate the forced photometry at the SN location. An example of such a contaminant is an active galactic nucleus (AGN).

3.2.4 Simulated interaction recovery fractions

To make sure the binning program works as expected and estimate its late-time signal recovery efficiency, I simulated an observing campaign using SIMSURVEY (Feindt et al., 2019a,b), a PYTHON package designed to simulate large-scale time domain surveys such as ZTF. To successfully simulate an observing campaign, the program needs to be told what, when, where, and how something is observed, and under what conditions. For this, I need a model of the SN Ia-CSM that is going to be observed, an explosion rate as a function of redshift, and a time range for these explosions to occur. I also require an observing log specifying which part of the sky is being observed at a specific time, the length of the observations, and the weather conditions during the observations. Lastly, I require details of the camera that is used to carry out the observations. The SN model needs to be in a similar format to the SNCOSMO models (a PYTHON package made for supernova cosmology, Barbary et al. 2021). This means that I need a set of spectra over the entire phase range, all having the same wavelength spacing and range. No such model exists, so I built my own model as described below.

3.2.4.1 The interacting SN model

I chose SN 2011fe as the template of a normal SN Ia as it is well observed and close by (in M101 at a distance of 6.4 Mpc, Shappee & Stanek 2011). Optical spectra of SN 2011fe were obtained between phases of -18 to 1017 d relative to the peak. I made a custom model using SNCOSMO and spectra found on WISereP (Yaron & Gal-Yam, 2012)¹, which are listed in Table 3.2. The spectra were flux-calibrated to match the observed coeval broadband magnitudes. Spectra up to 45 d after the peak were flux calibrated using a SALT2 (Guy et al., 2007) fit of the *PTF48g* and *PTF48R*-band photometry (Law et al., 2009; Rau et al., 2009), which is used to estimate the flux in the *g*- and *r*- bands at these phases. Spectra between 45 and 400 d are flux calibrated using a cubic spline interpolation of photometry in the *PTF48g*- and *PTF48R*-bands.

¹<https://wiserep.weizmann.ac.il>

The interpolation extends up to 600 d after the peak in the *PTF48R*-band, there is no *PTF48g* photometry used between 400 and 600 d after the peak. Three interpolated photometry points from the SALT2 fits were used as anchor points to connect these two parts of the calibration. [Dimitriadis et al. \(2017\)](#) show that there is a slight kink in the light curve tail around 600 d after the peak. This is replicated in the model by calibrating all spectra more than 600 d after the peak using a cubic spline interpolation of photometry from the Large Binocular Telescope (LBT; [Hill et al., 2006](#)) in the Bessel *R*-band ([Shappee et al., 2017](#)).

Table 3.2: Spectra used to make the SN 2011fe model. All spectra were taken from WISeREP ([Yaron & Gal-Yam, 2012](#)).

MJD	Phase (d)	Telescope	Instrument	Wavelength coverage (Å)	Reference
55798.0	−16.0	Lijiang-2.4m	YFOSC	3461 – 8956	Zhang et al. (2016)
55798.2	−15.8	Lick-3m	KAST	3416 – 10278	Nugent et al. (2011)
55799.0	−15.0	Lijiang-2.4m	YFOSC	3502 – 8958	Zhang et al. (2016)
55799.3	−14.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55800.2	−13.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55801.2	−12.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55802.3	−11.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55803.2	−10.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55804.2	−9.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55805.2	−8.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55806.2	−7.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55807.3	−6.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55808.2	−5.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55809.2	−4.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55811.4	−2.6	HST	STIS	1779 – 24965	Mazzali et al. (2014)
55812.0	−2.0	Gemini-N	GMOS	3497 – 9648	Parrent et al. (2012)
55813.2	−0.8	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55814.2	0.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55815.2	1.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55816.2	2.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55817.2	3.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55817.7	3.7	HST	STIS	1265 – 24965	Mazzali et al. (2014)
55818.2	4.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55821.2	7.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55823.2	9.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55826.2	12.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55828.2	14.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55829.0	15.0	Gemini-N	GMOS	3497 – 9643	Parrent et al. (2012)
55830.2	16.2	Keck1	LRIS	3227 – 10242	Stahl et al. (2020)
55831.2	17.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55832.0	18.0	Lijiang-2.4m	YFOSC	3577 – 8957	Zhang et al. (2016)
55833.2	19.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55835.3	21.3	HST	STIS	1731 – 10221	Mazzali et al. (2014)
55836.2	22.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55838.2	24.2	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55841.3	27.3	HST	STIS	1738 – 10221	Mazzali et al. (2014)
55855.2	41.2	HST	STIS	1738 – 10216	Mazzali et al. (2014)
55888.6	74.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55891.7	77.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)

3. ZTF SN IA DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY

MJD	Phase (d)	Telescope	Instrument	Wavelength coverage (Å)	Reference
55893.6	79.7	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55896.6	82.6	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55897.7	83.7	Keck1	LRIS	3164 – 10126	Stahl et al. (2020)
55901.6	87.6	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55903.6	89.6	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55911.0	97.0	XLT	BFOSC	3296 – 9693	Zhang et al. (2016)
55911.6	97.6	UH88	SNIFS	3296 – 9693	Pereira et al. (2013)
55913.5	99.5	Lick-3m	KAST	3427 – 10332	Stahl et al. (2020)
55914.0	100.0	WHT-4.2m	ISIS	3499 – 9491	Friesen et al. (2017)
55916.0	102.0	WHT-4.2m	ISIS	3498 – 9491	Law et al. (2009); Rau et al. (2009)
55917.0	103.0	WHT-4.2m	ISIS	3499 – 9492	Law et al. (2009); Rau et al. (2009)
55926.0	112.0	Lijiang-2.4m	YFOSC	3366 – 9069	Zhang et al. (2016)
55929.5	115.5	Lick-3m	KAST	3426 – 10170	Stahl et al. (2020)
55944.5	130.5	Lick-3m	KAST	3453 – 10088	Stahl et al. (2020)
55959.0	145.0	Lick-3m	KAST	3497 – 10000	Law et al. (2009); Rau et al. (2009)
55980.4	166.4	Lick-3m	KAST	3441 – 10250	Stahl et al. (2020)
55988.0	174.0	WHT-4.2m	ISIS	3495 – 9982	Mazzali et al. (2015)
56019.4	205.4	Lick-3m	KAST	3438 – 10324	Mazzali et al. (2015)
56040.4	226.4	Lick-3m	KAST	3437 – 10178	Mazzali et al. (2015)
56047.0	233.0	Lijiang-2.4m	YFOSC	3392 – 9053	Zhang et al. (2016)
56073.0	259.0	WHT-4.2m	ISIS	3495 – 9483	Mazzali et al. (2015)
56103.0	289.0	WHT-4.2m	ISIS	3423 – 10268	Mazzali et al. (2015)
56127.0	313.0	P200	DBSP	3197 – 10991	Law et al. (2009); Rau et al. (2009)
56162.2	348.2	Lick-3m	KAST	3487 – 10240	Mazzali et al. (2015)
56194.2	380.2	Keck1	LRIS	3232 – 10268	Stahl et al. (2020)
56277.0	463.0	Lijiang-2.4m	YFOSC	3379 – 9337	Zhang et al. (2016)
56778.5	964.5	Keck1	LRIS	3074 – 10320	Graham et al. (2015a)
56831.2	1017.2	LBT	MODS1	3098 – 10487	Taubenberger et al. (2015)

After flux calibration, the spectra were dereddened to remove dust extinction effects, using the [Cardelli et al. \(1989\)](#) extinction law with $A_V = 0.04$ mag ([Patat et al., 2013](#)). The spectra were rebinned, and any wavelength region that was not covered in all spectra was removed. Lastly, the model is corrected for distance, redshift and time dilation. The resulting model is that of a normal SN Ia, which exploded at a distance of 10 pc without any dust between the source and observer.

The best late-time detection of CSM in a normal/91T-like SN Ia was for SN 2015cp, where $H\alpha$ emission was identified in its spectra at 664 d after the light-curve peak ([Graham et al., 2019b](#)). To model potential CSM interaction signals similar to that of SN 2015cp, I add a narrow $H\alpha$ line with a Gaussian profile to the SN 2011fe model. Due to the rareness of interaction in otherwise normal SNe Ia, I do not have good constraints on the diversity of interaction signatures and simulate a broad parameter space. The interaction was chosen to start at 100, 200, 300, or 500 days after the peak, last for 100, 300, or 500 days, and have a similar brightness to the observed signal in SN 2015cp ([Graham et al., 2019b](#)), as well as 10 times weaker or 10 times stronger




Figure 3.3: Model spectrum at 300 days (SN 2011fe with the added $H\alpha$ line) is shown in magenta overlaid on a rest-frame spectrum of SN 2015cp at 694 days in grey. The model flux has been scaled to the distance of SN 2015cp for comparison. The green, orange, and red shaded regions are the bandwidths of the g -, r -, and i -bands, respectively. The transmission profiles are plotted in the same colours for each band. The model is also shown shifted to $z = 0.07$ (and offset up in flux), where the $H\alpha$ line has just started to be in the i -band.

than it. All possible combinations of these values are used, and a simulation without any interaction is also used as a control test, giving a total of 37 simulations.

Figure 3.3 shows an example spectrum of my model at 300 d and SN 2015cp at 694 d in the rest frame (Graham et al., 2019b). It also shows the model redshifted to $z = 0.07$, where the $H\alpha$ line is partly shifted into the i -band. Figure 3.4 shows the absolute magnitude ri -band light curves of the SN 2011fe model in the rest frame, as well as light curves with different strengths of $H\alpha$ emission. My interaction model will only generate a late-time interaction signal in the r -band (or i -band at $z > 0.06$) as I only add a $H\alpha$ emission line. This is enough to test the binning program but is likely too simple to reflect the actual late-time signal seen in SN 2015cp (which also showed O I and Ca II in the rest frame i -band) or potential other events.

3.2.4.2 Simulating the observing campaign

By specifying the object to observe, as well as the telescope details and observation schedule, an observation campaign can be performed using SIMSURVEY, resulting in a collection of observed light curves. For a deeper explanation of SIMSURVEY I refer the




Figure 3.4: r (orange) and i (red) absolute magnitude light curves of the SN 2011fe model used in the simulations in the ZTF bands as a function of phase from rest-frame B -band peak (Mazzali et al., 2014). The bumpiness in the models is because the underlying SNCOSMO model class interpolates in flux space but fails to find an exponential decay. The added rest-frame CSM interaction model based on $H\alpha$ emission (starting at a phase of 100 d) is shown with dotted lines for the r -band. Once the interaction becomes the dominant source, it smooths out the bumps from the underlying tail. The black line shows a radioactive decay with $t_{1/2} = 50$ d, typical of a declining normal SN Ia tail.

reader to Feindt et al. (2019b). The parameters used as input are:

- Model: SN 2011fe + $H\alpha$ line (Section 3.2.4.1).
- Sky distribution: $RA \in [0^\circ, 360^\circ]$, $Dec. \geq -30^\circ$ (The area covered by ZTF, Bellm et al. 2019a).
- Volumetric rate: The SN Ia rate is $2.4 \times 10^{-5} \text{ Mpc}^{-3} \text{ yr}^{-1}$ for $z \leq 0.1$ (Frohmaier et al., 2019). SIMSURVEY uses this to calculate the amount of SNe to generate at a given redshift interval.
- SN peak time distribution: $58\,195 \leq \text{modified Julian date (MJD)} \leq 58\,487$ (between 18 March 2018 and 4 January 2019).
- Galactic extinction: dust maps from Schlegel et al. (1998).

- Host galaxy extinction: [Cardelli et al. \(1989\)](#) extinction law, with $E(B - V)$ drawn from an exponential distribution with exponent $\lambda = 0.11$ ([Stanishev et al., 2018](#)), the same way as host extinction was added in the original SIMSURVEY paper ([Feindt et al., 2019b](#)).
- Telescope specifications: ZTF P48 camera, 4×4 grid of CCDs with four readout channels each, resulting in 64 separate output channels ([Dekany et al., 2020](#)).
- Survey plan: ZTF observation logs between $58\,197 \leq \text{MJD} \leq 59\,211$ (between 20 March 2018 and 28 December 2020), ensuring all simulated SNe are followed for a minimum of about 2 years after the peak.

In each SIMSURVEY run, 10^5 SNe Ia are simulated to produce observed light curves and meta-data such as redshift, observed peak date, etc. To ensure that the SNe Ia are similar around the peak to those recovered, I require that the SN Ia light curves must have at least three detections of $\geq 5\sigma$ and are brighter than 19 magnitude at the peak. This reduces the sample to $\sim 40,000$ objects per simulation. These are sent through the binning and filtering program (Section 3.2.1) as if they were real observed light curves to determine the recovery efficiency.

The volumetric rate used as input in the simulations favours more distant SNe, which results in very few SNe at extremely low redshift values and hence larger uncertainties. To mitigate this, I split $0 \leq z \leq 0.015$ into bins of size 0.001 and simulated an additional 100 SNe in each bin using the same parameters as in the original simulations. Introducing these additional events does not impact the recovery efficiencies because I am comparing the number of recovered events relative to the input number in each redshift bin and therefore, am insensitive to the input rate of events.

3.2.4.3 Simulated interaction recovery

My aim is to determine from my simulations how many SNe Ia with signatures of late-time interaction similar to that of SN 2015cp would have been detected by my pipeline. For each of the simulations, I binned the SNe based on their redshift and




Figure 3.5: Fraction of SNe Ia for one of my simulations (interaction occurring between 500 – 1000 d after the peak) where the CSM interaction was recovered per redshift bin of size 0.002. The simulations are shown for interaction strengths of zero (grey), similar to SN 2015cp (red), and 10 times stronger than SN 2015cp (blue). In the simulation without CSM interaction, the recovery fraction should be interpreted as the fraction of false positives. The simulations with normal ZTF quality reference images are shown with dots and fitted sigmoid functions with solid lines. Simulations where one magnitude deeper reference images were assumed are shown in triangles, with their fitted sigmoid functions in dashed lines.

looked at the fraction of SNe that were reported by the pipeline to show late-time excesses. Figure 3.5 shows the recovery fractions as a function of redshift for an example simulation when the interaction starts at 500 d and lasts for 500 d for simulations of no CSM interaction, late-time interaction with the same strength as SN 2015cp, and interaction 10 times as strong as SN 2015cp (strong interaction). As expected, the recovery fraction drops off with increasing redshift for both the SN 2015cp equivalent strength and the interaction that is 10 times stronger, with the strong interaction recoverable out to a higher redshift.

The recovery fraction of the simulations with CSM interaction does not reach 100 per cent in the lowest redshift bins. This is because the radioactive tails of these SNe Ia tend to be bright out to hundreds of days after the explosion. Therefore, depending on the cadence and uncertainties of the simulated photometry, the SN light can dominate

over the CSM interaction and the CSM interaction signal does not alter the shape of the SN decay tail enough to be flagged as CSM interaction.

In the simulation without CSM interaction, the recovery fraction is non-zero at small redshifts, meaning that some objects are falsely identified as having a late-time excess. For these very bright and high S/N SN Ia light curves, the decaying tail model for normal SNe Ia proves to be too simple. My analysis pipeline detects real deviations of the SN light curve evolution from the simple decay tail model. This only occurs at the lowest redshifts and nearby SNe Ia are rare, with only 0.6% of my observed SN Ia sample at $z \leq 0.01$ ($d \approx 43$ Mpc). This means that contamination of my sample due to normal SNe Ia tails that cannot be fit by the simplified tail fit model is very low.

I fitted a sigmoid function to the recovery fractions of each simulation, using the total amount of objects in each bin as its weight (figure 3.5). A sigmoid function is an oversimplification (the recovery fraction is underestimated at the low redshifts) but it allows me to easily estimate the redshift limit where CSM interaction can be recovered. I define my redshift limit where CSM interaction can be recovered as z_{50} : the redshift where 50% of the SN interactions are recovered. These values are listed for all simulations in Table 3.3.

As discussed in section 3.2.4.1, I simulated 36 models with interaction signatures starting at 100, 200, 300 and 500 d after the peak, lasting for 100, 300 and 500 d, and with strengths the same as SN 2015cp, 10 times weaker and 10 times stronger. I also simulated a model without any late-time CSM interaction. For the models with an interaction strength similar to SN 2015cp (Graham et al., 2019b), when the interaction is short and early (starting 100 days after the peak and lasting for 100 days), the interaction cannot be distinguished from a normal SN Ia decaying light curve and the recovery fraction is as low as the no interaction simulation. When the interaction is longer, or if it starts later, the light curve flattens enough to be identified as deviating from a normally declining SN Ia tail. This pushes the redshift boundary where 50% of the interaction would be recovered by ZTF to $z_{50} = 0.0105 \pm 0.0003$ ($d \approx 45 \pm 1$ Mpc) for the longest and latest interaction (500 - 1 000 days after peak).

3. ZTF SN IA DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY

Table 3.3: Redshift values where 50% of the simulated SNe were found to have CSM interaction. Strength is the strength of the $H\alpha$ line compared to the strength detected in SN 2015cp. Start shows how many days after the peak the interaction begins, and duration is in days as well. I fitted sigmoid functions to the results of each simulation in order to find the redshift where 50% of the interactions were recovered, assuming the reference images were of the same depth as the ones used in ZTF or 0.5 or 1 mag deeper. These values are shown in $z_{50} \pm \sigma_{z_{50}}$, and χ^2_{red} shows the quality of the fit.

strength	start	duration	ZTF references			0.5 mag deeper			1 mag deeper		
			z_{50}	$\sigma_{z_{50}}$	χ^2_{red}	z_{50}	$\sigma_{z_{50}}$	χ^2_{red}	z_{50}	$\sigma_{z_{50}}$	χ^2_{red}
0.0	—	—	0.0038	0.0009	0.72	0.0042	0.0019	1.83	0.0046	0.0032	3.47
0.1	100	100	0.0041	0.0008	0.83	0.0049	0.0011	1.82	0.0049	0.0033	3.41
0.1	100	300	0.0050	0.0007	0.79	0.0055	0.0014	1.79	0.0059	0.0025	3.42
0.1	100	500	0.0053	0.0006	0.77	0.0062	0.0011	1.76	0.0068	0.0018	3.44
0.1	200	100	0.0045	0.0007	0.78	0.0050	0.0016	1.80	0.0054	0.0028	3.38
0.1	200	300	0.0053	0.0006	0.78	0.0058	0.0014	1.79	0.0063	0.0022	3.44
0.1	200	500	0.0056	0.0006	0.77	0.0066	0.0010	1.79	0.0073	0.0016	3.47
0.1	300	100	0.0050	0.0006	0.80	0.0056	0.0012	1.80	0.0059	0.0023	3.43
0.1	300	300	0.0052	0.0006	0.80	0.0061	0.0044	1.78	0.0067	0.0017	3.46
0.1	300	500	0.0053	0.0005	0.82	0.0060	0.0010	2.03	0.0069	0.0016	3.53
0.1	500	100	0.0045	0.0007	0.82	0.0048	0.0014	2.08	0.0054	0.0021	3.52
0.1	500	300	0.0050	0.0005	0.83	0.0056	0.0010	2.09	0.0062	0.0017	3.53
0.1	500	500	0.0050	0.0005	0.82	0.0057	0.0009	2.09	0.0063	0.0017	3.52
1.0	100	100	0.0041	0.0007	0.86	0.0045	0.0013	2.10	0.0046	0.0027	3.58
1.0	100	300	0.0091	0.0005	0.85	0.0096	0.0012	1.98	0.0108	0.0017	3.40
1.0	100	500	0.0111	0.0004	0.84	0.0129	0.0008	2.08	0.0139	0.0013	3.50
1.0	200	100	0.0067	0.0007	0.80	0.0077	0.0012	2.05	0.0083	0.0021	3.51
1.0	200	300	0.0104	0.0005	0.93	0.0122	0.0008	2.06	0.0130	0.0014	3.31
1.0	200	500	0.0116	0.0004	0.84	0.0129	0.0008	2.11	0.0140	0.0018	3.63
1.0	300	100	0.0086	0.0004	0.78	0.0094	0.0008	1.95	0.0100	0.0013	3.62
1.0	300	300	0.0105	0.0004	0.75	0.0118	0.0007	2.02	0.0128	0.0015	3.61
1.0	300	500	0.0109	0.0004	0.78	0.0125	0.0008	2.11	0.0139	0.0013	3.76
1.0	500	100	0.0081	0.0003	0.83	0.0090	0.0008	1.94	0.0094	0.0016	3.69
1.0	500	300	0.0104	0.0003	0.78	0.0121	0.0007	2.10	0.0133	0.0012	3.79
1.0	500	500	0.0105	0.0003	0.80	0.0122	0.0007	2.11	0.0135	0.0012	3.80
10.0	100	100	0.0091	0.0014	0.93	0.0097	0.0028	2.13	0.0087	0.0058	3.45
10.0	100	300	0.0259	0.0007	1.36	0.0299	0.0009	1.79	0.0328	0.0012	1.83
10.0	100	500	0.0297	0.0005	1.21	0.0347	0.0007	1.83	0.0383	0.0010	1.84
10.0	200	100	0.0204	0.0009	1.69	0.0223	0.0012	2.40	0.0217	0.0019	2.54
10.0	200	300	0.0314	0.0004	1.45	0.0351	0.0007	1.98	0.0377	0.0010	2.40
10.0	200	500	0.0326	0.0005	1.50	0.0374	0.0007	2.05	0.0415	0.0009	1.95
10.0	300	100	0.0221	0.0023	1.21	0.0233	0.0011	2.01	0.0231	0.0019	2.71
10.0	300	300	0.0311	0.0004	1.30	0.0350	0.0007	1.90	0.0375	0.0010	2.27
10.0	300	500	0.0328	0.0004	1.42	0.0381	0.0007	2.22	0.0423	0.0009	2.21
10.0	500	100	0.0199	0.0011	1.80	0.0224	0.0014	2.47	0.0226	0.0019	3.04
10.0	500	300	0.0323	0.0004	1.23	0.0371	0.0007	2.12	0.0406	0.0009	2.19
10.0	500	500	0.0323	0.0004	1.22	0.0373	0.0007	2.12	0.0407	0.0009	2.12

If the CSM interaction is 10 times weaker than that of SN 2015cp, the decaying SN Ia tail generally dominates over the interaction signature and the light curve shows little deviation from a normal decaying SN Ia tail. Even in the best-case scenario of the longest and latest CSM interaction simulation, the 50% recovery threshold lies at $z_{50} = 0.0050 \pm 0.0005$ ($d \approx 22 \pm 2$ Mpc). In the simulations where the interaction is 10 times stronger compared to that of SN 2015cp, the shortest and earliest interaction (lasting from 100 to 200 days after peak) has $z_{50} = 0.0091 \pm 0.0014$ ($d \approx 39 \pm 6$ Mpc). For the longest and latest interaction, the 50% recovery rate is at $z_{50} = 0.0323 \pm 0.0004$ ($d \approx 139 \pm 2$ Mpc).

3.2.4.4 Impact of reference image depth

The mean limiting magnitude of the ZTF reference images is ~ 21.8 mag, and as discussed in (Strotjohann et al., 2021), this is the limiting factor for recovering faint signals from binned light curve data. To test the improvement of deeper reference images, the assumed limiting magnitude was changed to be 0.5 and 1 mag deeper. The recovery fraction for one magnitude deeper is shown for comparison in figure 3.5. As expected, deeper reference images allow the interaction signatures to be detected to higher redshift, although the increases in z_{50} values are modest (see Table 3.3). For example, for the latest onset and longest interaction duration interaction, z_{50} increases from $z_{50} = 0.0323 \pm 0.0004$ to $z_{50} = 0.0407 \pm 0.0009$ ($d \approx 139 \pm 2$ Mpc to $d \approx 174 \pm 4$ Mpc).

3.3 Results

I run my custom detection pipeline on the ZTF DR2 light curves in the same way as it was performed on the simulated light curves in section 3.2.4. In 1932 light curves nothing is detected in any of the 16 trials discussed in section 3.2.1, in 432 light curves the late-time detections are attributed to declining SN Ia tails, and in 1020 light curves the detections were not considered robust (< 4 successful trials). These are the three largest cuts in my sample, as can be seen in Table 3.1, and leave me with 134 light

curves that pass the pipeline. In section 3.3.1, I describe the light curves of these events and discuss how some light curves fit into known classes of events (e.g. known Ia-CSM, late-time SN Ia tail detections). In section 3.3.2, I describe the additional tests that were performed on the remaining promising ten events to determine if their late-time excesses are due to CSM interaction or other scenarios.

3.3.1 Initial summary of detected events

As can be seen in Table 3.1, the result of the pipeline is a list of 134 objects that require visual inspection after passing the detection cuts of positive, $> 5\sigma$ detections in adjacent light curve bins in at least four of the 16 bin size and placement combinations. In 47 of these cases, by visual inspection I identified that an incorrect baseline caused false positives. In five cases, the peak date estimation failed and estimated the peak to be over 100 days before the actual SN explosion. Because of this, the SN itself was detected as a late-time signal. Furthermore, in 29 cases there was evidence of the host galaxy being improperly subtracted or showing signs of activity, which interfered with the forced photometry at the SN location. Finally, in 20 cases the tail fit test was unable to show the nature of the tails due to various reasons (e.g., the fits did not converge properly or there was a gap in the observations while the tail was visible). After this step, 33 objects were remaining in the sample.

Further details of these remaining 33 SNe Ia are shown in Table 3.4 and can be split up into four main groups: i) known Ia-CSM events, ii) transient siblings, where a second transient event occurs near the identified SN Ia causing its light to (partially) be picked up during forced photometry at the first SN location, iii) nearby objects whose tail could not be fitted by my simple model, and iv) objects that do not fall in the first three groups. In the following sections, I describe the first three of these groups that are not due to potential CSM interaction at late times.

Table 3.4: List of objects that passed the initial visual inspections.

Name	IAU name	Redshift	Type ^(a)	Peak MJD	Peak mag.	Excess phase (d)	Excess band	Excess mag.	Group ^(b)
ZTF18aaykjei	SN 2018crl	0.09690 ± 0.00002	Ia-CSM	58297.3	18.40 ± 0.04	100 – 400	<i>r</i>	20.1 – 21.9	Known Ia-CSM
ZTF18abuatfp	SN 2018gkx	0.13643 ± 0.00001	Ia-CSM	58382.1	18.71 ± 0.05	100 – 475	<i>gri</i>	19.6 – 21.5	Known Ia-CSM
ZTF18actuhrs	SN 2018evt	0.02442 ± 0.00001	Ia-CSM	58476.5	16.19 ± 0.01	100 – 525	<i>gri</i>	16.6 – 21.8	Known Ia-CSM
ZTF19aaoqst	SN 2019agi	0.05958 ± 0.00001	Ia-CSM	58511.5	18.46 ± 0.04	100 – 450	<i>gri</i>	19.2 – 21.9	Known Ia-CSM
ZTF19abidbqp	SN 2019ibk	0.04014 ± 0.00001	Ia-CSM	58688.5	18.56 ± 0.05	100 – 1125	<i>gr</i>	19.1 – 21.5	Known Ia-CSM
ZTF19acbjddp	SN 2019rvb	0.1832 ± 0.0004	Ia-CSM	58790.1	18.90 ± 0.06	100 – 400	<i>gr</i>	20.4 – 22.0	Known Ia-CSM
ZTF20aatxryt	SN 2020eyj	0.0294 ± 0.0004	Ia-CSM	58939.2	17.24 ± 0.01	100 – 450	<i>gr</i>	19.0 – 21.8	Known Ia-CSM
ZTF20abbbsfs	SN 2020kre	0.13530 ± 0.00001	Ia-CSM	58998.2	19.09 ± 0.04	175 – 425	<i>gr</i>	19.8 – 21.4	Known Ia-CSM
ZTF20abmlrxr	SN 2020onv	0.0940 ± 0.0004	Ia-CSM	59052.4	17.85 ± 0.01	100 – 500	<i>gr</i>	18.8 – 21.6	Known Ia-CSM
ZTF20abqkbfz	SN 2020qzx	0.0968 ± 0.0004	Ia-CSM	59094.3	18.18 ± 0.04	100 – 450	<i>gri</i>	19.5 – 21.9	Known Ia-CSM
ZTF20acmutv	SN 2020uem	0.043 ± 0.001	Ia-CSM	59173.5	16.38 ± 0.01	100 – 525	<i>gr</i>	17.4 – 21.4	Known Ia-CSM
ZTF20aciwcuz	SN 2020xtg	0.06122 ± 0.00001	Ia-CSM	59189.5	17.45 ± 0.02	100 – 500	<i>gri</i>	18.0 – 21.9	Known Ia-CSM
ZTF20acyroke	SN 2020aeh	0.12665 ± 0.00004	Ia-CSM	59217.4	19.01 ± 0.06	100 – 250	<i>r</i>	19.8 – 20.9	Known Ia-CSM
ZTF18aasdted	SN 2018big	0.01814 ± 0.00001	Ia-norm	58268.4	15.64 ± 0.01	450 – 550	<i>gr</i>	20.4 – 21.5	Sibling
ZTF19aaysiwt	SN 2019hnt	0.0926 ± 0.0004	Ia	58651.2	18.46 ± 0.05	525 – 625	<i>gr</i>	20.5 – 21.7	Sibling
ZTF19acihfz	SN 2019tjz	0.055 ± 0.003	Ia-norm	58795.1	18.00 ± 0.03	950 – 1050	<i>r</i>	19.4 – 20.7	Sibling
ZTF20abzetdf	SN 2020tft	0.071 ± 0.002	Ia-norm	59113.5	18.2 ± 0.1	725 – 800	<i>r</i>	19.4 – 20.2	Sibling
ZTF20acehyxd	SN 2020uud	0.0346 ± 0.0005	Ia-norm	59129.3	18.72 ± 0.03	300 – 350	<i>r</i>	20.2 – 21.5	Sibling
ZTF19aatlmbo	SN 2019ein	0.0072 ± 0.0001	Ia-norm	58617.2	blinded ^(c)	100 – 425	<i>r</i>	18.9 – 21.8	Kinky tail
ZTF20abqvsik	SN 2020rcq	0.00246 ± 0.00001	Ia-norm	59144.5	blinded ^(c)	100 – 400	<i>i</i>	16.5 – 22.2	Kinky tail
ZTF20abrjmgj	SN 2020qxp	0.00356 ± 0.00001	Ia-91bg	59088.1	blinded ^(c)	100 – 375	<i>r</i>	17.9 – 21.8	Kinky tail
ZTF20abwrcmq	SN 2020sck	0.01643 ± 0.00001	Iax	59099.4	16.25 ± 0.01	100 – 450	<i>gri</i>	19.1 – 21.9	Kinky tail
ZTF20achlced	SN 2020uxz	0.00867 ± 0.00008	Ia-norm	59142.4	blinded ^(c)	100 – 400	<i>gr</i>	17.1 – 21.9	Kinky tail
ZTF19acwrqtv	SN 2019vzf	0.059 ± 0.004	Ia-99aa	58829.1	18.00 ± 0.07	150 – 400	<i>gr</i>	20.7 – 21.8	Other – AGN ^(d)
ZTF20aahptds	SN 2020awr	0.07649 ± 0.00002	Ia-norm	58888.5	18.54 ± 0.07	300 – 1050	<i>i</i>	20.2 – 20.5	Other – data issue ^(d)
ZTF20aazwuin	SN 2020kzd	0.082 ± 0.001	Ia-norm	58997.4	18.86 ± 0.03	250 – 850	<i>gr</i>	21.3 – 22.0	Other – data issue ^(d)
ZTF18abtqevs *	SN 2018grt	0.042 ± 0.003	Ia-norm	58372.3	18.56 ± 0.02	1350 – 1450	<i>r</i>	21.1 – 21.3	Other
ZTF19aanyuyh	SN 2020pkj	0.02456 ± 0.00001	Ia-norm	59060.4	16.90 ± 0.01	100 – 175	<i>r</i>	20.8 – 21.1	Other
ZTF19abfvhlx *	SN 2019ldf	0.05646 ± 0.00001	Ia-norm	58686.5	17.90 ± 0.04	1050 – 1225	<i>ri</i>	20.1 – 21.1	Other
ZTF19ablekwo	SN 2019mse	0.088 ± 0.004	Ia-norm	58715.4	18.37 ± 0.02	450 – 700	<i>gri</i>	19.8 – 20.6	Other
ZTF19abzwtiu	SN 2019rqn	0.075 ± 0.003	Ia-norm	58760.3	18.62 ± 0.03	950 – 1050	<i>i</i>	20.6 – 21.4	Other
ZTF20aaiifyfx	SN 2020ahm	0.06001 ± 0.00001	Ia-norm	58873.5	18.10 ± 0.02	750 – 1025	<i>gri</i>	19.9 – 21.9	Other
ZTF20abjfufv *	SN 2020tfc	0.031 ± 0.001	Ia-norm	59116.3	17.22 ± 0.01	550 – 800	<i>gri</i>	18.9 – 21.4	Other

* The three final objects with suggested detections of late-time interaction.

^(a)Type is based on spectral classifications from [Dimitriadis et al. \(2025\)](#) and [Burgaz et al. \(2025a\)](#), Ia-CSM are those interacting with CSM that is observed around peak, Ia-norm are those that are most consistent with a normal SN Ia, Ia are those without a sub-classification but are consistent with being a SN Ia, and SN 2020sck is a Iax ([Dutta et al., 2022](#)).

^(b)The objects are split into four groups: Ia-CSM that were previously known, identified siblings, SNe Ia that are detected due to their early-time tail fits showing a kink (Kinky tail), and the ‘Other’ category that includes ten events with potential flux excesses, including three (SN 2019vzf, SN 2020awr and SN 2020kzd) that are subsequently ruled out (see section 3.3.2).

^(c)The peak magnitudes of four SNe Ia in the sample are blinded because of their planned use in H_0 constraints ([Rigault et al., 2025a](#)).

^(d)SN 2019vzf is ruled out as a true excess due to AGN variability at the position, in SN 2020awr the excess is only detected in the *i* and is ruled out by the scene modelling analysis, and SN 2020kzd is detected in the *gr*-bands on a complex galaxy environment and not detected in the scene modelling analysis (see section 3.3.2.1).

3. ZTF SN Ia DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY




Figure 3.6: Binned late-time observations of the recovered known SNe Ia-CSM. All objects are shown in absolute magnitude and over the same time range for easy comparison. All objects are detected beyond 100 days after the peak without using the binning technique. I do not show these individual data points to increase readability. The tail fits are shown as solid lines with the hashed region denoting their 1σ uncertainties. For comparison, SN 2020ssf (ZTF20abytpc) in the top left corner is a normal SN Ia with a normally declining tail with $t_{1/2,g} = 53 \pm 1$ days and $t_{1/2,r} = 26 \pm 1$ days. The fitted tails for the SNe Ia-CSM are significantly shallower.

3.3.1.1 Known Ia-CSM

The first group are the 13 known Ia-CSM: those objects that already had a Ia-CSM classification. These objects started interacting relatively soon after the explosion and remained active long enough to be picked up by my pipeline beyond the 100 day threshold. Figure 3.6 shows the light curves of the recovered SNe Ia-CSM in absolute magnitude (uncorrected for extinction). Even if the peak identified by my code is not the real peak due to it not being observed (e.g., for SN 2018evt, SN 2019agi, and SN 2019ibk), the CSM interaction persists for long enough to be picked up by my pipeline.

Ten of the known Ia-CSM SNe are presented in [Sharma et al. \(2023\)](#), who search for SNe Ia-CSM discovered in the ZTF Bright Transient Survey from May 2018 to May 2021 (BTS; [Fremming et al., 2020](#); [Perley et al., 2020](#)). They find two objects that are not in my sample: SN 2020abfe (ZTF20acqikeh) and SN 2020aekp (ZTF21aaabwzx). SN 2020abfe is in the DR2 sample, but due to a combination of a gap in the observed light curve and the interaction not altering the declining tail sufficiently, the tail fit procedure is unable to distinguish it from a normal declining SN Ia tail. SN 2020aekp was first detected after the final date for objects to be included in my sample. Two of the events in my sample (SN 2020eyj and SN 2020kre) are not presented in [Sharma et al. \(2023\)](#). SN 2020eyj was excluded as [Sharma et al. \(2023\)](#) focussed on interaction with H-rich material, and this object showed He emission lines suggesting interaction with He-rich material ([Kool et al., 2023](#)). SN 2020kre is not in the BTS sample and, therefore, not included in [Sharma et al. \(2023\)](#). However, it was confirmed with spectroscopy to have $H\alpha$ emission in its peak spectra.

Out of the 13 events, SN 2020aeuh is an outlier due to its distinct light curve. While the other 12 known Ia-CSM events detected in my sample have decline tails whose slopes are significantly shallower than for a normal SN Ia or steepen over time, SN 2020aeuh brightens significantly, having a double-peaked nature with the second peak at around 100 days after the first. Even though the light curve suggests the SN to be interacting, no H emission lines ($H\alpha$ or other) are observed. Kool et al. (in prep)

present a detailed analysis of this object.




Figure 3.7: Binned late-time observations in flux space of the five events with a detected sibling, with the flux normalised to the found peak flux. All objects are plotted on the same flux scale for easy comparison, except for SN 2018big, as its late-time detections are much weaker compared to the original SN peak magnitude due to the larger distance offset between the siblings.

3.3.1.2 Siblings

Siblings are transients that occur in the same host galaxy, and can be useful for understanding differences in local environments (e.g., [Biswas et al., 2022](#); [Graham et al., 2022](#)). In some cases, the siblings occur in (almost) the same place in the sky, only differing in explosion time. This can be either due to the two transients being physically close together, or a projection effect due to the inclination of the host. However, the result is the same: forced photometry at the location of one sibling will result in a (partial) recovery of the other. Assuming that the first transient is a SN Ia in my sample and the second transient is fainter, my pipeline will flag the late-time rebrightening as a late-time excess in one of my objects.

Careful examination of the images using SNAP and cross-referencing using Fritz (an alert broker, [van der Walt et al. 2019](#); [Duev et al. 2019](#); [Kasliwal et al. 2019](#); [Coughlin](#)




Figure 3.8: Light curves of SN 2018big and its sibling SN 2019nvm in magnitude space using bins of 25 d. The g (green) and r (orange) bins follow the tail of SN 2018big until it disappears in the noise. About 450 d after the peak of SN 2018big, new detections are identified in the binned photometry. The individual observations remain upper limits, but their shape hint at the true nature of these late-time detections.

Table 3.5: Objects with a detected sibling transient.

Primary name	IAU name	Sibling name	IAU name ^(a)	Type	Date ^(b)	Offset (") ^(c)
ZTF18aasdted	SN 2018big	ZTF19abqhobb	SN 2019nvm	IIP	58 714	3.7
ZTF19aaysiwt	SN 2019hnt	ZTF20acwpads	-	-	59 186	1.2
ZTF19acihfxz	SN 2019tjz	ZTF18aanhpii	-	-	59 774	1.2
ZTF20abzetdf	SN 2020tft	-	-	Ia	59 867	< 1
ZTF20acehyxd	SN 2020uvd	ZTF21abouuow	SN 2021udv	Ia	59 422	3.1

^(a) The siblings ZTF19aaysiwt and ZTF20acwpads share an IAU name. For ZTF20abzetdf, the siblings are too close together ($< 1''$ separation) to be automatically recognised as separate events, causing them to share both ZTF and IAU names. ZTF18aanhpii is a sibling transient in 2022 on top of the host nucleus, resulting in the internal name being from 2018.

^(b) MJD of the first detection of the sibling.

^(c) Angular separation on the sky between the siblings.

et al. 2023) and the Transient Name Server (TNS)¹ showed that there are five objects in my shortlist whose late-time detections are due to a sibling. In all five cases the sibling is bright enough to be detected in individual difference images, which resulted in an alert from the ZTF pipeline for each image where something was detected. In contrast to forced photometry where the PSF fit is forced at a given location, the alert photometry attempts to find the best location for the PSF fit in each image, resulting in a position measurement in each image. By averaging over these measurements, the SN position can be found with sub-pixel precision. In some cases, such as SN 2018big and SN 2019nvm, these points cluster around two locations, showing that the two events are spatially apart. In other cases, such as SN 2020tft and its sibling, both events were spectroscopically classified near their peak, showing that the type and / or phase at the time of the observations cannot match with a single event.

Figure 3.7 shows the binned light curves of these objects in flux space. In each light curve there is a sudden significant spike in the detected fluxes in all observed bands, which falls back down again after a short period of time. Table 3.5 lists the name and type of each sibling if known, as well as their sky separation. In some cases the siblings are close enough together that they are not automatically recognised as separate events, resulting in them having the same name. In the case of SN 2019tzj the sibling (ZTF18aanhpii) exploded close to the nucleus, which had some spurious detections in 2018. This caused the sibling to have a 2018 ZTF name, even though it exploded in 2022.

Figure 3.8 shows the detection of a sibling (SN 2019nvm) in the late-time light curve of SN 2018big in magnitude space. The photometry pipeline forces the point spread function (PSF) fit at the position of SN 2018big. As the position of SN 2019nvm is slightly offset ($\sim 4''$), only some of the total flux of SN 2019nvm is captured in the fit. Besides these five siblings, I identified two other pairs of siblings with SNAP: SN 2019gcm and SN 2021fnj, and SN 2020jgs and SN 2021och. These siblings were too far apart to be picked up with the forced photometry ($9.6''$ and $10.8''$, respectively) but

¹<https://www.wis-tns.org/>

were found while inspecting using SNAP. For a complete list and study on the siblings found in the ZTF DR2, I refer the reader to [Dhawan et al. \(2024\)](#).

3.3.1.3 Kinky tails

This group consists of five objects where the simple tail model, based on the typical decline rate of SNe Ia based on SN 2011fe ([Dimitriadis et al., 2017](#)), failed to fit the observations at later times. The reason for this failure in four of the events is that there is a slow-down in the r - and i -band decline rates at $\sim 200 - 250$ days after the peak, which the model does not take into account. The fifth event, SN 2020sck ([Dutta et al., 2022](#)), is a known SN Iax, a subclass known for having lower ejecta velocities and luminosities, suggesting that the explosion did not necessarily fully disrupt the WD ([Jordan et al., 2012](#); [Kromer et al., 2013](#)). This event was flagged because of a slow-down in its decline rate roughly 80 days after the identified peak. The change in slope is visible in all bands and significantly longer than the assumed $t_{1/2} = 50$ d of normal SNe Ia. The presence of a bound remnant has been suggested to be the cause of similar late-time signatures seen in other SN Iax ([Kawabata et al., 2018](#); [McCully et al., 2022](#); [Camacho-Neves et al., 2023](#)).

The four SNe that deviate from the simple tail model are very nearby (< 39 Mpc) compared to the majority of the ZTF DR2 sample, and when using the binned observations they are bright enough to be detected up to (nearly) a year after their first detection. As [Rigault et al. \(2025a\)](#) are using these nearby events to calibrate their H_0 measurement, their peak magnitudes are currently blinded. Figure 3.9 shows the g -, r -, and i -band light curves of these objects in magnitude space. The R -band light curve of SN 2011fe ([Zhang et al., 2016](#)), which had a similar change in decline slope, is also shown for comparison. As discussed in [Dimitriadis et al. \(2017\)](#), the radioactive decays produce γ -rays and positrons, as well as X-rays and electrons that can be thermalised, depositing their energy in the expanding SN ejecta. As the ejecta expand over time they become more transparent, shortening the delay between thermalisation of the deposited energy and the emission of optical radiation. This results in the SN

tail slope changing as the opacity changes.




Figure 3.9: Five objects with kinks in their tails that start to deviate from the assumed decline rate at $\sim 200 - 250$ days after the peak are shown in magnitude space as a function of days since the peak. As most of these objects have their peak magnitude blinded, no scaling is shown. A normal radioactive decay model was fitted to these tails, shown as solid straight lines with their 1σ uncertainty as dashed regions. But as the ejecta opacity changes over time so does the half-life time of the tail, causing a kink seen in the bins which is not reproduced by the model. The arbitrarily normalised R -band light curve of SN 2011fe (known not to have CSM interaction from detailed spectral studies) (Zhang et al., 2016) is shown in blue, showing the same shift in decline slope at a similar phase.

In order to be able to observe a change in decline slope such as this, a SN has to be both bright and well observed during the time it is visible. The other SNe Ia in my sample at similarly low redshifts have gaps in their observations or the change in slope is not strong enough for the tail fits to fall outside the allowed range of reduced chi-squared values ($\chi^2_{\text{red}} > 5$) and therefore, are not flagged by the pipeline due to this.

3.3.2 Additional tests of promising events

After performing the tests discussed in the previous sections on each event, there are ten objects remaining with an unexplained late-time excess. Light echoes produced by SN light scattering off of nearby dust clouds were considered but were ruled out as these are typically ≥ 10 mag fainter than the SN at its peak (Patat, 2005; Graur et al., 2016).

Table 3.6: Results of the additional tests for the ten promising objects. The host separation is given in " and converted to kpc using the redshift given in Table 3.4. In the final five columns the similarity of each SN Ia with a late-time excess is compared to different transient classes to see if the late-time signal can be interpreted as another transient. This can be excluded based on an inconsistent colour (1), duration (2), and/or an excessive amount of host extinction (3) needed to obtain the observed magnitudes.

Name	Red. error ^(a)	Host separation (")	$E/(B - V)_{\text{host}}$ (mag.) ^(b)	AGN ^(c)	Approx. mag. ^(d)	SN Ia?	Ib?	Ic?	IIP?	TDE?
SN 2018grt ^(*)	no	0.36 ± 0.03	0.32 ± 0.02	$0.21 - 0.36$	no	no (23)	no (23)	no (2)	no (3)	no (13)
SN 2020pkj	no	0.52 ± 0.02	0.28 ± 0.01	$0.23 - 0.36$	no	yes	yes	yes	no (23)	no (13)
SN 2019ldf ^(*)	no	0.65 ± 0.04	0.78 ± 0.05	≤ 0.03	no	no (123)	no (123)	no (123)	no (123)	no (13)
SN 2019mse	no	0.59 ± 0.05	1.09 ± 0.10	0	no	no (2)	no (2)	no (2)	no (2)	yes
SN 2019rqn	no	1.39 ± 0.08	2.21 ± 0.13	≤ 0.04	no	no (13)	yes	yes	no (1)	no (13)
SN 2019vzf	no	3.89 ± 0.08	4.86 ± 0.10	0	yes	-	-	-	-	-
SN 2020awr	yes	19.60 ± 0.04	31.65 ± 0.07	-	-	-	-	-	-	-
SN 2020alm	no	0.67 ± 0.07	0.85 ± 0.09	≤ 0.05	no	no (123)	no (12)	no (23)	no (123)	yes
SN 2020kzd	yes	4.67 ± 0.04	8.07 ± 0.07	-	-	-	-	-	-	-
SN 2020tfc ^(*)	no	0.21 ± 0.02	0.14 ± 0.01	≤ 0.12	no	no (123)	no (12)	no (23)	no (123)	no (3)

(*)My three final objects with suggested detections of late-time interaction.

^(a)Reduction error: Comparison between the standard light curve reduction and scene modelling identified issues with the baseline correction (see section 3.3.2.1).

^(b)Estimated assuming the main peak is a normal SN Ia. When correcting for Milky Way extinction and distance is enough to exceed an absolute g -band magnitude of -19.3 , I quote a host $E/(B - V) = 0$ mag.

^(c)The presence of an AGN was estimated using the WISE colours of the host and the criteria of Hviding et al. (2022).

^(d)Mean absolute magnitude of the r -band late-time excess after correction for Galactic and host extinction, averaged over the host extinction range considered.

I performed several additional tests on these ten events to try to determine the origin of their late-time excess. The first is using SMP as described in section 3.1.1. The second is testing for coincidence with an AGN and third is more detailed comparisons with known transient classes. These tests are discussed below and a summary is shown in Table 3.6.

3.3.2.1 Scene modelling photometry

To test for issues in data processing, the ten events were run through the scene modelling pipeline of Lacroix et al. (in prep.) by Leander Lacroix as discussed in section 3.1.1. Two SNe Ia (SN 2020awr and SN 2020kzd) were found to have issues with their difference imaging and forced photometry light curves, causing false detections. The detections in both these events are close to the detection limit and, therefore, are impacted by even small errors in the baseline placement. The difference imaging (left) and scene-modelling (right) light curves of these two events are shown in figure 3.10.

For SN 2020awr, ~ 300 days after the SN peak, the i -band observations jump up to detections at ~ 20.3 mag. In contrast to other objects where baseline issues were found, here the offset occurs only for part of the light curve, while the pre-SN baseline has no visible issues. I could not identify a clear reason for this when inspecting the images with SNAP, and the SN is too far away from the host nucleus for it to be host activity. For the scene-modelling version of the photometry, the jump in the i -band observations has disappeared completely, showing that there was indeed an unidentified issue with the i -band data for this object. Since the late-time detections are determined to be spurious, this object is ruled out from having late-time detections.

In the case of SN 2020kzd, late-time detections are present in the g - and r -bands for hundreds of days (see figure 3.10). The SN is in a complex environment with three galaxies close to its sky position, which likely complicates the image subtraction and baseline correction. When SMP is performed on the events, the detections disappear and the average flux at late times is consistent with zero, showing that the binned forced photometry light curve likely suffered from a wrongly determined baseline correction.




Figure 3.10: Two objects whose late-time detections were revealed to be caused by the photometry extraction. The colours are as in figure 3.9. The left side shows the binned forced photometry light curve, and the right side shows the binned SMP light curve. Bins with 5σ detections are shaded solid, while the non-detections are hashed.

3.3.2.2 AGN contamination

If a SN explosion site is coincident with a host galaxy that has an AGN, host activity is a likely cause of the late-time detections. Using data from the *Wide-field Infrared Survey Explorer* (WISE; Wright et al. 2010), Hviding et al. (2022) present a criterion to test if a galaxy hosts an AGN based on the WISE $W1-W2$ and $W2-W3$ colours, which I apply to my events. One object (SN 2019vzf) is 4.86 kpc off its host centre, and its host is a known AGN with WISE colours of $W1-W2 = 0.55 \pm 0.03$ and $W2-W3 = 2.90 \pm 0.04$ mag. In SNAP, the late-time signal appears to cover both the SN and AGN locations. The AGN contamination is too strong to put any meaningful constraints on the late-time flux at the SN location. I kept the object in my sample until now to test if it is possible to use scene modelling to reduce the AGN contamination. However, this is not possible. I attribute the late-time signal to host activity and disregard it in future discussions. The forced photometry and SMP light curves of SN 2019vzf are shown in figure 3.11.

3.3.2.3 Presence of a sibling close to the SN location

To test if a previously unidentified sibling transient is causing the late-time detections of the remaining seven objects, I compared their late-time light curves to known

3. ZTF SN IA DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY




Figure 3.11: Objects whose late-time detections are explained in the additional tests. The top four rows show the light curves of the objects where a previously undetected sibling transient as an explanation for the late-time observations could not be ruled out, with forced photometry light curves in the left-hand panels and the scene modelling photometry light curves in the right-hand panels. The 5σ detections are shown as bins with solid uncertainty regions, and bins with hashed uncertainty regions are non-detections. The object whose late-time detections are caused by the host galaxy AGN, SN 2019vzf, is shown in the bottom row. The colours are as in figure 3.9 with g -band in green, r -band in orange, and i -band in red.

classes of transients, including a SN Ia, core-collapse SNe (Type Ib, Type Ic, Type IIP) and two tidal disruption events (TDEs). First I have estimated the amount of potential host extinction from the main SN peak by assuming it was a normal SN Ia with a typical g -band peak of -18.8 to -19.3 mag after correcting for the distance to the SN and for Milky Way extinction. These estimated host extinction values are given in Table 3.6, assuming $R_V = 3.1$. The bright end of the absolute peak magnitude gives an upper limit for the host extinction, and the faint end gives a lower limit. After correcting for this range of additional host galaxy extinction, the late-time excesses have mean absolute r -band magnitudes of -15.4 to -17.5 mag (see Table 3.6).

Table 3.7: Details of the comparison transients used to test if late-time detections could be explained by another transient at a similar sky position. The first column shows the assumed type of transient, the second shows the transients used to represent each type, the third has the approximate absolute extinction-corrected r -band peak magnitudes, and the fourth is the reference for each event.

Type	Name	Peak abs. M_r (mag.)	Reference
SN Ia	SN 2011fe	$-18.4 - -19.4$	Mazzali et al. (2014)
SN Ib	SN 2019yvr	-17.9	Kilpatrick et al. (2021)
SN Ic	SN 2021krf	-17.3	Ravi et al. (2023)
SN IIP	SN 2020jfo	-17.8	Ailawadhi et al. (2023)
SN IIP	SN 2017gmr	-18.7	Andrews et al. (2019)
TDE	AT 2018hco	-22.1	van Velzen et al. (2021)
TDE	AT 2018zr	-20.1	van Velzen et al. (2021)

After estimating the allowed extinction for each primary SN Ia, I initially assumed that if the late-time excess is due to another transient, then it will have the same extinction along the line-of-sight. For these other transients, I used examples of a SN Ia, Ib, Ic, two IIPs, and two TDEs to compare against, with details of the comparison objects described in Table 3.7. The transients chosen to represent their category are all in the typical magnitude range for their type. Two TDEs and two SN IIPs were chosen to represent the upper and lower end of the range of peak magnitudes expected for these transient types. I use my model for SN 2011fe to represent SNe Ia and test the lower and upper edge of the range of normal SN Ia peak r -band magnitude.

These comparison objects were chosen as they have well-sampled light curves in the ZTF filters and literature values for the host extinction. I correct the light curves of the comparison objects using their literature redshifts and their extinction values before correcting for the redshift and extinction of each SN Ia in my sample with a potential late-time excess. I then compare this transient light curve to the found late-time detections to see how well they match. A good match will have a similar magnitude, colour, and duration.

It could be the case that the suspected sibling was in the same line-of-sight direction, but had a different amount of extinction due to, for instance, exploding behind a cloud that adds additional extinction. I check this by adding enough extinction to match the r -band detections between the different comparison events and the observed late-time

detection and again check if the colour and duration match up, as the observed colour is affected by the extinction. For the TDE comparisons, I allow a host galaxy $E(B - V)$ of up to one magnitude, as was estimated for the ZTF TDE sample of [Hammerstein et al. \(2023\)](#).

For four events (SN 2019mse, SN 2019rqn, SN 2020alm, and SN 2020pkj), the late-time detections are consistent with at least one of the comparison classes, as detailed in Table 3.6 and shown in figure 3.11. I describe them individually in the following sections. The three remaining events (SN 2018grt, SN 2019ldf, and SN 2020tfc) cannot be explained by the presence of a sibling transient and are discussed further in section 3.3.2.4.

SN 2019mse. SN 2019mse has late-time detections starting at 450 days after the peak and lasting ~ 250 days in the *gri*-bands with an absolute *r*-band magnitude of -17.5 mag during the excess. Careful re-examination of the difference images shows that the late detections are slightly offset (about one pixel) from the SN location and appear to be on the host nucleus location instead. However, with its WISE colours being $W1 - W2 = 0.25 \pm 0.04$ and $W2 - W3 = 2.44 \pm 0.09$, the host is determined to not contain an AGN. The scene modelling version of the light curve shows a similar late-time excess, showing that this is not an artefact from the chosen photometric analysis method.

The late-time signal is detected in all three bands, and its behaviour is very similar in all of them (see figure 3.11). Its rise and decline time scales, absolute magnitude, and colours agree well with the ranges seen for TDE. Together with the observation that the late-time detections are at the host nucleus location, this suggests that a nuclear transient explains the late-time signal adequately.

SN 2019rqn. In the case of SN 2019rqn, there is a short period of detections at 950 – 1 050 days in the *r*-band after a gap in the observations, declining and fading below the detection significance within 100 days of the first detection. Nothing is detected

at a $\geq 5\sigma$ level in the g - or i -band observations. The i -band SMP light curve had its host contribution not completely subtracted, causing a flux offset in the data points. I therefore do not consider the SMP i -band further. The data points in the g - and r -band light curve have slightly larger uncertainties in the SMP version, causing the main SN light curve to fall below the 5σ threshold at an earlier epoch, resulting in fewer individual points visible (figure 3.11). Similarly, larger uncertainties for the SMP prevent the detection of a late-time r -band signal.

Assuming that a sibling exploded during the gap in the observations between 870 and 950 days after the peak, both a SN Ib and SN Ic with $E(B - V) \leq 0.3$ mag extinction in the r -band can fit their tail to match the observed detections in the r -band without being excluded by the g -band and i -band non-detections. Therefore, I cannot rule out a sibling as the source of the detected late-time signal, and remove it from my sample.

SN 2020alm. The late-time signal in SN 2020alm is seen in all three bands, beginning at ~ 750 days after the peak and lasting for at least 300 days. There is a gap of 80 days in the observations immediately before the period of activity. The detections slowly rise to a plateau. My initial analysis only included data for SN 2020alm up to $\sim 1\,000$ days after the peak. However, when this object was identified as having late-time detections, the light curve pipeline was rerun and it was found to be still bright at later times, with significant detections in the r - and i -bands, but no $\geq 5\sigma$ detections in the g -band. The i -band SMP light curve contained a significant flux offset, as the host was not fully subtracted. I therefore do not further consider this band. The binned SMP r -band light curve fails to reproduce the late-time detections found in the forced photometry light curve. However, the g -band detections are recovered in the SMP. The SN is close to the host nucleus at $0.66''$ (0.85 kpc) offset at the redshift of the SN, but the WISE colours of $W1 - W2 = 0.06 \pm 0.04$ and $W2 - W3 = 2.42 \pm 0.12$ place it far outside the AGN region.

As SN 2020alm was still active while my analysis was ongoing, Two spectra (which have been stacked) were obtained by Dr. Lluís Galbany using the Optical System for Imaging and low-Intermediate-Resolution Integrated Spectroscopy (OSIRIS+) instrument on the Gran Telescopio CANARIAS (GTC) at Roque de los Muchachos in La Palma on 26 July 2023 using the R1000R grism, 1277 days after the estimated peak date of the SN. As the spectrum is heavily dominated by the host galaxy, I subtracted a host galaxy spectrum taken by the Sloan Digital Sky Survey (SDSS, [York et al. 2000](#); [Adelman-McCarthy et al. 2006](#); [Gunn et al. 2006](#); [Smee et al. 2013](#)) in 2003, to remove the host contamination. This was done after re-sampling the host spectrum to have the same wavelength spacing as the new spectrum. This left only the spectrum causing the late-time photometry detections and is shown in figure 3.12. I confirmed successful host subtraction by checking for residual Na ID and Mg I $\lambda 5175$ absorption lines, as these are not expected to be due to the late-time signal but purely from the host galaxy. I found that no residual features of these lines were present.

The resulting spectrum shows an excess that is stronger towards longer wavelengths. This is consistent with the broadband photometry finding a brighter excess in the redder bands, while the g -band remains within the noise after binning the observations. There is some excess in the narrow [N II] $\lambda\lambda 6548, 6583$, $H\alpha$ and [S II] $\lambda\lambda 6716, 6730$ emission lines, but I lack the resolution to check if this is significant. There is no visible additional $H\alpha$ component in the spectrum, which would be indicative of CSM interaction. Integrating the spectrum over the r - and i -band efficiencies gives a colour of $r - i = 0.6$ mag, which is within 3σ of the value found in the latest photometry bin.

[Hammerstein et al. \(2023\)](#) show that TDEs generally have a $g - r$ colour of zero and can have featureless spectra. I approximate a TDE by a line of constant flux density and add Milky Way extinction (using the SFD89 dust maps in the direction of the object; [Schlegel et al., 1998](#)) in order to estimate the amount of extinction needed to generate a red excess similar to the spectrum. I find that the general shape of the spectrum can be approximated with $0.6 < E(B - V)_{\text{host}} < 1$ mag. This would mean an absolute r -band magnitude of $-18.8 > M_r > -19.8$ mag. [Hammerstein et al. \(2023\)](#)




Figure 3.12: Spectrum of the late-time signal in SN 2020alm in its rest frame. The top panel shows the late-time spectrum obtained on 26 July 2023 using OSIRIS+ on the GTC and the SDSS spectrum obtained in 2003. The bottom panel shows the late-time excess, obtained by subtracting the SDSS host galaxy spectrum from the observed late-time spectrum. A smoothed spectrum is shown in blue. The smoothing was done using a rolling kernel of size 5 to average over the values. The red lines are a simple TDE model with Milky Way and some amount of host galaxy extinction applied (the amount is shown in the legend) in order to get the approximate shape of the observed spectrum. Narrow emission and absorption lines that were notable in the unsubtracted spectrum are marked with dashed lines. The grey regions are affected by sky lines and should be ignored. The late-time spectrum of SN 2020alm is available in electronic form at the CDS via anonymous ftp to [cdsarc.u-strasbg.fr](ftp://cdsarc.u-strasbg.fr) (130.79.128.5), via <http://cdsweb.u-strasbg.fr/cgi-bin/qcat?J/A+A/> or upon request to me.

show that both this amount of host extinction and late-time brightness are possible for TDEs. They also show that it is possible for a TDE to rise and fall back down within the 80 day gap in observations, although this is seen in fainter TDEs than corresponding to my estimated absolute magnitude range.

The TDE sample of [Hammerstein et al. \(2023\)](#) did not contain a single object that matches my late-time detections in duration and luminosity in SN 2020alm. However, it is possible to combine parts of different TDEs to make a TDE that peaked and decayed within the 80 day gap and levelled out by the time it became observable again. Based on this, a TDE is a plausible explanation for the late-time signal detected in this object.

SN 2020pkj. In the case of SN 20120pkj, the first r -band bin with a detection is the end of the normally declining tail, but after the first bin the detections rise slightly in the next r -band bin (figure 3.11). None of the binned photometry for the g - and i -bands give significant detections. Unfortunately, there is a gap in the observations immediately after the r -band detections preventing me from following its evolution closely at these phases. When it became observable again at > 200 days after the peak, no significant detections were found in any band. The duration of the transient is at least 75 days and could be up to 150 days. In the SMP light curve of this object, a similar rise at the same epochs is recovered, though it is found in the g -band instead of the r -band. This is likely due to the low significance of these detections at just 5.8σ in the r -band in the forced photometry and 5.2σ in the g -band in the SMP, showing that they are just on the detection limit of my binning technique.

With a redshift of $z = 0.02456$, this object is nearby enough for the end of the tail to be visible in the bins at 100 days. This could explain why the first bin is a detection, but the slight increase is still unexpected for a normal SN Ia decline tail. While the second bin is marginally consistent with my tail fit, a slight but steady increase can be seen in the unbinned flux values as well, suggesting that the brightening is real.

Assuming that the detected part of the late-time signal is the brightest part of a sibling transient, the tested transients need significantly more extinction of ~ 3 to 6

mag in the r -band (depending on the comparison transient) than was found for the main SN Ia peak. However, the ~ 80 day gap is long enough for a Type I sibling SN to peak higher during the gap and dim again before observations resumed. The other transient could have just started when it stopped being observable, and already have declined below the detection limit when the location became observable again. Since there is no constraint on the magnitude, this could work with any amount of extinction. Therefore, I cannot rule out conclusively that the late-time detections are due to another transient.

3.3.2.4 Late-time interaction candidates

Finally, I present the three objects (SN 2018grt, SN 2019ldf, and SN 2020tfc) whose late-time detections could not be explained by any of the explanations discussed above. Their light curves are shown in figure 3.13, the colours of the late-time detections in figure 3.14, and the SMP light curves for SNe 2019ldf and 2020tfc in figure 3.15. A scene modelling analysis could not be performed for SN 2018grt as it exploded early in the ZTF survey. The forced photometry pipeline uses images obtained during the ZTF commissioning phase as templates but the SMP does not, and therefore does not have enough images for the scene modelling to work. I present these three objects as late-time interaction candidates as there are no spectra of these objects at these late times to confirm the presence of features consistent with CSM interaction.

By using the estimated mean absolute r -band magnitude after removing all extinction effects (see Table 3.6) I can estimate the required $H\alpha$ flux assuming that it is the source of all the detected flux in the r -band, and that the emission line has the same width as was used during the simulations. The identified strength of the signal is compared to the $H\alpha$ signal detected in SN 2015cp for each of the events.

3. ZTF SN IA DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY




Figure 3.13: Three candidate objects, shown in magnitude and flux space. All three have significant detections ($\geq 5\sigma$) after a period of observations consistent with zero flux. From the alternative explanations, the best-fitting alternate transients are shown in dotted lines. For SN 2018grt this is the Type IIP SN 2017gmr, for SN 2019ldf and SN 2020tfc this is the TDE AT 2018hco.




Figure 3.14: Colour curves of the three candidate objects, together with the colours of the best-fitting alternate transients. The top and bottom panels show $g - r$ and $r - i$, respectively. The first bin for each object starts at zero days, but the bins can be shifted horizontally in an attempt to better fit the colour curve of the transient compared against (given that this is allowed by the rest of the light curve). Bins whose mean observation dates are closest to each other are used to calculate the colour, provided that these bins overlap in time. If there is a detection in only one band used to calculate the colour while the other is a non-detection, the result is a lower or upper limit.

SN 2018grt. The late-time signal of SN 2018grt is only detected in the r -band. The g -band stays around zero flux, and there are no i -band observations at these phases. The first detection in the r -band begins at 1350 days after the peak with a magnitude of 21.4 ± 0.2 mag (absolute magnitude of -16.4 mag). It varies little over the ~ 100 day period where it is detected, after which it returns to zero flux within 50 days. The SN is close to the host nucleus with an offset of $0.35''$ (0.32 kpc), but its host colours place it well outside the expected AGN region. Checking the difference images with SNAP shows that the host nucleus and SN location differ by ~ 1 pixel.

The late-time detections are ~ 2.5 mag below the main SN peak. If the late-time signal is due to another SN Ia or a TDE, it would require a significantly higher extinction value than was found for SN 2018grt itself. In addition, this object shows a sudden drop in the r -band, which is not normal behaviour for a TDE. Most of the

3. ZTF SN IA DR2: SEARCHING FOR LATE-TIME INTERACTION SIGNATURES IN TYPE IA SUPERNOVAE FROM THE ZWICKY TRANSIENT FACILITY




Figure 3.15: Binned SMP light curves of two of the three candidate objects. As no bins $\geq 5\sigma$ are recovered in SN 2018grt, it is not shown here. Both SN 2019ldf and SN 2020tfc do, however, still have robust detections in some bands. The best fitting alternate transients, shown in dotted lines, are the same as in figure 3.13. The i -band of SN 2020tfc is not shown as the background was not subtracted completely, resulting in a significant flux offset.

other transient types cannot reproduce the plateau followed by the sharp decline only detectable in the r -band, apart from a SN IIP, where the plateau of SN 2017gmr has nearly the same time span. However, to fit the observed magnitude with a SN IIP, it would require a host $E(B - V)$ that is three times higher than what was found for SN 2018grt, making this scenario less likely to be the case. Therefore, I conclude that late-time CSM interaction is a plausible scenario for the late-time signal in this object.

If I assume $0.21 \leq E(B - V)_{\text{host}} \leq 0.36$ mag and that the r -band signal is produced only by an $H\alpha$ emission line with a similar width to the one used in my simulations in section 3.2.4, I estimate the strength of the emission to be much stronger than SN 2015cp, at 60 to 100 times its emission strength. However, there are examples of SNe Ia-CSM with interaction strengths this strong. For example, SN 2020eum was within this range (Sharma et al., 2023).

SN 2019ldf. SN 2019ldf has late-time detections in the r -band beginning at 1 050 days after the peak and lasting for about 100 days, with an additional increase in brightness towards the end. There is a single 5σ detection in the i -band, but there are a number of lower significance detections coeval with the r -band detections. These detections are directly after a long period without observations due to the object being behind the Sun. Nothing is detected in the g -band during the time of the rise in the r -band. The binned SMP light curve recovers these late-time r -band detections and a single i -band detection, showing that these detections are not specific to the photometry method.

I compared the properties of the late-time detections to those of my comparison transient objects. Even if I assume that a SN exploded during the gap in observations in order to avoid needing a significantly larger $E(B - V)_{\text{host}}$ value, most SNe evolve too much over a period of 100 days to explain the detections. In addition, detections would also be expected in the g -band, which have not been found.

A TDE could fit the detections if it was intrinsically bright but heavily extinct, as this could explain the red colour and absence of signal in the g -band. However, SN 2019ldf is offset from the host nucleus by 0.65 " (0.78 kpc) and inspection of the difference images using SNAP shows that the late-time signal is more consistent with the SN location than the host nucleus location, disavouring the TDE explanation. Therefore, I conclude that the late-time signal could be due to late-time CSM interaction.

The late-time detections persist until the end of the observation window. To determine if there were still signs of interaction once it was visible again, I obtained g - and r -band photometry with the EFOSC2 imaging spectrograph (Buzzoni et al., 1984) on the ESO New Technology Telescope (NTT) in La Silla, Chile on 2023 May 19 as part of the extended Public ESO Spectroscopic Survey of Transient Objects+ (ePESSTO+; Smartt et al. 2015).

To examine whether the SN is still detected in the images from May 2023, image subtraction was performed by Dr. Steve Schulze. Due to the lack of reference images

in the g - and r -band filters from EFOSC2, images from the DESI Legacy Imaging Surveys Data Release 9 (Dey et al., 2019) were used. After aligning the images, They were subtracted from each other with the High Order Transform of Psf ANd Template Subtraction code version 5.11 (HOTPANTS; Becker, 2015). The brightness in the difference images was measured using aperture photometry. The photometry was calibrated against stars from DESI Legacy Imaging Surveys. The EFOSC2 and DESI Legacy Imaging Surveys filters are not identical, which might add an unknown systematic error to the reported photometry. The 5σ upper limits are $m_g = 24.7$ mag and $m_r = 24.3$ mag at 1397 days after the peak, with no detection in either band. This means that the signal has disappeared at this time and thus could have lasted, at most, for about 500 days.

Assuming that the r -band signal is entirely due to the $H\alpha$ emission, I estimate it to be ~ 60 times as strong as the late-time interaction found in SN 2015cp. However, this assumption is very simplistic, as it completely disregards the rise in the i -band. Therefore, it is only a rough estimate.

SN 2020tfc. This object has late-time gri -band detections, beginning at 550 days after the peak and lasting for at least 250 days. While the r - and i -bands are at more or less the same magnitude, the g -band detections are about 1.3 mag fainter. This immediately poses an issue for any alternate transient considered, as either there is a low amount of extinction and a weak g -band signal, or there is a high amount of extinction and the intrinsic signal is even brighter in the i -band. This, combined with the fact that the signal lasts for several hundreds of days with little variation, disfavours a SN as an alternate transient explanation. The intrinsic colour also heavily disfavours a TDE as these objects tend to have a similar intrinsic brightness in the g -, r -, and i -bands. In the binned SMP light curve the late-time detections were confirmed in the g -band. The SMP i -band data points have a large scatter, most likely due to uncertain background removal because of a low number of available images for the SMP template. Therefore, I do not consider them further in my analysis. In the r -band SMP light

curve, the data points are below my 5σ cut-off for detections, although one is very close to my limit at 4.6σ . However, the g -band detections seen in the forced photometry are confirmed by SMP suggesting that a real signal is present at late times in at least this band and at lower significance in the r -band.

Similar to SN 2020alm, the late-time signal was ongoing during my analysis, but unfortunately, there is no archival host galaxy spectrum available to compare to. As the host dominates the late-time signal (the SN is at a distance of $0.28''$ (0.21 kpc) from the host nucleus) and cannot be removed using difference imaging or SMP, as was done for the photometry, this prevented me from taking a spectrum of the late-time signal. With all alternate explanations ruled out or severely challenged by observations, the late-time CSM interaction remains as a plausible explanation for these late-time detections.

If I assume that the r -band signal is due to $H\alpha$ emission only, the interaction is estimated to be 110 to 150 times as strong as the interaction found in SN 2015cp. This is by far the strongest of the three, but again, this simple assumption is unrealistic as it completely ignores the measured g - and i -band signal that suggests a contribution from a continuum or other spectral lines to the late-time signal.

3.3.3 Late-time CSM interaction rates based on my candidate objects

I identified three objects (SNe 2018grt, 2019ldf, and 2020tfc) with potential late-time CSM interaction signatures. For SN 2018grt and SN 2019ldf, these detections were in the r -band (with no i -band data available for SN 2018grt and low significance i -band detections for SN 2019ldf). For SN 2020tfc, late-time detections were found in all three bands (gri). In my initial simulations to determine my recovery efficiency (section 3.2.4), I made the assumption that any CSM interaction would be dominated by $H\alpha$ emission that would be present in only the r -band up to $z \sim 0.07$ and in the i -band beyond that. This was based on the dominant interaction signatures seen in

SNe Ia-CSM and also the one event with late-onset interaction, SN 2015cp (Graham et al., 2019b).

In Graham et al. (2019b), the Ca II NIR triplet emission was also identified in its spectra, along with emission consistent with Mg I $\lambda 5175$. Harris et al. (2018) speculated that although the Ca II NIR region was very noisy, the Ca II emission may be at a similar strength to the $H\alpha$ emission. For SN 2019ldf at a redshift of $z = 0.05646$, the Ca II NIR triplet, if present, would be partially shifted out of the i -band, which could potentially explain its low-significance i -band detections. SN 2020tfc is at a lower redshift ($z = 0.031$), so the Ca II NIR triplet would fall completely in the i -band and could potentially be the source of the i -band detections. The strength of the potential Mg I $\lambda 5175$ line in SN 2015cp was weak and not well constrained (Graham et al., 2019b) but could potentially result in a weak signature in the g -band. However, without spectral confirmation, this is only speculation.

The other difference between my initial simulations for determining recovery efficiency and the observed detections in my three candidate late-time CSM objects is that the late-time CSM interaction signature is much stronger in the observed events. In my three candidates, the late-time signal is between 2 to 3 magnitudes below the SN peak magnitude, while in the simulations with the strongest interaction, the signal was about 4.4 magnitudes below the peak.

It is clear from my candidate events that there is significant diversity (detected bands, timescales and strengths) in their potential CSM signatures and without knowledge of their underlying spectra (e.g., emission lines that are present), it is difficult to develop spectral models and run simulations covering this full diversity. Therefore, although very simplistic and ignoring other potential emission lines occurring in the g - and i -bands as were seen in SN 2015cp (Graham et al., 2019b), I focussed on $H\alpha$ emission dominated models and the corresponding r -band detection efficiency to estimate a rate of late-time $H\alpha$ emission dominated CSM interaction based on my three candidate objects. I use the same underlying spectral model as detailed in section 3.2.4 (SN 2011fe combined with $H\alpha$ emission) in further SIMSURVEY simulations but con-

Table 3.8: Parameters used for rate estimation simulations for each object.

Object	Start epoch (d)	Duration (d)	Strength (SN 2015cp)	N_{sample}	late-time CSM fraction	late-time CSM rate ($\text{Gpc}^{-3} \text{yr}^{-1}$)
SN 2018grt (worst)	1 375 *	100	60	748	$0.0084^{+0.0183}_{-0.0041}$	203^{+438}_{-97}
SN 2018grt (best)	1 275 *	200	100	988	$0.0015^{+0.0032}_{-0.0007}$	36^{+76}_{-18}
SN 2019ldf (worst)	1 050 *	200	60	1 931	$0.0019^{+0.0036}_{-0.0009}$	45^{+87}_{-22}
SN 2019ldf (best)	875 *	500	60	2 505	$0.0023^{+0.0038}_{-0.0011}$	54^{+91}_{-26}
SN 2020tfc (worst)	550	250	100	3 439	$0.0004^{+0.0009}_{-0.0002}$	10^{+22}_{-5}
SN 2020tfc (best)	450	500	150	3 493	$0.0003^{+0.0008}_{-0.0002}$	8^{+20}_{-4}

For each of the three SNe I assumed a worst and best-case scenario for detecting them in my sample, giving upper and lower limits for a simulated observing campaign. The efficiency curves generated from these simulations were then used to determine the intrinsic fraction of SNe Ia with late-time CSM interaction in an MCMC process, with the assumption that only one object was recovered in a sample size that is the same as the amount of DR2 objects with observations after the start epoch (N_{sample}). The last two columns show the found late-time CSM interaction fraction of the total SN Ia rate and the late-time CSM interaction rate.

* When the start epoch of the late-time excess is > 750 days, 750 d is used as the start epoch because of limitations of the available ZTF survey plan. See section 3.3.3 for more details.

strain the strength and timescale of the CSM signature based on my three events. In Table 3.8, I show the range of start epoch, duration and strength (dependent on assumed host galaxy extinction) simulated for each event. For each of the three objects, I take the worst case scenario (shortest, weakest, latest start epoch) and the best case scenario (longest, strongest, earliest start epoch) allowed by the constraints from the data to estimate a worst and best case scenario recovery efficiency for each object.

Since the available ZTF survey plan for my simulations only spans 1 004 days, I cannot set the interaction signal to start at phases after this, as the interaction would never be observed. Therefore, I choose to start the interaction in my simulations no more than 750 days after the peak. This is done to ensure there is a decent chance to observe an event until after the interaction has started, while still being able to apply a good baseline correction. As the interaction is by far the dominant source of light at these epochs, changing the start date of the interaction in this way has little effect on the total brightness of the SN. However, it does affect the number of objects observed by allowing more objects to have observations at later times when the interaction

occurs. This is a redshift-independent effect but will impact the resulting recovery fraction function. To account for this, I only consider objects with observations after the start of the interaction and remove those where the simulation limits prevent any possible late-time detection. While changing the interaction start date has increased the number of objects that satisfy this condition, there are still some objects that do not due to, e.g., being in a sparsely observed field.

Not all objects in the DR2 have detections up to the time after the peak where the interaction was detected in the three final candidates. Similar to the simulation, in order to get a proper estimate of the rate I can only consider the objects that could have been found interacting (N_{sample}), and remove those without observations this late after the peak. This gives me an effective sample size that is smaller than the full DR2, and is listed in Table 3.8.

The expected number of observed objects with late-time interaction in a sample $N_{\text{recov}}(N_{\text{sample}}, \eta(z))$ is a function of the efficiency and the intrinsic rate, which in turn is a fraction of the total SN Ia rate. If I assume the SN Ia rate from [Frohmaier et al. \(2019\)](#) as in my initial simulations and that the fraction of these showing late-time interaction, f , is redshift independent, then the only z dependency is in the efficiency of my pipeline. In a similar approach to [Prajs et al. \(2017\)](#) and [Frohmaier et al. \(2018\)](#), I run a Markov chain Monte Carlo (MCMC) simulation with 10^7 realisations for each scenario to find the fraction of interacting SNe that best explains my findings.

In each realisation of the MCMC, I draw N_{sample} objects and assign a redshift to each object according to the distribution from [Frohmaier et al. \(2019\)](#). For each object I also draw two random numbers from a flat distribution between zero and one. The first random number is used to decide if the object was interacting, and the second to decide if any interaction would have been recovered. Interaction is true if the random number is below the chosen value for f for that realisation. Recovery is true if the second random number is below $\eta(z)$ at the z of the object (e.g. if the recovery efficiency is 0.7 at the SN redshift, a drawn number below 0.7 will result in a detection if the object was interacting). The number of recovered objects with interaction can be found by




Figure 3.16: Fraction of MCMC realisations per bin that resulted in one object being recovered in a sample with the same size as the effective DR2 sample size as a function of the late-time CSM interaction fraction. The best and worst-case scenario for each object is shown in blue and red, respectively. A skewed normal distribution fit with a 1σ uncertainty band is shown for each scenario, and the dashed and dotted lines give maximum and 68% confidence interval of these distributions, respectively. The distributions continue on the right side of each plot.

counting the objects for which both interaction and recovery are true, while the actual number of interacting objects can be found by counting the objects where interaction is true.

Figure 3.16 shows per bin the fraction of realisations that resulted in one object being recovered as a function of late-time CSM interaction fraction with the best and worst scenario parameters for the three discovered events. I approximate them by fitting a skewed normal distribution and use the fit to estimate the peak of the distribution and the 68% confidence interval on either side. These values are quoted as a fraction of the total SN Ia rate and as the late-time CSM interaction rate in the last two columns of Table 3.8.

3.4 Discussion

In this section, I first discuss the overall rate of CSM interaction signatures found in my analysis (section 3.4.1) and some of the properties of the three SNe Ia displaying these unexplained late-time detections (section 3.4.2).

3.4.1 Late-time interaction is rare in SNe Ia

In my study of the ~ 3500 SNe Ia in the ZTF DR2 with photometry at >100 days after the peak, I identified three objects (SN 2018grt, SN 2019ldf, and SN 2020tfc)

for which late-time CSM interaction is the best explanation for their detected late-time flux excesses. However, as each of these has significantly different parameters for the interaction, they have to be treated separately when trying to estimate the rate of late-time CSM interaction in SNe Ia (section 3.3.3). Except for the worst-case scenario of SN 2018grt, the identified rate for all cases agree with one another within the uncertainty, with the fraction of SNe Ia showing late-time $H\alpha$ -dominated CSM interaction between $0.03^{+0.08}_{-0.02}\%$ and $0.23^{+0.38}_{-0.11}\%$. This is similar to what was found by Sharma et al. (2023) in their Ia-CSM sample and agrees with the upper limit set by Graham et al. (2019b) with their discovery of the late-time interaction in SN 2015cp. Some objects were disregarded as I could not definitively determine that they were not due to a sibling transient. Therefore, there could be a handful of additional objects displaying strong late-time $H\alpha$ -dominated CSM interaction that were not included in my rate estimates, but this would double the rate at most. Therefore, I estimate that the intrinsic rate of late-time $H\alpha$ -dominated CSM interaction in SNe Ia is $< 0.5\%$.

The identified late-time signals that are consistent with CSM interaction occurred between 1.5 to over 3.5 years after the original SN peaked. If I assume that the SN ejecta have a velocity of the order of 10^4 km s^{-1} , the distance at which the CSM shell resides is of the order of 10^{17} cm . As CSM moves slowly, this means it must have been ejected from the system a long time ago. The further out the shell is from the progenitor system, the more mass is contained in even a thin shell. This is especially true when one considers that in order to get a detectable signal from the interaction, the CSM density cannot be too low. At these distances, light travel time effects also become significant, and a short interaction with a thin shell will be smeared out over a long time. Unfortunately, since I have only partial constraints on the interaction timescales and CSM mass estimation from an interaction signature is far from straightforward, I do not attempt to provide a CSM mass estimate.

An interesting aspect of the late-time detections in my three candidates is their inferred strength. While my model is likely an oversimplification by trying to explain the entire r -band signal using only a narrow $H\alpha$ line, it is clear that the signal is

much stronger compared to the one found in SN 2015cp, especially when there are other band(s) at a similar magnitude. Interaction this strong is not unprecedented, though. SN 2020aekp, for instance, starts to plateau around 50 days after peak at an absolute magnitude $M \sim -18.5$ in all three ZTF bands (Sharma et al., 2023) and holds this plateau for several hundreds of days. Most of the known SNe Ia-CSM presented in Sharma et al. (2023) have detections in multiple bands, and the spectra presented show more than just an $H\alpha$ line coming from the CSM interaction. CSM around a SN Ia also does not have to be spherically symmetric, especially in a scenario where material has been stripped and partially lost from the donor star in the progenitor system. Smith et al. (2015) show for SNe IIn that if the CSM mostly resides in a disk, only the ejecta travelling in the direction of the disk will be slowed down and interact while ejecta at higher inclinations continue to expand normally.

The colour curves for the late-time detections in my three objects (Fig. 3.14) show that $g-r \sim 1.7$ mag for SN 2020tfc, with limits of $g-r > 0.5$ and 1.5 mag for SN 2018grt and SN 19ldf, respectively. For the two objects with i -band observations, $r-i < 0.5$ mag. The SN Ia-CSM sample of Sharma et al. (2023) had late-time (> 300 days) $g-r$ colours of $0-0.5$ mag, which are similar to SN 2018grt but bluer than SN 2019ldf and SN 2020tfc. Kool et al. (2023) find $g-r \sim -1$ mag and $r-i \sim 0.5$ mag at late times for the He-interacting Ia-CSM SN 2020eyj. When taking into account the host galaxy extinction that they identify, it becomes even bluer and is much bluer than my three objects.

3.4.2 Properties of SNe Ia with late-time interaction

The three SNe Ia with late-time flux excesses that cannot be explained by other scenarios have SALT2 light curve fitter x_1 and c values that are generally typical of normal SNe Ia (Rigault et al., 2025a). However, the c value of SN 2018grt of 0.61 ± 0.03 is at the high end and would be excluded from cosmological analyses. I also identified significantly more host extinction for this object than for the other two based on their luminosities at the peak.

All three candidates are found at small projected distances from their host nuclei. While the small sample size might be used to explain part of this observation, the mean projected distance of all objects in the ZTF DR2 is 6.3 kpc (Rigault et al., 2025a). The chance for my three objects to be at most 0.8 kpc from the host is $< 1\%$, suggesting that objects with late-time interaction have a preference for small host separations, the detections are caused by nuclear variability, or they are caused by bad host subtraction. The location of the late-time signal is too close to the host nucleus to distinguish which location the late-time signal is more consistent with by using SNAP. The host galaxy colours are inconsistent with being AGN and the properties of the late-time detections cannot be easily explained by known nuclear transient classes such as TDEs. It is possible that by going to deeper limiting magnitudes with my light curve binning I have identified previously unstudied nuclear variability, but this cannot be confirmed. Imperfect template subtraction can occur in galaxy centres, resulting in a dipole artefact. However, these are easily recognisable with SNAP, and this possibility was ruled out for my three candidates. Therefore, I conclude that a likely interpretation is that there may be something intrinsic to SNe Ia exploding in these environments that produces late-time CSM interaction but further samples are required to confirm.

The morphologies of the host galaxies of the candidate events are broadly elliptical or spheroidal in nature with no signs of spiral arms. The galaxies have stellar masses of 4×10^9 , 6×10^{10} , and $2 \times 10^{10} M_{\odot}$ for SN 2018grt, SN 2019ldf, and SN 2020tfc, respectively. These host masses are all within the bulk of the masses of the ZTF DR2 SN Ia sample (Rigault et al., 2025a). The hosts of the known H-rich Ia-CSM sample have stellar masses of $\sim 3 \times 10^8$ to $3 \times 10^{10} M_{\odot}$, consistent with my objects (Sharma et al., 2023). SN 2020eyj, the He-interacting and radio detected SN Ia-CSM (Kool et al., 2023) also had a very small offset from its host of $0.57 \pm 0.02''$ (0.36 ± 0.01 kpc). Its host only has a WISE $W3$ upper limit so it cannot be excluded from having an AGN present. It is a compact star-forming galaxy with a mass of $\sim 6 \times 10^7 M_{\odot}$, which is lower than the typical Ia-CSM range (Sharma et al., 2023).

The hosts of the three candidate events all have WISE $W1 - W2 \approx 0$ mag. SN 2020tfc has $W2 - W3$ of ~ 0.8 mag, while the other two have only limits with $W2 - W3 > 1.2$ mag. Based on Fig. 2 of [Irani et al. \(2022\)](#), this places them in the overlap region between elliptical and spiral host properties from the galaxy sample of [Lintott et al. \(2011\)](#). The vast majority of Ia-CSM sample of [Sharma et al. \(2023\)](#) has $W2 - W3$ colours of > 1 mag, again broadly consistent with my sample. However, the morphologies of my candidate events do not show evidence for spiral arms and their overlap between ellipticals and spirals in the $W2 - W3$ parameter space suggests they may come from different, older stellar populations than the known Ia-CSM events ([Kool et al., 2023](#); [Sharma et al., 2023](#)). However, to prove that these are different populations requires a larger sample.

3.4.3 Limitations of the analysis

Since there is no good model of a Type Ia SN interacting with CSM at late times that could be used for my SIMSURVEY simulations, I had to make my own based on several assumptions. My main assumptions were that the SN Ia looks normal until the moment the interaction starts, and I assumed that this interaction was with hydrogen-rich material showing itself primarily as an $H\alpha$ emission line. This is motivated by the fact that $H\alpha$ is found in most SNe Ia-CSM and that the late-time interaction in SN 2015cp was confirmed through the observation of an $H\alpha$ emission line. However, [Graham et al. \(2019b\)](#) also found other emission lines that are associated with the SN, such as the Ca II triplet near 8500 Å and a tentative detection of Mg I $\lambda 5175$, suggesting that my model is too simplistic. Similarly, for SN 2020tfc, I identified late-time detections in the *gri*-bands, which cannot be explained by $H\alpha$ emission alone.

This method of binning the late-time light curves allows me to push the detection limit down to a limiting magnitude $m = 21.44$ mag. Since the references used in ZTF have a mean limiting magnitude of $\overline{m}_{\text{lim}} \sim 21.8$ mag, this becomes the leading uncertainty preventing my binning technique from going deeper. I have shown through my simulations that deeper references will allow for the binning technique to go deeper

and recover fainter interaction signatures in a larger redshift volume. I have also used SMP to generate light curves that did not rely on forced photometry to test for issues coming from the photometry measurement. In most cases the late-time detections were recovered with both methods, showing their robustness. However, the SMP has been found to have issues with identifying a baseline flux when the number of images used in SMP template creation is small. This highlights the requirement for a long off time to model the underlying galaxy light sufficiently.

The pipeline focuses more on the historical light curve of the SNe than their current state. Since I require two adjacent bins to be detections in order to avoid false positives and the smallest bins I use are 25 d, the interaction needs to be active for over a month before it is picked up. Only after this can it be followed up with spectroscopy if it is identified fast enough and the interaction signature is still occurring. Unfortunately, in the cases of my three candidates, this could not be done due to a combination of recovering the interaction late, the objects being blocked by the Sun for a section of the year, and not having a suitable reference spectrum of the host galaxy to subtract and isolate the interaction spectrum. Adapting the pipeline to run in real-time would help to decrease the time until detection, but if the host galaxy contribution is strong, an archival reference spectrum is required.

3.4.4 By-product pipeline detections

The conservative, catch-all approach outlined in my work poses the challenge of identifying exactly why each object is flagged by the pipeline. My approach allows for different kinds of objects that become interesting at late times to be caught. I recovered nearly all known Ia-CSM SNe through their non-normal SN Ia tails, as well as a number of known and unknown siblings, along with evidence of a change in the declining tail slope of close-by, normal Type Ia SNe around 200 days after the peak. In my sample, 4% of the objects needed to be visually inspected, upon which the cause of the (false) positive became clear quickly for most objects. Some of the checks, such as using the WISE colours to identify an AGN host close to the SN location, have been

automated, reducing the workload in future attempts to use this method for searching for late-time signals, especially when using large datasets such as the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST; [Ivezić et al., 2019](#)).

The pipeline, as presented in this paper, is specifically tailored for finding SNe Ia that deviate from normal behaviour at > 100 days after the peak. However, the method of binning late-time observations can be used for any type of transient, but the checks to ensure that the resulting detections are not due to expected behaviour are SN Ia specific. When other types of transients (e.g. SN Ib, Ic, IIP) are used as input, the result will likely vary. For instance, a long-lasting plateau in an SN IIP will be unexpected by the pipeline as it does not follow the declining tail of a normal SN Ia. However, modifying the pipeline to work for different kinds of transients mostly requires revision of these checks.

3.5 Conclusions

I have presented a search of the ZTF SN Ia DR2 light curve sample for signatures of late-time CSM interaction. This is the first systematic search for signatures of late-time interaction in SNe Ia in a large optical survey. I made a custom pipeline to calibrate and bin the light curve data at more than 100 days after the peak, using bins with sizes between 25 and 100 days. My analysis was based on searching for $H\alpha$ emission (as was seen in SN 2015cp and Ia-CSM) in the r -band at lower redshift and i -band above a redshift of $z = 0.07$. I performed simulations with SIMSURVEY to determine the efficiency of my search, as well as the intrinsic rate of potential $H\alpha$ -dominated CSM interaction in the sample. My main conclusions are:

1. My pipeline returned 134 SNe that were potentially interesting based on their late-time light curves. Visual inspection of these objects was performed using my visualisation program, SNAP, to inspect the difference images, removing 101 objects as false positives.

2. Of the remaining 33 objects, I identified 13 out of the 14 known Ia-CSM objects in the DR2, five siblings close to the position of the original SN Ia, four very nearby events whose late-time light curves were not captured by my simple radioactive tail model and one Iax with a late-time excess in all three bands, consistent with the presence of a bound remnant.
3. Out of my final shortlist of ten candidate events, I identified three SNe Ia (SN 2018grt, SN 2019ldf, and SN 2020tfc) that displayed late-time detections beginning 550 - 1350 d after peak and lasting at least 100 - 250 d, which could not be explained by data issues, AGN activity, or other transient events exploding at a similar location.
4. For SN 2018grt, these late-time detections were only in the r -band (no coeval i -band data was available). For SN 2019ldf, detections were made in the ri -bands and for SN 2020tfc in all three bands, suggesting potential contributions with Ca II NIR emission or other Mg I as was identified in SN 2015cp (Graham et al., 2019b).
5. The r -band magnitudes of the late-time interaction are -16.5 , -16.4 , and -16.8 mag for SN 2018grt, SN 2019ldf, and SN 2020tfc, respectively. At their respective redshifts, this corresponds to $H\alpha$ interaction strengths of 60 - 150 times that of SN 2015cp (depending on the extinction correction used). The strong nature of this signal could suggest I might only have found the high end of the late-time interaction strength distribution.
6. Using SIMSURVEY simulations of the ZTF survey, I estimated the intrinsic rate of strong $H\alpha$ -dominated late-time (> 100 days after the SN peak) interaction to be occurring in $<0.5\%$ of SNe Ia. This translates to absolute rates of 8_{-4}^{+20} to 54_{-26}^{+91} Gpc $^{-3}$ yr $^{-1}$, assuming a constant SN Ia rate of 2.4×10^{-5} Mpc $^{-3}$ yr $^{-1}$ for $z \leq 0.1$.

The rarity of late-time interaction (occurring in $< 0.5\%$ of SNe Ia) highlights the importance of a large dataset of objects that have been observed for multiple years. The late-time detections occurred at different epochs for each object (from 550 – 1350 days after the peak), showing that the phase at which SNe Ia will start to show signs of CSM interaction is highly variable. Therefore, the only viable strategy is to keep observing SNe even after they have faded beyond the detection limit and binning the late-time light curves. The interaction strength that the binning technique is sensitive to is heavily dependent on both the science and reference image depth. Future improvements to the analysis would include the use of deeper reference images to detect fainter signatures of late-time interaction, as well as running the pipeline in real-time so that additional photometry and spectroscopy can be obtained to further characterise the late-time excesses. The deeper magnitude limits of LSST would be ideal for this study but cannot be immediately performed when the survey starts because of the requirements of deep reference images, as well as the need to wait up to more than three years after the SN Ia peak for the interaction to occur.

ZTF-observed late-time signals of pre-ZTF transients

In chapter 3 I presented my pipeline to search for faint signals in light curves by binning observations together to push beyond the single observation magnitude limit. This pipeline was then used to search for signs of late-time CSM interaction in the ZTF SN Ia DR2 data set, finding three candidate objects in a sample of 3 628 objects. For a late-time signal to be detected in this sample requires both the SN and the late-time signal to occur after the start of ZTF, which sets an age limit of at most 5.5 years on any theoretically identifiable CSM interaction.

However, the binning procedure does not rely on having the original SN as part of the light curve that is to be binned. As long as the coordinates of the SN are known, a forced photometry light curve spanning the entire duration of ZTF can be constructed and put through the pipeline (although a slight modification of the baseline correction will be needed, see section 4.2.1). This means that by using SNe that exploded before the start of ZTF, much later epochs can be investigated for signs of late-time signals. On top of this, as most SNe are expected to have faded before the start of ZTF, the procedure becomes identical for any type of transient, allowing me to widen my search beyond SNe Ia.

This chapter is based on Terwel et al. (submitted), which has been submitted to A&A. The sample selection, data reduction, analysis and interpretation were performed by me. Where relevant, I have noted the specific parts that were carried out by collaborators. In Section 4.1, I build my sample of objects whose first detection was in the decade before the start of ZTF. In Section 4.2, I introduce the changes I made to the pipeline to adapt it to this sample, as well as the post-binning analysis. In Section 4.3, I describe the objects flagged by the binning program and split them into different groups with different origins. These are discussed in Section 4.4, and I conclude in Section 4.5.

4.1 Data

My first aim is to compile a list of all transients discovered in the decade before ZTF started. To build my sample, I started with all transients in the Open Supernova Catalog (OSC¹, Guillochon et al. 2017) that were discovered between 1 January 2008 up to 1 January 2018, giving me 22 790 objects. I end my sample a few months before the start of the full ZTF survey in March 2018 to remove most pre-ZTF transients that were still visible by the start of ZTF. For each object I require a name, sky position, redshift, and classification of its type. I require the redshift to make an informed estimate of the absolute luminosity of any potentially detected late-time signal. For the vast majority of transients a classification means a spectrum was obtained that enabled typing, but the size of the sample precludes checking of each individual classification. However, further investigation of identified interesting events is performed where required. I also include objects classified as a ‘SN candidate’ as most of these are likely SNe but were never spectroscopically classified, although this group of objects is likely to be contaminated with some non-SN transients. These requirements cut my sample down to 8 865 objects. By querying WISEREP² (Yaron & Gal-Yam, 2012) for objects reported in the same date range, I attempted to recover objects that were incomplete or absent from the OSC. Out of the 12 955 OSC objects that did not have all the required information and no WISEREP match, 1 137 had no reported sky position while the other 11 818 had a sky position but no reported redshift and/or classification.

From the combined sample, I removed objects at $\text{Dec} \leq -32^\circ$ as ZTF is unable to observe below $\text{Dec} \approx -30^\circ$. After these updates (including WISEREP and with the Dec cut) my sample increased to 8 914 objects. The Dec cut is liberal to ensure that all object locations that could have been observed with ZTF remain in my sample. Objects outside the ZTF footprint but still in my sample are cut automatically when obtaining the ZTF light curves. I do not make a cut on proximity to the host nucleus as I want to be as inclusive as possible. Instead, I accept the possible contamination

¹<https://github.com/astrocatalogs/supernovae>

²<https://www.wiserep.org>




Figure 4.1: The final sample split into nine classes. As there is a big difference in class sizes, the smaller classes have been put together in the left chart as ‘other’ and are split up in the right chart.

by nuclear activity which, as will be shown later, is in most cases easily identifiable.

While cross-matching between OSC and WISEREP to remove duplicate entries, I noted several objects with different names and discovery dates but located very close together in the sky. While uncommon, it is not impossible for multiple SNe to occur at the same sky position in short succession. In chapter 3 I found several such sibling transients at very small separations where both SNe were detected by ZTF. Since I just look at the late-time light curve at the sky position of each pre-ZTF transient, sibling transients like these are investigated together as they have the same sky position.

My late-time light curves at the position of each transient in the list were constructed, using the method detailed in section 4.2, between 9 December 2023 and 24 January 2024 using all available ZTF data at the time. This results in an effective range of over 5.5 years of ZTF observations for each object in my sample. For 207 sky positions that lay close to the edge of the ZTF survey, no ZTF observations were obtained and these were removed from my sample, leaving me with 8 707 transients. The ZTF light curves at these sky positions are generated using FPBOT¹ (Reusch, 2023).

¹<https://github.com/simeonreusch/fpbot>




Figure 4.2: Sample size as a function of redshift for each class of objects. All objects with $z > 1$ are put together in the bin starting at $z = 1$. The histograms are split into two plots for better readability. In all classes of objects my sample is biased towards lower z .

I group the transients in my sample into one of nine classes based on the most precise common denominator of the reported classifications. For example, an object that was reported as both a SN Ia and SN Ib is grouped as a ‘SN I’, and an object that was reported as some type of SN I and SN II is grouped as a ‘SN’. SLSNe are handled the same and grouped with normal SNe into SN I or SN II, though I can recover their SLSN classification if the object proves interesting later. Objects whose classification included the interacting classes SNe Ibn and SNe IIn are kept separately as I expect these may have a higher chance of having a late-time signature that can be picked up by ZTF. Figure 4.1 shows the nine classes in my sample (‘Ia’, ‘II’, ‘Ib/c’, ‘IIn’, ‘Ibn’, ‘I’, ‘SN’, ‘SN candidate’, and ‘not SN’) and their relative sizes after a cut on the number of points used in the baseline correction (see section 4.2.1).

For completeness I include a ‘non SN’ class that consists of other types of transients. This class consists of 60 variable stars which include cataclysmic variables (CVs), luminous blue variables (LBVs), and novae, 36 nuclear transients (including TDEs and AGN), 73 other transients (including gap transients, impostor-SNe, kilonovae), and

140 Long Gamma-Ray Bursts (LGRBs). Figure 4.2 shows the amount of objects in each class as a function of redshift. My sample is biased towards lower redshifts as it is magnitude-limited. I also have more objects that have been discovered in years where large surveys such as PTF, iPTF and ATLAS were active.

4.2 Analysis

I use an adapted version of the pipeline introduced in chapter 3 to test for late-time flux excesses in my sample of pre-ZTF transients. In brief, the pipeline first applies a baseline correction to ensure that the light curve has zero flux when no signal is expected (e.g. Yao et al., 2019; Miller et al., 2020). It then bins the post-SN observations together in bins of 25, 50, 75, or 100 days to recover signals that are below the noise level of individual observations. To test if bin placement has an effect on the result, the binning is performed multiple times with shifted bin phase locations. Only signals that are sufficiently insensitive to bin placement are considered real. Binning observations increases the depth at which signals can be recovered at the cost of time sensitivity. As the smallest bins are 25 days in the observer frame, this means that I cannot detect details at similar or smaller time scales, such as rise times or short-term variations. In section 4.2.1, I describe the modifications of my detection pipeline from chapter 3. In section 4.2.2, I discuss the identification of false positives and in section 4.3.4.1, I describe a method of obtaining a rough estimate of potential CSM masses to explain any observed late-time signals.

4.2.1 Detection pipeline modifications

My main modifications to the pipeline presented in chapter 3 are in the method of the baseline correction and the removal of the tail fitting procedure that was put in place to remove false positives from detecting the end of a normal SN Ia tail decline. In chapter 3 the baseline correction was done by using the pre-SN observations, but this cannot be done in my current sample since the date of explosion is not contained

4. ZTF-OBSERVED LATE-TIME SIGNALS OF PRE-ZTF TRANSIENTS

Table 4.1: Details of the three baseline regions used in my pipeline. The first column gives the start and end MJD of each baseline region, and the second column gives its length in days. The last three columns give the number of transients that had at least 30 points in the baseline to provide a robust estimate of the baseline, in the g , r , and i -band, respectively.

Baseline region (MJD)	Length (d)	No. passing cuts		
		g	r	i
58300 – 58664	365	6817	7198	744
59031 – 59395	365	7251	7258	3447
59915 – 60335	422*	4792	5350	2457

* The light curves were generated between 9 December 2023 and 24 January 2024, meaning that the final baseline for each object is between 374 and 420 days long. The final two days are a buffer to ensure all data is used.

within the ZTF data time frame. Instead, I have to consider that a period with flux excess could occur at any time during the ZTF observations.

Therefore, I run the pipeline three times, using different time frames for the baseline corrections. These time frames are chosen at the start, middle, and end of my ZTF light curves (see Table 4.1). I choose the baseline time frames to each be a year long to ensure that the targets are likely to have been observed, even if there are gaps in the observations due to the target location not being observable the entire year. I also choose not to include the first few months of ZTF data in the first baseline region as the final calibrations were not completed until July 2018 (Bellm et al., 2019b). By using different baseline regions I ensure that I bin the entire light curve multiple times (since the region used for the baseline corrections cannot be included in the light curve binning). If one of the baseline regions overlaps with the late-time signal the result can be a false baseline correction. By using multiple baseline regions, these cases can more easily be identified.

To ensure that there were enough points in each baseline region for a robust correction, I ignore all detections where the band in which the detection was found had < 30 points in the baseline. With the best ZTF cadence of one data point every two days this translates to requiring two months of good observations. On the other hand, if the observations are spread over the entire baseline region, the requirement translates into averaging one observation every 12-14 days. Both give a good estimate of the baseline,

either by monitoring it closely over a shorter duration or more globally over a longer time period. Table 4.1 shows the number of objects that meet this condition within each baseline region in each of the observational bands. In total, 989 objects never meet this condition in any band and are effectively removed from my sample, reducing the sample to 7718 objects.

4.2.2 Removing false positives

The pipeline outputs a list of 360 objects that have 5σ or greater binned detections in a band in at least four out of 16 attempts (four bin sizes, each shifted four times to avoid spurious detections caused by specific bin placement, see chapter 3 for more details). These are inspected visually to determine if the detections are due to observational issues, software issues, or if it is likely astrophysical in nature. A large fraction of the flagged light curves were deemed false positives after visual inspection. The main causes of the binning program wrongly finding bins with detections are due to issues in the difference imaging processing or the baseline correction.

To investigate these cases, I use the SuperNova Animation Program (SNAP)¹ to inspect the difference images directly. Details on SNAP are presented in section 3.2.3. Many false positives are identified as due to being close to another source by inspection of the difference images. Usually, this other source is the host galaxy (nucleus) but in some cases it can also be a foreground star. This can lead to an issue with an improper subtraction of the bright source, which results in a residual, generally a dipole, at the position that is picked up by my pipeline. To identify these cases, I use SNAP to visually inspect the images and remove spurious detections. This removes 155 events from my sample.

The other main group of false positives (88 cases) are flagged due to issues in the baseline determination. These can present as extremely large corrections, in some cases several orders of magnitude larger than the signals I expect to find. A baseline correction of $\mathcal{O}(10^5)$ can make a 17.5 mag signal appear or disappear. While in some

¹https://github.com/JTerwel/SuperNova_Animation_Program

cases big corrections are needed (see section 4.3.1), in most cases such a correction led to noisy light curves, preventing me from probing beyond the individual ZTF image mag limit. In other cases the baseline is not constant but seems to vary over time or suddenly jump, making it very difficult or impossible to apply a proper baseline correction.

One potential cause of these issues could be due to ZTF continuing to improve the reference images by rebuilding them and stacking more observations, including observations from during the survey (Masci et al., 2019). To be able to compare observations from before and after this is done, one should remake all the difference images using the updated references. However, this is not feasible in a survey as large as ZTF, and since the offsets are usually below the noise threshold of the individual images the issue is of little importance to most of the survey science outputs. Only when attempting to go beyond the single image noise limit by using, for instance, my binning method, this issue becomes noticeable enough and can lead to baseline offsets and varying or jumping baselines.

In 19 cases there were enough points in the baseline for the object to stay in my sample but the light curve was sampled sparsely with large gaps without observations surrounding sparse detections, making it impossible to determine the validity of these detections as there are no proper non-detections close in time to compare against.

4.3 Results

I have identified 98 transients with potential late-time excesses in their light curves that require further investigation. In section 4.3.1, I describe the 63 objects whose detections in ZTF can directly be linked to the pre-ZTF transient. These transients were still bright enough at the start of the ZTF survey to be detectable. In section 4.3.2, I describe 12 objects whose found ZTF signal is due to a sibling transient occurring at nearly the exact same sky position. A small fraction of my sample consists of non-SN transients. 14 of these are returned by the pipeline, and these are discussed in

section 4.3.3. Finally, in section 4.3.4, I describe the nine objects that required an individual, deeper investigation of the signal found in ZTF.

4.3.1 Pre-ZTF transients still active in ZTF

I chose to limit my sample to those objects that were first detected before 2018 to reduce the number of transients still visible at the start of ZTF. While this three-month gap is enough for most transients to fade away, some super-luminous SNe, Type II-Plateau SNe, interacting classes like SNe Ibn and IIn, and even very nearby SNe Ia that exploded before ZTF started, may still be active at the start of ZTF.

As these objects were active while the initial set of ZTF reference images were being produced, none of these objects have been found by ZTF even when the SN was still bright enough to be detected by ZTF once observing began. As the transient is in the reference images, it will cause an over-subtraction and leave an imprint, or ghost, at its location in the difference images. These are easily recognisable through visual inspection of the difference images using SNAP. Another clear sign is a significant baseline correction that is consistent between the different baseline regions after the transient has faded away. The baseline correction corrects for the flux offset created by the ghost, revealing the tail as observed by ZTF in the light curve.

I find 63 transients whose ZTF detections are consistent with ongoing transient flux. These are listed in Table 4.2 and some example light curves are shown in Fig. 4.3. The light curves show pre-ZTF data taken from the literature for SN 2017gmr (Andrews et al., 2019), SN 2017erp (Brown et al., 2019), SN 2015da (Tartaglia et al., 2020), SN 2017egm (Zhu et al., 2023), and SN 2017dio (Kuncarayakti et al., 2018). SN 2015da is an extremely slowly declining SLSN IIn event (Tartaglia et al., 2020; Smith et al., 2024) and its light curve extends in the binned ZTF data to approximately eight years after discovery. In SN 2015da, the pre-ZTF and binned ZTF light curves do not overlap exactly in magnitude at epochs when data from the different surveys are available (MJD 58300 – 58600), showing that my pipeline underestimates the brightness of the transient. This is because some SN flux is still present even in my latest baseline

4. ZTF-OBSERVED LATE-TIME SIGNALS OF PRE-ZTF TRANSIENTS

region during mainly 2023, as can be seen in the observations presented in [Smith et al. \(2024\)](#). Therefore, the baseline correction is too small causing the binned detections to be slightly too low.




Figure 4.3: Examples of pre-ZTF SNe whose light curves have been recovered in the binned ZTF light curves. For each object I show the pre-ZTF g (dark green), r (dark yellow), and i (dark red) data points, as well as the binned ZTF detections (green, orange, red, respectively). The binned ZTF observations are in 25 day bins, with the width showing the beginning and end of each bin, and the shaded region shows its 1σ magnitude uncertainty. The pre-ZTF data was taken from [Andrews et al. \(2019\)](#) (SN 2017gmr), [Brown et al. \(2019\)](#) (SN 2017erp), [Tartaglia et al. \(2020\)](#) (SN 2015da), [Zhu et al. \(2023\)](#) (SN 2017egm), and [Kuncarayakti et al. \(2018\)](#) (SN 2017dio). None of the light curves are corrected for host extinction.

Table 4.2: List of the 63 pre-ZTF objects whose tail is still visible in ZTF using my binning method after a baseline correction. For transients that were not spectroscopically classified the quoted redshift is that of the host galaxy.

Name	Type	z	Discovery date	visible until (mjd)	Object Papers
AT 2017fwg	Candidate	0.113627	01-08-2017	58325	
AT 2017gpv	Candidate	0.025563	04-09-2017	59480	
AT 2017gpy	Candidate	0.025187	04-09-2017	58250	
AT 2017hat	Candidate	0.020454	01-10-2017	58275	
AT 2017igk	Candidate	0.040847	16-11-2017	58250	
AT 2017ihh	Candidate	0.027521	03-10-2017	58300	
AT 2017iho	Candidate	0.053421	17-11-2017	58250	
AT 2017ims	Candidate	0.030094	15-11-2017	58300	
AT 2017iru	Candidate	0.018156	29-11-2017	58375	
ATLAS17nbe	Candidate	0.013476	05-11-2017	58400	
ATLAS17oai	Candidate	0.099205	23-12-2017	58300	
ATLAS18eas	Candidate	0.033103	30-12-2017	58325	
ATLAS18mmr	Candidate	0.077507	15-12-2017	58325	
DES16X2bkr	SN II	0.159	21-09-2016	58400	

Name	Type	z	Discovery date	visible until (mjd)	Object Papers
SN 2015da	SLSN IIn	0.007222	09-01-2015	59475	Tartaglia et al. (2020) ; Smith et al. (2024)
SN 2016bkv	SN II	0.002	21-03-2016	58475	Hosseinzadeh et al. (2018) ; Nakaoka et al. (2018) ; Deckers et al. (2021)
SN 2016cyi	SN IIn	0.044	25-06-2016	58475	
SN 2016ieq	SN IIn	0.066	14-11-2016	58400	
SN 2017aym	SN IIP	0.005928	13-01-2017	58700	
SN 2017dio	SN Ic	0.037	26-04-2017	58300	Kuncarayakti et al. (2018) ; Shi et al. (2024)
SN 2017dpu	SN II	0.018	30-04-2017	58250	
SN 2017eby	SN Ia-CSM	0.081	01-04-2017	58350	
SN 2017egm	SLSN I	0.030721	23-05-2017	58300	Chen et al. (2017) ; Nicholl et al. (2017) ; Wheeler et al. (2017) ; Bose et al. (2018) ; Izzo et al. (2018) ; Maund et al. (2019) ; Hatsukade et al. (2020) ; Saito et al. (2020) ; Tsvetkov et al. (2022) ; Lin et al. (2023) ; Zhu et al. (2023) ; Li et al. (2024)
SN 2017emq	SN Ia	0.005247	03-06-2017	58250	
SN 2017erp	SN Ia	0.006174	13-06-2017	58375	Brown et al. (2019)
SN 2017err	SLSN IIn	0.107	12-06-2017	58275	
SN 2017faa	SN II	0.01845	27-06-2017	58300	
SN 2017fgc	SN Ia	0.007722	11-07-2017	58400	Burgaz et al. (2021) ; Zeng et al. (2021a)
SN 2017fvr	SN IIP	0.012539	01-08-2017	58450	
SN 2017gas	SN IIn	0.011	10-08-2017	58850	
SN 2017ghw	SN IIn	0.076	25-08-2017	58450	
SN 2017glx	SN Ia-91T	0.011294	03-09-2017	58400	
SN 2017gmr	SN IIP	0.005037	04-09-2017	58475	Andrews et al. (2019) ; Nagao et al. (2019) ; Utrobin et al. (2021)
SN 2017gvb	SN IIn	0.030344	18-09-2017	59000	
SN 2017gww	SN II	0.01748	26-09-2017	58375	
SN 2017gxq	SN Ia	0.008406	17-09-2017	58450	
SN 2017hbg	SN II	0.016	25-09-2017	58375	
SN 2017hca	SN II	0.013403	28-09-2017	58450	
SN 2017hcc	SN IIn	0.0173	02-10-2017	58500	Prieto et al. (2017) ; Kumar et al. (2019) ; Smith & Andrews (2020) ; Chandra et al. (2022) ; Moran et al. (2023) ; Mauerhan et al. (2024)
SN 2017hfv	SN Ia	0.028199	10-10-2017	58200	
SN 2017hix	SN Ic	0.012	13-10-2017	58375	
SN 2017hlt	SN Ia	0.027	10-10-2017	58275	
SN 2017hmi	SN Ia	0.0398	18-10-2017	58250	
SN 2017hpa	SN Ia	0.015654	25-10-2017	58400	Dutta et al. (2021) ; Zeng et al. (2021b)
SN 2017hqj	SN IIP	0.009	27-10-2017	58300	
SN 2017hro	SN II	0.015	28-10-2017	58800	
SN 2017igf	SN Ia	0.005624	12-11-2017	58350	
SN 2017ijr	SN Ia	0.04	20-11-2017	58300	
SN 2017ijx	SN Ia	0.027729	18-11-2017	58300	
SN 2017ivh	SN II	0.008	05-12-2017	58300	
SN 2017ivu	SN IIP	0.006528	11-12-2017	58400	
SN 2017ivv	SN II	0.022	12-12-2017	58450	Gutiérrez et al. (2020)
SN 2017ixg	SN Ia	0.0277	14-12-2017	58400	
SN 2017ixv	SN Ic-BL	0.007302	17-12-2017	58350	
SN 2017ixx	SN II	0.041	17-12-2017	58300	
SN 2017ixz	SN IIb	0.024	14-12-2017	58250	
SN 2017iyd	SN IIb	0.0285	13-12-2017	58275	
SN 2017jav	SN Ia	0.01517	19-12-2017	58275	
SN 2017jbj	SN II	0.013492	20-12-2017	58400	
SN 2017jeh	SN Ia	0.020961	26-12-2017	58275	
SN 2018L	SN Ia	0.02582	25-12-2017	58325	
SN 2018bq	SN Ia	0.025628	30-12-2017	58275	
SN 2018fd	SLSN I	0.263	11-10-2017	58450	



Figure 4.4: Light curve of AT 2017gpv with the axes representing absolute magnitude (left), apparent magnitude (right), the rest frame days since first detection (top) and mjd (bottom). Single epoch detections by ZTF, Gaia, Pan-STARRS and ATLAS are shown, as well as the binned ZTF observations after the transient faded below the single epoch noise limit. The dashed lines are the *gri*-band light curves of iPTF2014hls, which have been corrected for time dilation but not extinction (data taken from [Arcavi et al. 2017](#); [Sollerman et al. 2019](#)).

4.3.1.1 AT 2017gpv - a 14hls-like event

AT 2017gpv, shown in figure 4.4, is an unusual event similar in nature to iPTF14hls ([Arcavi et al., 2017](#); [Sollerman et al., 2019](#)) and SN 2020faa ([Yang et al., 2021](#); [Salmaso et al., 2023](#)) that was detected by my pipeline in the ZTF data. It was originally identified by Gaia, which detected it repeatedly during the first 500 days after its discovery. It has also been detected by ATLAS, which has a rich *o*-band light curve between 100 and 300 days after discovery that shows a plateau followed by a shallow decline that is interrupted by a rebrightening event. After the baseline correction the first detections in ZTF are visible even in the non-binned data, beginning around 200 days after first detection and lasting for around 400 days. With the binned observations the object can be recovered for another 1 000 days before it fades below the noise limit. Pan-STARRS also detected this object in the *w* band, and given its deep detection limits, it detected the transient up to over 1 400 days after discovery.

Due to the long duration of detections of AT 2017gpv in ZTF, only the last baseline region can be trusted as it has the least contribution from the late-time signal, though given the slow decline of the transient it is likely that there could still be some excess

present at the time of the last baseline region. As expected, the baseline corrections are significant, and SNAP clearly shows the transient in the reference images, as well as a significant ghost in all difference images. While it was picked up by the ZTF alert system and given an internal name (ZTF18acueiall), it was not recognized as a real transient.

Unfortunately, AT 2017gpv was never spectroscopically classified. The transient is at a distance of $6.88''$ from the host nucleus, removing host variability as a possible explanation, as well as limiting the amount of host extinction expected. The long time it was detectable and the bumpy nature of its light curves look similar to iPTF14hls (Arcavi et al., 2017; Sollerman et al., 2019), a very peculiar SN II that is also shown in Fig. 4.4. Some similar events to iPTF14hls have been identified, but they are rare (Yang et al., 2021; Soraisam et al., 2022). The light curve of iPTF14hls spans over 600 days and is very bumpy. In both cases the explosion epoch is badly constrained, but the peak found in iPTF14hls matches up quite well with the bump around MJD 58250 in AT 2017gpv. Overall, AT 2017gpv looks like a fainter and somewhat faster decaying version of iPTF14hls. The late-time data of iPTF14hls (Sollerman et al., 2019) extends to $\sim 1\,200$ days after discovery and shows a sharp decline after $\sim 1\,000$ days, which is not seen in AT 2017gpv.

iPTF14hls is in my initial sample, but because it is $\sim 1\,000$ days older than AT 2017gpv, it was not detected in ZTF. Assuming that these two events evolved similarly, iPTF14hls would have been close to the detection limit at the start of ZTF and faded below that before it could be picked up by my pipeline.

4.3.2 Siblings

Siblings are two (or more) transients that occur in the same host galaxy. While siblings can occur at any location in a galaxy (see e.g. Graham et al., 2022; Dhawan et al., 2024, for a sample of ZTF-detected siblings), a subset of these occur with a small enough sky separation that part of the light of one sibling can be detected when performing forced photometry at the sky position of the other sibling. In chapter 3, five such

sibling pairs were found, with both transients detected within ZTF. With my current sample being larger and spanning a bigger time range over which the second transient can be observed, it is reasonable to expect a larger number of same-location sibling transients. I arbitrarily define siblings as transients detected in ZTF *without* the use of additional binning to push the detection limit, that are distinctively separate in time from their pre-ZTF transient sibling counterpart and at a separation small enough to be picked up in forced photometry centered at the location of the pre-ZTF transients. The ZTF transients do not have to be classified. I find 12 pairs of transients that satisfy these conditions, which I verified through the Fritz broker [van der Walt et al. \(2019\)](#); [Coughlin et al. \(2023\)](#). These are shown in Table 4.3.

In three cases, the ZTF-detected transient was classified but mistakenly associated with a pre-ZTF transient in WISEREP or OSC that subsequently got the same classification. For AT 2017gcd, a decaying transient was observed in four epochs of unforced Pan-STARRS photometry spread out over 125 days. This is enough to conclude that the 2017 transient was real and likely some kind of SN, but not necessarily a SN Ia.

In the two other cases (AT 2017keg/SN 2019tka and SN 2013ld/SN 2021rgw), the pre-ZTF detections are spurious and are unlikely to be true sibling pairs. AT 2017keg was reported by ATLAS in 2019 ([Tonry et al., 2019](#)) when SN 2019tka was found. Due to SN 2019tka being close to the host nucleus, likely spurious detections from two years before were present in the ATLAS light curve, resulting in the automated discovery report stating the wrong discovery date. SN 2019tka was reported by ZTF ([Forster et al., 2019](#)), which did not have such earlier detections as it only started operating in 2018. SN 2013ld was reported in 2021 by Pan-STARRS1. Again, its sibling SN 2021rgw was on top of the host nucleus, which has had several epochs of minor variability. Even the internal ZTF name is from 2018, showing that ZTF also detected minor changes at the host nucleus location. When SN 2021rgw was found, ZTF issued an alert with the discovery date in 2021, resulting in SN 2021rgw ([Nordin et al., 2021](#)) while Pan-STARRS1 used the first unforced detection epoch in 2013 as the discovery date, resulting in SN 2013ld ([Chambers et al., 2021](#)).

Table 4.3: Pre-ZTF transients with a sibling transient detected in ZTF in single exposures. The first four columns give the name, type, discovery mjd, and redshift of the pre-ZTF transient, and the next five columns give the same information about its ZTF sibling when available. The second to last column gives an estimate of the peak r -band (unless specified otherwise) absolute magnitude of the ZTF transient assuming the pre-ZTF transient redshift, except for AT 2018iml as this is likely a foreground CV. The last column gives the separation between the siblings. A ? type means the transient was never spectroscopically classified. IAU names in parentheses are ZTF transients wrongly associated with a pre-ZTF transient.

Pre-ZTF transient				ZTF transient					
Name	Type	MJD	Redshift	ZTF name	IAU name	Type	MJD	Peak mag	Sep (")
AT 2017gcd	??	57968	0.028593	ZTF19acdghq	SN 2019aavr	Ia-norm	58763	−18.8	1.02
AT 2017keg	??	58068	0.05	ZTF19acihng	SN 2019tka	Ia-norm	58781	−19.1 ^(g)	1.69
SN 2013ld	??	56522	0.02741	ZTF18abavruc	SN 2021rgw	Ia	59393	−18.9 ^(g)	0.57
ASASSN-14ba	Ia pec/91T	56796	0.032668	ZTF22aaawghw	AT 2022csd	II	59629	−16.8	1.29
SN 2017acp	II	57785	0.0215	ZTF19aavqics	SN 2019gxo	II	58633	−18.1	3.15
SN 2014gz	II	56678	0.02558	ZTF21abcpbqd	SN 2021nof	II	59362	−17.2 ⁽ⁱ⁾	0.30
SN 2009hz	II	55046	0.0253	ZTF18acotwcs	AT 2018iml	CV?*	58439	—	0.33
iPTF15wk	Ia	57097	0.23	ZTF21aaeebxm	—	?	59226	−20.0 ^(g)	0.31
SN 2016bsc	Ia	57500	0.05	ZTF23aaeljse	(SN 2016bsc)	?	60043	−17.9	1.61
SNF20080522-001	Ia	54608	0.04872	ZTF22aalbuig	AT 2022kuh	?	59724	−17.3	1.64
PS15ctg	Ia	57328	0.078	ZTF22abamjrf	AT 2022rol	?	59806	−18.7	1.03
iPTF15eot	II	57357	0.039	ZTF19aakvysq	AT 2019bll	?	58541	−17.5	2.19

* Classified ZTF SNe that were wrongfully associated with the pre-ZTF name resulting in a misclassification. The pre-ZTF transient itself was never actually classified in the cases of AT 2017keg and SN 2013ld. The case of AT 2017gcd / SN 2019aavr has been fixed on TNS.

** Not officially classified but based on evidence gathered from the ZTF forced photometry light curve.

^(g) Peak absolute magnitude in the g -band.

⁽ⁱ⁾ Peak absolute magnitude in the i -band.

For the other nine sibling pairs listed in Table 4.3, the pre-ZTF sibling has a classification (five SNe Ia and four Type II SNe) obtained at the time of discovery. Three of their paired ZTF siblings were spectroscopically classified as Type II SNe at the time of discovery in ZTF. For AT 2018iml, the sibling of Type II SN 2009hz, the ZTF detections were obtained in two periods (November 2018 and March 2024) and they are best matched to a CV. The remaining five had no ZTF-era classification. In some cases the ZTF sibling is close enough to get wrongly associated with the pre-ZTF sibling and obtain the same IAU name.

Assuming the five unclassified ZTF transients occurred in the same host galaxy as their classified sibling, I can use the known redshifts to estimate the absolute brightness of the unclassified transients and show that they are in the range of SNe and unlikely to be caused by late-time CSM interaction. ZTF21aaeebxm is found around 1 800 days

after its sibling, iPTF15wk, was first detected, and has an inferred absolute g -band magnitude of -20.0 at its peak at just $0.31''$ from iPTF15wk. To get such a strong signal this long after the explosion would require an unreasonably large CSM and thin shell, making CSM interaction an unlikely explanation. Its location is consistent with its host nucleus, making a nuclear transient origin likely.

For ZTF23aaeljse, ZTF22aalbuig, ZTF22abamjrf, and ZTF19aakvysq, their absolute peak magnitudes are in the range typical of SNe (-17.3 to -18.7 mag), and SNAP shows that there is a noticeable separation between the pre-ZTF and ZTF-detected signals ($1.03''$ to $2.19''$), disfavouring any explanation that would require the two events to be at the same spatial position, such as CSM interaction. The inferred CSM masses to explain these events in this way are also unrealistically high. Their locations are inconsistent with their host nuclei. Therefore, I conclude that a sibling transient, likely an unclassified SN, is the most plausible explanation for these events.

4.3.3 Non-SN sources of flux excesses in ZTF

I identify some late-time excesses at the positions of my sample that are astrophysical but are due to known AGN activity or due to stellar variability. Ten transients (listed in Table 4.4) in my sample have been found to have genuine long-term variability due to an AGN that lasts for the whole time period of ZTF and is picked up by my pipeline. In each case, the source is a known AGN (verified through cross-referencing with SIMBAD, ZTF, and checking the AGN criterion from Hviding et al. 2022 using the *Wide-field Infrared Survey Explorer* WISE Wright et al. 2010). Their light curves are shown in figure 4.5. While several of the pre-ZTF transients in this category have been classified as SNe II, some have high redshift, which suggests that these are AGN misclassified as SNe II.

I also recovered four long-term variable stars that are shown in figure 4.6 and listed in Table 4.4. The nature of these sources make it impossible to create a template that always subtracts them completely without leaving residual flux or a ghost as there may be no region without variability. However, if the variations in the source’s magnitude

are large enough over a long period of time, it will easily be picked up by the binning algorithm.




Figure 4.5: Binned flux of the recovered AGN. The flux is calibrated to a zero point at mag 30. The dashed lines show the baseline value for each object. Since the observations used to determine the baseline cannot be binned, I show the unbinned data in the baseline region of each object. This region is marked by the black lines. Even though not all objects are properly sampled over the entire lifetime of ZTF, they all clearly show variability over long timescales. As these objects are always varying, it is impossible to do a baseline correction without the transient present. This causes some light curves to go below their baseline. Note that the values for AT 2017kas have been divided by 10 as the variability is very large.




Figure 4.6: Binned flux of the recovered variable stars. The flux is calibrated to a zero point at mag 30. The dashed lines show the baseline for each object. Since the observations used to determine the baseline cannot be binned, I show the unbinned data in the baseline region of each object. This region is marked by the black lines. All three objects vary around their central values, but the amplitude varies widely between the ZTF bands. Note that the values for PTF10qpf have been divided by 100 as the variability is extremely large.

Table 4.4: Pre-ZTF transients detected in the binned ZTF light curves due to AGN or variable star activity. The original type and z values are as they were recorded on the OSC or WISEREP. The late-time type gives the reason for the ZTF detections.

Name	Pre-ZTF type	z	Late-time type
Gaia14adg	II	0.154	AGN
SN 2016fiz	II	0.05	AGN
SN 2017avb	II	0.096	AGN
LSQ12biu	II _n	0.136	AGN
SN 2017bcc	SLSN-II	0.133	AGN
PSN J0151 ⁽¹⁾	SLSN-II?/AGN	0.26	AGN
ATLAS17khl	AGN	0.06	AGN
PS17bgm	AGN	0.358	AGN
LSQ12ehj	AGN	0.12	AGN
AT 2017kas	SN candidate	0.031328	AGN
LSQ12fgx	Variable star	0	Variable star
SNhunt44	LRV? ⁽²⁾	0.0005864	Variable star
AT 2016ijb	SN candidate	−0.000781	Variable star
PTF10qpf	Variable star	0	Variable star

⁽¹⁾ PSN J0151 = PSN J01510869+3155215

⁽²⁾ Potential luminous red variable

4.3.4 Final shortlist of late-time interaction

The final nine objects in my sample cannot be put in any of the groups (ongoing SN flux, sibling transient, nuclear activity, variable star) above. Each of these events is discussed individually below. Table 4.5 shows the general information of these events, and their light curves are shown in figure 4.7. The bin sizes and placement chosen for the plots are those that result in the clearest and cleanest signals. The baseline regions used are those that have the least amount of transient flux in them, which means they are the furthest away in time from the excess. As was done in chapter 3, I compare the ZTF detections to several classes of transients in an attempt to explain them as a previously unidentified sibling transient. I also estimate the potential CSM mass required to explain the late-time signature using the analysis method described below for objects where a CSM interpretation cannot be ruled out.

Table 4.5: Objects with detections whose origin was not immediately clear. The first five columns give information on the original transient, and the last six columns give an overview of the binned ZTF detections. Start gives the time after discovery of the first detections in ZTF, and duration gives the length of these detections. Both are in rest frame days and rounded to 10 days. The g , r , and i columns give the range of detections in absolute magnitude if detected. Inferred cause of excess gives the type of object that best explains the detections. The last two columns state whether the excess is consistent with the SN and host nucleus location, respectively.

Pre-ZTF					ZTF excess						
Name	Type	Discovery date	z	Host sep (")	Start (d)	Duration (d)	g band abs. mag	r band abs. mag	i band abs. mag	Inferred cause of excess	Excess consistent with SN host nucleus
SN 2017ige ⁽¹⁾	Ia	17-11-17	0.02431	0.99	1130 ⁽²⁾	90	-13.5 – -13.3	-14.3 – -13.6	-	Sibling	yes no
SN 2016cob	Ia-91T	26-05-16	0.02961	0.7	2080	480 ⁽³⁾	-15.3 – -14.7	-16.5 – -16.1	-17.2 – -16.4	CSM/Nuclear	yes yes
SN 2017fh	Ia	17-07-17	0.032188	0.0	230 ⁽²⁾	180 ⁽³⁾	-15.6 – -14.9	-16.0 – -15.6	-	CSM/Nuclear	yes yes
SN 2017fby	Ia	01-07-17	0.043513	0.73	1780	190 ⁽³⁾	-	-16.5 – -16.1	-	CSM/Nuclear	yes yes
PTF10jtp	Ia	04-06-10	0.067	1.11	2670	200 ⁽³⁾	-	-16.5 – -16.3	-	Nuclear	yes yes
AT 2017fwf	Cand.	01-08-17	0.033707	1.67	1690	410	-	-14.7 – -14.0	-	Nuclear	no yes
PScl30283	Ia	29-01-11	0.07622	0.44	3750	480 ⁽³⁾	-	-17.8 – -16.9	-	Nuclear	yes yes
PTF13cow	Ia	07-08-13	0.086	0.56	2980	560 ⁽³⁾	-16.3 – -16.1	-17.4 – -17.1	-	Nuclear	no yes
iPTF15aow	Ia	06-05-15	0.07597	0.92	1970	1050 ⁽³⁾	-16.0 – -15.9	-17.2 – -16.4	-18.0 – -17.1	Nuclear	yes yes

⁽¹⁾ SN 2017ige has two separate periods with detections. The second is shown in the table while the first is consistent with the radioactive tail phase of the SN Ia. It starts at 80 days after discovery (the start of the ZTF survey) and lasts for 150 days declining from -14.6 mag to -13.8 mag over this time in the g -band.

⁽²⁾ Detected from the start of ZTF.

⁽³⁾ Start or end with a gap in the binned observations or at the edge of available data.




Figure 4.7: The nine objects whose late-time detections required a more rigorous investigation. Each object is shown in absolute and apparent magnitude (left and right axes respectively) and rest frame days since first detection and mjd (top and bottom axes respectively). I show the binned ZTF data with the late-time detections as well as pre-ZTF detections from other surveys (Pan-STARRS, Gaia, ASASSN, PTF) where available. For SN 2017ige a heavily extinct ($(E(B-V))_{\text{host}} = 1.2$ mag) light curve of SN 2020jfo is shown with dashed lines matching the late-time ZTF r -band detections. The objects are marked with the explanation for their late-time detections.

4.3.4.1 Estimating CSM masses

To determine if the objects in my final shortlist are potentially due to CSM interaction signatures, I can make a rough estimate of the CSM mass required to generate such a signal. In this section, I describe the method used by [Graham et al. \(2019b\)](#) to put constraints on the CSM mass of their targeted SNe at similar epochs to mine, based on the observed near-ultra-violet (NUV) (non-)detections using the models of [Harris et al. \(2016\)](#). While CSM is line-dominated, especially by $H\alpha$, it is much more difficult to estimate a CSM mass without making many assumptions about the state of the CSM. Even though the statistical errors can be large, the resulting CSM masses for some of the objects in section 4.3.4 is still unreasonable.

[Harris et al. \(2016\)](#) assume a CSM shell of constant low density $\rho_{\text{CSM}} < 10^{-17} \text{ g cm}^{-3}$ extending from R_{in} to R_{out} such that the fractional radius F_R of the shell is $F_R \equiv R_{\text{out}}/R_{\text{in}}$. $F_R = 1.1$ represents a thin nova-like shell, with higher values representing thicker shells. The ejecta impacting the CSM are assumed to have a density profile $\rho_{\text{ej}}(r, t)$ described by

$$\rho_{\text{ej}}(r, t) = \frac{0.124 M_{\text{ej}}}{(v_t t)^3} \left(\frac{r}{v_t t} \right)^{-10}, \quad (4.1)$$

where M_{ej} is the ejecta mass, t is the time since explosion, r is the distance to the explosion site, and v_t is the transition velocity between the outer and inner ejecta, with typical values for SNe Ia being $v_t = \mathcal{O}(10^9 \text{ cm s}^{-1})$. After the shock the CSM is compressed, heated, and starts emitting light. This shocked CSM is assumed to be fully ionized H-dominated gas with the electrons and ions in equipartition.

Assuming that the time of the late-time detections corresponds to the peak of the CSM interaction light curve, [Graham et al. \(2019b\)](#) showed that the Bremsstrahlung spectral luminosity L_ν at frequency ν at the moment the interaction shock reaches the outer edge of the shell is

$$L_\nu \approx 1.63 \times 10^{-31} T^{-1/2} t_r^{-3} M_{\text{CSM}}^{17/7} e^{-\frac{h\nu}{kT}} y(F_R) \text{ erg s}^{-1} \text{ Hz}^{-1}, \quad (4.2)$$

where T is the temperature of the shocked material in Kelvin, $t_r = t/(1+z)$ is the time after explosion in seconds in the rest frame of the SN at which the ejecta reach the outer edge of the CSM shell and the interaction is assumed to be at its strongest, M_{CSM} is the CSM mass in grams, and $y(F_R) = F_R^{-3/7}(1 - F_R^{-3})^{-10/7}$ is the dependence on the fractional radius of the shell.

Graham et al. (2019b) derive Eq. 4.2 to use in the NUV and assume temperatures around 10^8 K. In these conditions the exponential term is near unity and can be ignored. Since I am targeting optical wavelengths, if the mechanism described here is the source, lower temperatures are needed where the exponential term cannot be ignored. The other assumptions made in Graham et al. (2019b) still hold for optical wavelengths, so I can use Eq. 4.2 to estimate the M_{CSM} required to explain any detected late-time signals using

$$10^{-0.4M} = \frac{L}{L_0} = \frac{L_\nu c}{L_0 \lambda}, \quad (4.3)$$

with M the absolute magnitude of the detected signal, c the speed of light, $L_0 = 3.0128 \times 10^{35} \text{ erg s}^{-1}$ the zero point luminosity, and λ the effective wavelength of the band I am considering (4746.48 Å, 6366.38 Å, and 7829.03 Å for g , r , and i , respectively). By combining Eq. 4.2 and Eq. 4.3 I get

$$\frac{M_{\text{CSM}}}{M_\odot} \approx 3 \times 10^{-8 - \frac{14}{85}M} \left(\frac{T}{\text{K}} \right)^{\frac{7}{34}} \left(\frac{t/(1+z)}{\text{days}} \right)^{\frac{21}{7}} \left(\frac{\lambda}{\text{Å}} \right)^{\frac{7}{17}} e^{\frac{7hc}{17kT\lambda}} y(F_R)^{\frac{-7}{17}}. \quad (4.4)$$

I can obtain M , t , and z directly from my shortlist of transients with potential CSM interaction. For t the start of the signal should be taken as this assumes the least delay due to a non-negligible light-crossing time across the CSM shell. The remaining parameters are the fractional radius of the shell and the temperature.

It is also possible to estimate the required efficiency η for converting the kinetic energy from the ejecta E_{kin} in radiated energy E_{rad} to sustain the observed late-time signal by requiring $E_{\text{kin}} = \eta E_{\text{rad}}$, similar to what was done in Fraser et al. (2013a). When expanded, this can be written as

$$E_{\text{kin}} = \frac{M_{\text{CSM}} v_{\text{ej}}^2}{2} = \eta L_0 t_{\text{rad}} 10^{-0.4M} = \eta E_{\text{rad}}, \quad (4.5)$$

when assuming the CSM is hitting the ejecta at v_{ej} in the rest frame of the ejecta, as sweeping up CSM is equivalent to the CSM slowing down in this frame. t_{rad} is the time over which the interaction signal is detected. Solving Eq. 4.5 for M_{CSM} leads to

$$\frac{M_{\text{CSM}}}{M_{\odot}} \approx 2.6 \times 10^{-3-0.4M} \eta \frac{t_{\text{rad}}}{\text{d}} \left(\frac{v_{\text{ej}}}{\text{km s}^{-1}} \right)^{-2}. \quad (4.6)$$

Assuming $\eta = 1$, $v_{\text{ej}} = 5000 \text{ km s}^{-1}$, and that an absolute magnitude $M = -16$ is being observed for $t_{\text{rad}} = 100 \text{ d}$, an $M_{\text{CSM}} \approx 0.03 M_{\odot}$ is required to sustain this signal. The efficiency is likely much lower as it needs to account for the CSM possibly being patchy, not all of the kinetic energy being transferred, and not all of the transferred energy being radiated away. However, even when $\eta = \mathcal{O}(10^{-2})$ an observable signal can still be achieved with a massive but feasible amount of CSM.

$F_R = 1.1$ if I assume the CSM shell to be nova-like. Equation 4.4 easily gives very massive CSM shells, especially if there is a significant delay before the onset of the CSM interaction. The required M_{CSM} can be somewhat lowered by assuming that the actual thickness $R_{\text{in}} - R_{\text{out}}$ of the CSM shell remained constant after its creation, meaning that F_R decreases as the shell travels outwards. Since the SN ejecta move at a constant velocity, the distance of the CSM is proportional to the delay time of the interaction signal. Assuming a shell has $F_R = 1.1$ if the late-time signal occurs 300 days after the explosion, that same shell will have $F_R \approx 1.01$ if the interaction starts 3 000 days after the explosion. The main reason for allowing $F_R < 1.1$ is to see how thin a shell would need to be to give a reasonable M_{CSM} . I assume temperatures of 10^5 to 10^7 K , lower than the 10^8 K assumed by [Graham et al. \(2019b\)](#), but more suitable for explaining optical emission.

Equation 4.4 assumes that the magnitude is corrected for extinction. To get the lowest M_{CSM} possible I assume there is only Milky Way extinction in the line of sight of any of my targets. Any extinction in the host galaxy would make the intrinsic colour of

these objects bluer allowing for higher temperatures, but also raise the intrinsic absolute magnitude, resulting in a higher overall estimate for M_{CSM} . The estimates should therefore be seen as lower limits.

SN 2017ige. SN 2017ige was initially discovered and classified as a SN Ia in late 2017, only a few months before the start of ZTF. As it was relatively close by, the radioactive tail phase was identified in the ZTF binned data out to ~ 200 days, before it faded below the noise limit. This tail matches well with the declining light curve seen in the early (< 100 days) Gaia data (see figure 4.7). SNAP shows that the SN is in the ZTF reference images, leaving a small ghost at its location in the difference images.

At 1000 days after the SN tail faded below the noise limit, a small excess is detected in both the ZTF g - and r -bands that lasts around 90 days. In the difference images a small excess can be seen semi-overlapping the negative imprint left by SN 2017ige during this excess. Such a dipole signal would usually suggest imperfect image subtractions, but in this case it can also be interpreted as a separate transient slightly offset from the location of SN 2017ige being over-subtracted due to the presence of a SN in the reference images.

The late-time excess has been detected by Pan-STARRS in their i -band, given the internal name PS21aos, and linked to SN 2017ige based on the sky position, confirming that the excess found in ZTF at late times is real. The Pan-STARRS i -band detections are brighter than what was found in the ZTF g - and r -bands, likely due to a combination of the ZTF detections being underestimated as a result of the ghost of SN 2017ige, as well as the late-time excess being intrinsically red ($g - r \approx 1$ mag).

I use the duration and shape of the identified excess to test if it could be explained as a sibling transient. A light curve similar to that of SN 2020jfo (Sollerman et al., 2021; Ailawadhi et al., 2023), a Type IIP SN with a relatively short plateau, fits the excess quite well in duration and absolute magnitude when a host extinction of $E(B - V) = 1.2$ mag is added. Given that the sky position at which the forced photometry was




Figure 4.8: M_{CSM} as a function of the r -band absolute magnitude M for different assumptions of t (panels), T (shaded region), and F_R (colours). The five objects for which I estimate M_{CSM} in section 4.3.4 have been marked in the panel that has t the closest to the detected signal. Decreasing shell thickness and temperature to mitigate M_{CSM} is effective, though there is a limit to far this can be done while keeping realistic values for F_R .

performed is only $0.99''$ from the host nucleus of an edge-on galaxy, such a high amount of extinction is plausible. The original peak of the light curve of SN 2017ige in 2017 was not caught so a comparison with a similar extinction estimate cannot be made. Although the late-time excess is close to the nucleus, inspection of the images shows that it is offset from the host centre as well as the original SN position. Therefore, I conclude that a Type II SN sibling with a relatively high value of extinction is an adequate explanation for this late-time excess.

SN 2016cob. SN 2016cob was classified as a SN Ia-91T at early times. There is light curve data from Pan-STARRS and ASASSN around peak and extending to just past 50 days after the peak. In the binned ZTF data, this event jumps up from non-detections to detections in all three bands around ~ 2150 days after first detection (figure 4.7). The i -band data is included for completeness but does not have a reliable baseline estimate (>30 points per region) so the values should be treated with caution. The detections are present in all bands at the same time and extend for at least 480 days at absolute magnitude of -15 to -17 mag. No SN-like transient can explain this long-lived flux excess.

Following the calculation in section 4.3.4.1, if I assume a CSM shell with a fractional radius of $F_R \approx 1.005$ at a temperature of 10^5 K, this gives $M_{\text{CSM}} \approx 6 M_\odot$. In figure 4.8 I show the M_{CSM} as a function of the r -band absolute magnitude, M_r , for three different times of the onset of the CSM interaction (300, 2000, 3000 days), three different F_R values (1.1, 1.01, and 1.005) and temperatures of 10^5 to 10^7 K. A CSM mass of $M_{\text{CSM}} \approx 6M_\odot$ is large, but around the same amount as suggested to explain the interaction signatures seen in SN 2002ic (Hamuy et al., 2003b). Based on this, I put SN 2016cob forward as a candidate for showing late-time CSM interaction, though I note that it would require a very thin shell ($F_R = 1.005$).

SN 2016cob is $0.7''$ from the host nucleus, putting the nucleus on the same or an adjacent pixel in the ZTF observations. This leaves room for a host variability interpretation, though the host has no AGN according to its WISE colours. SNAP does not show any reduction or subtraction issues that could explain the ZTF detections as a spurious source either.

The absolute magnitude, duration, and nuclear location are consistent with the ambiguous nuclear transient (ANT), ASASSN-20hx/AT 2020ohl (Hinkle et al., 2024). ANTs are events that cannot be easily classified into AGN activity or TDE (Wiseman et al., 2025). AT 2020ohl displayed a plateau or very slow decline in its optical light curves for > 250 days relative to peak with an approximate plateau magnitude in the gri -bands of -17.5 mag. In figure 4.9, I compare the light curves of the shortlisted transients that are consistent with their host nuclei compared to AT 2020ohl. The flux excess at the position of SN 2016cob is shown as blue squares and is slightly fainter than AT 2020ohl. AT 2020ohl had an observed rise-time of 30 days (Hinkle et al., 2024), slightly larger than my smallest bins. This would explain why no rise is detected, as my time resolution is too poor to detect time variations of this scale. I can only observe a sudden appearance of a very flat light curve. Both late-time CSM interaction and an ANT are adequate explanations for the identified signal, and without additional information I cannot decisively point at one of these two explanations.




Figure 4.9: Binned r -band observations of the objects whose late-time detections are consistent with a galaxy nucleus (see Table 4.5) and are potentially transient events unrelated to the original SN. Upper limits are shown with downward arrows. The r -band light curve of the slowly evolving ambiguous nuclear transient AT 2020ohl (ASASSN-20hx, [Hinkle et al. 2022](#)) is also shown for comparison. As the late-time signal in AT 2017fwf was found to be significantly offset from the SN location, I show the processed light curve at the host nucleus location instead.

SN 2017frh. SN 2017frh was discovered by ASASSN and spectroscopically classified as a SN Ia. An excess was detected in the binned ZTF g - and r -bands at absolute magnitudes of -15.3 and -15.8 , respectively. In the i -band only upper limits were obtained. The signal was visible from the start of ZTF at a phase of 240 days after discovery and lasted for 180 days before disappearing behind the Sun. Nothing is detected after its return. The detections are found regardless of the used baseline region.

These detections are inconsistent with a normal SN Ia at these phases, which would have a significantly lower absolute magnitude ($M \approx -12$ mag) and also keep fading over time. If it were some type of sibling transient it would have to plateau at -15 to -16 mag for at least 180 days. No known class of SN shows this behaviour without the help of CSM interaction. The Luminous Red Nova AT 2021biy ([Cai et al., 2022](#)) does show a plateau of a similar duration but is much fainter than the late-time signal in SN 2017frh.

The late-time detections in SN 2017frh have a similar absolute magnitude and duration to the three late-time CSM interaction candidates that were presented in chapter 3. However, the interaction in SN 2017frh starts significantly earlier (~ 240 days) than for those events. If I assume $F_R = 1.1$, a M_{CSM} of $1.8 M_{\odot}$ would be required

to explain both the g - and r -bands assuming $T = 10^5$ K. Raising the temperature to $T = 10^7$ K gives $M_{\text{CSM}} \approx 4 M_{\odot}$ for the g -band and $M_{\text{CSM}} \approx 4.5 M_{\odot}$ for the r -band. SN 2017frh is shown in the left-hand panel of figure 4.8. If the F_R is reduced to 1.01, then the CSM mass is lowered to between 0.5 and $1.2 M_{\odot}$. These mass estimates are in line with the suggestions of previous interacting SN events (Dilday et al., 2012; Inserra et al., 2016). For these reasons, I classify SN 2017frh as a candidate SN with late-time CSM interaction.

The late-time excess at the position of SN 2017frh is also on top of the host nucleus, which could mean that the late-time signal is instead related to the host. However, the duration of the detected signal is the shortest of those shown in figure 4.9, although there is an observing gap at the end of the detected signal so no strict limit can be placed. Late-time CSM interaction signatures have been found in SNe Ia at similar phases (Graham et al., 2019b). Therefore, given the relatively low CSM mass required and the time frame of the interaction, I prefer the CSM interpretation for this late-time signal.

SN 2017fby. SN 2017fby was discovered by ASASSN and spectroscopically classified as a SN Ia. The ZTF binned r -band detections, at an absolute magnitude of -16.3 mag, are found when the sky location returns from behind the Sun at 1 860 days after discovery and lasts for 200 days until it becomes unobservable again. The light curve stays roughly constant over this time frame. There are hints of a g -band excess as well but only one bin is above the 5σ threshold. After it comes back from behind the Sun nothing is detected anymore.

With a much stronger r -band signal compared to the g -band, one could argue for this being a reddened previously unknown sibling transient whose rise was missed due to it occurring while the sky position was too close to the Sun to be observed. However, the long plateau in the light curve of at least 200 days rules this out, as Type IIP SNe have plateaus that are generally significantly shorter than this (Barbon et al., 1979; Anderson et al., 2014).

CSM interaction with H-rich material could explain the strong r -band signal through strong $H\alpha$ emission. If I assume the detections could come from a CSM shell with a fractional radius as low as $F_R \approx 1.005$, this gives a required CSM mass of $M_{\text{CSM}} \approx 5.5 M_{\odot}$ according to the calculation in section 4.3.4.1. This is around the upper end of the CSM masses that have been estimated for known SNe Ia-CSM, such as SN 1997cy (Chugai & Yungelson, 2004), SN 2002ic (Chugai & Yungelson, 2004; Inserra et al., 2016), and SN 2012ca (Inserra et al., 2016). Therefore, I note SN 2017fby as a SN with potentially late-time interaction with a thin, massive shell of CSM.

Checking the difference images with SNAP shows a clear bright spot at the SN location during the period of detections and no clear issues before or after it, confirming the excess to be real. At a distance of 0.73" (0.66 kpc at the host redshift) from the host nucleus, a TDE or other nuclear transient would likely appear on the same or adjacent pixel in the ZTF images. Like for SN 2016cob, this makes it difficult to say whether late-time CSM interaction or an ANT is the best explanation for the signal (figure 4.9).

PTF10jtp. PTF10jtp was discovered by PTF and classified spectroscopically as a SN Ia. Detections were made in the ZTF binned r -band light curves starting at 2850 days after discovery. It had a ~ 200 day r -band plateau around -16.4 mag before the sky position became unobservable. When it came back from behind the Sun, there were 5σ upper limits that were about one magnitude deeper than the previous detections (not shown in figure 4.7).

The g -band has upper limits at approximately -15.5 mag during this entire period, while the i -band bins are not deep enough for a constraining upper limit to be placed. The baseline correction in r is significant and comparable between the later two baseline regions, but is somewhat smaller in the first region as it partially includes the period in which the excess was detected. This suggests that the excess was present in (some of) the reference images, and the late-time excess was present for longer than was found with the binning procedure. SNAP confirms this, as a slight excess can be seen in the difference images, followed by a small ghost at later times.

As this excess occurs over eight years after the SN, any SN Ia radioactive tail contribution at this magnitude can immediately be ruled out. Although there are only detections in one band, the limits placed on the colour together with the > 200 day plateau limit the possibilities for a sibling, even with moderate values of extinction.

An unreasonably thin shell of CSM is needed to explain detections that are this late and still have an absolute magnitude of $M \approx -16.5$ mag while keeping the total mass of to be at most $M_{\text{CSM}} \approx 5 M_{\odot}$ in line with literature events. A more realistic shell thickness results in a CSM mass of at least $10 M_{\odot}$. Therefore, I rule out CSM interaction as a likely explanation for this late-time flux excess.

The sky location is $1.11''$ from the host nucleus, and SNAP also shows that the excess is slightly offset from the SN location in most frames and consistent with the host nucleus. The long duration of the excess, combined with the central location, suggests host activity is the cause, although it is not an AGN according to its WISE colours. Therefore, a nuclear transient is the most likely explanation for these late-time detections. Its light curve is shown compared to the ANT AT 2020ohl in figure 4.9, where it is seen to be plausibly consistent in duration and absolute magnitude.

AT 2017fwf. This object was only detected by Gaia at early times with no spectrum obtained, preventing even a speculative classification based on the photometry. It is detected in ZTF binned r -band data starting at 1750 days after discovery and lasts for 420 days. The light curve rises from -14 mag to nearly -14.7 mag before slowly fading again. Despite having upper limits around the same magnitude as r , nothing was detected in g . The i -band limits are shallower and less constraining. SNAP shows that the excess is real, but located at the nucleus location $1.67''$ from the SN. Therefore, CSM interaction can be ruled out. The long timescale of the transient also makes a SN-like transient unlikely.

Performing forced photometry at the host nucleus location reveals its properties much more clearly, being visible for nearly 700 days and reaching -15.2 mag at its brightest (see figure 4.9). This is longer by nearly 300 days and brighter by nearly

0.5 mag than the values given in Table 4.5 for the forced photometry at the SN position. The host galaxy is not an AGN according to its WISE colours. The excess is still very faint for a TDE unless it is heavily reddened, although a reddened TDE would likely have been bright enough in the i -band to still be visible. However, an ANT could be possible, though the signal is ~ 3 mag fainter than AT 2020ohl but also visible for over twice as long. Faint nuclear variability could also be possible, and there is some scatter in the light curve, but with my 100-day bins for the light curve, small-scale variability cannot be identified. Therefore, I can only conclude that some sort of nuclear variability is present.

I note that a sibling transient is also detected by ZTF in this host galaxy (SN 2020ackb, SN IIP), though it is at a distance of $5.23''$ and has no effect on the forced photometry at the location of AT 2017fwf. The time at which the sibling was visible also does not correspond with the detected host variability.

PSc130283. PSc130283 was detected by Pan-STARRS and spectroscopically classified as a SN Ia. At 4030 days after the discovery, ZTF-detections are found in the r -band at an absolute magnitude of -17 mag before brightening again to -17.8 mag in the next bin and staying there with variation ≤ 0.5 mag for the remainder of the light curve (at least 520 d). The non-detections before the start of the excess were up to 1.5 mag lower, suggesting that the excess started suddenly rather than a gradual brightening. During the entire late-time excess, the g -band stays with 5σ upper limits at -16 mag. At a redshift of $z = 0.07622$ the $H\alpha$ line is shifted into the overlap between the r - and i -bands. If the late-time signal is due to interaction with H-rich CSM, these detections could be a sign of strong $H\alpha$ emission pushing the brightness in these bands.

Detections this late ($\gtrsim 3750$ rest frame days) and bright ($M \sim -18$ mag) cannot give a reasonable M_{CSM} in Eq. 4.3.4.1 unless no extinction, a temperature around $10^{4.5}$ K, and $F_R < 1.0005$ is assumed, giving an unreasonably thin shell. Therefore, I rule out CSM interaction as a likely cause of the late-time flux excess.

SNAP shows a clear excess that is consistent with the SN location, and no ghost or residuals before the excess starts or image defects that could explain these detections. However, the flux excess is only $0.44''$ from the host nucleus, although the host is not an AGN according to its WISE colours. The host has a history of small variability that causes sparse detections (including a detection that put its discovery date over 400 days before the SN explosion), which points to the ZTF detections in the binned data being host-related. Out of all excesses plotted in figure 4.9, PSc130283 is the brightest. It is very similar in brightness to AT 2020ohl though a bit longer in duration. These properties, combined with its relatively sharp rise suggest nuclear variability, such as an ANT, could explain these detections.

PTF13cow. PTF13cow was discovered by PTF and classified as a SN Ia. The late-time detections in the ZTF *gri*-bands begin 3 230 days after the discovery at absolute magnitudes of -16.0 mag and -17.3 mag in the *g*- and *r*-bands, respectively. There are detections in the *i*-band, but the baseline is too small for these to be considered any further here. The *r*-band detections last for at least 610 days. The *g*-band is only detected at the start but is significantly fainter and disappears below the detection threshold earlier. This long timescale rules out SN-like transients as a cause for the excess.

SNAP shows that there is a clear excess in the images, and despite the host nucleus distance being only $0.56''$ from the SN location, the SN location and host nucleus are on different pixels. The excess has a preference to be at the host nucleus location, suggesting that it is the cause of the signal. However, the host nucleus is not an AGN according to its WISE colours. The duration of over 560 d and a peak absolute *r*-band magnitude of -17.4 mag of the identified late-time excess is similar to AT 2020ohl, meaning that an ANT could explain these detections.

***iPTF15aow*.** *iPTF15aow* was discovered by the intermediate PTF survey (*iPTF*) and also detected in Pan-STARRS, and classified as a SN Ia. The late-time detections begin in the *r*-band at 2 120 days after discovery at -16.4 mag. There are only upper limits in the *g*- and *i*-bands at this time, but at later times some detections in these bands are made. The detections last for at least 1 050 days and slightly brighten over time to -17.2 mag in the last bins. The long-lived nature rules out SN-like transients at the same position as the original SN Ia.

Even when assuming a scenario where the *r*-band interaction stays at the level it was discovered at for the entire duration it was detected, $> 5 M_{\odot}$ of CSM is required in a very thin, low-temperature shell. However, the signal increases in strength over time, which would increase the required CSM mass. On top of that, no known SN Ia with CSM interaction has interaction as long as the late-time excess is found for or with the strength increasing over time, further disfavoured the CSM explanation.

The host nucleus is close by at $0.92''$ but from SNAP the transient does not look to significantly favour the SN or host nucleus location over the other. The host is not an AGN according to its WISE colours ($W1 - W2 = 0.077$ mag, $W2 - W3 = 1.564$ mag).

Pan-STARRS has some detections in the unforced photometry at the SN location, the first being at the time of *iPTF15aow*'s discovery. However, it has no internal name. In the years after the first detection, there are some detections in their *i*-band, but these are likely spurious because of a significant number of upper limits in between them that are not shown in figure 4.7. The number of Pan-STARRS *i*-band detections increased significantly after $\sim 2\,900$ days after discovery, which corresponds to the brightest part of the binned ZTF detections.

The duration, shape, and close proximity to the host nucleus of the excess point to it being most likely host-related. It has a similar brightness to AT 2020ohl, but has a several times longer duration and is still ongoing. The fact that it continues to brighten is also very unusual. None of these signatures are consistent with what would be expected from a late-time CSM interaction.

4.4 Discussion

In this study I found 98 cases with detections that cannot be explained as false positives due to observational, reduction, or software issues. These objects can be split into four groups: i) ongoing signatures of bright and/or nearby transients that were still detectable at the start of ZTF, ii) sibling transients at nearly the exact same location in the sky, iii) known variable sources such as AGN and variable stars, and iv) nine late-time flux excesses that required a more in-depth examination. Of these nine, I concluded that one was a sibling SN, five cases were nuclear transients close to the SN locations, two objects where it is unclear whether the ZTF signal is host or CSM-related, and one SN whose ZTF-detections are most consistent with late-time CSM interaction.

4.4.1 Rarity of late-time signals from SNe

I started my search for late-time signals based on the positions of 7718 transients discovered before ZTF. Table 4.6 shows the number of objects that had detections in ZTF, split over the six main types of SNe I distinguish between in my sample. In the sections below, I discuss the main conclusions for each of the likely classes (siblings, nuclear, CSM) of the late-time signals detected.

4.4.1.1 Siblings

Within the uncertainties, the rate of siblings I found is the same for the different groups of objects, with the SN Ia and Type II SN classes having rates of sibling discovery of 0.1 – 0.2%. The percentage is similar, given the large uncertainties associated with low numbers of events, to the percentage of siblings found in chapter 3 over a smaller time frame. These rates can be taken as lower limits because of observing gaps, which would result in transients being missed. In total I found 10 siblings, which gives a lower limit on the rate of siblings of $0.13 \pm 0.04\%$. SN Ia sibling pairs along the same line of sight may be useful for constraining the origin of reddening in SN Ia light curve

fitting for cosmology (Dhawan et al., 2024), but a larger sample would be required for meaningful constraints.

Table 4.6: Number of recovered signals for each SN type in my sample. As it is unclear whether SN 2016cob and SN2017fby should be counted as nuclear transients or due to CSM interaction, they have been counted in both.

Type	No.	Siblings		Nuclear		CSM	
		No.	%	No.	%	No.	%
Ia	4991	6 ^(a)	0.1	7	0.1	3	0.1
II	1311	4 ^(b)	0.3	0	0	0	0
Ib/c	339	0	0	0	0	0	0
IIn	179	0	0	0	0	0	0
I	58	0	0	0	0	0	0
Ibn	17	0	0	0	0	0	0

^(a) This includes the unconfirmed but likely sibling at the position of SN 2017ige determined from my binned light curve analysis.

^(b) This includes one sibling pair (SN 2009hz/AT 2018iml) where the ZTF transient was a likely CV based on its light curve.

4.4.1.2 Serendipitous nuclear transient detections

I identify five late-time signals that were most likely due to host activity that was (partially) picked up in the forced photometry at the SN location. This was determined due to a combination of the shape and duration of the late-time light curve, as well as the inferred M_{CSM} from the absolute magnitude and delay between the main SN and late-time detections being unrealistically high in the best-case scenario. I also identify signals at the positions of two events (SN 2016cob and SN 2017fby), where the inferred cause could be CSM or nuclear activity, bringing the total to seven potential nuclear events (see Table 4.6).

In figure 4.9, I show the late-time excess light curves for which nuclear transients cannot be ruled, along with the late-time light curve of SN 2017frh for completeness, although I prefer a CSM interpretation for it. The light curves are shown as absolute magnitude in the r -band against the time since the first binned data where a flux excess is identified. I also show the ANT AT 2020ohl (Hinkle et al., 2022) for comparison. AT 2020ohl is slightly brighter than most of these transient events, apart from PSc130283.

My events are also longer lasting, ranging from ~ 200 days up to ~ 1100 days. This could mean that there is a previously unknown population of lower luminosity ANTs that require deeper (binned) observations to be detected.

[Wiseman et al. \(2025\)](#) showed that their sample of ANTs all have changes in their WISE $W1$ - and $W2$ -band observations, with a slight delay compared to the optical. The WISE light curve for my potential nuclear transients were generated by Prof. Phil Wiseman and investigated by me, but unfortunately, these light curves use the total measured flux from the source. As the hosts of my nuclear transients have an apparent magnitude of around 15 mag in these bands and the optical variability I have identified is ≥ 5 mag fainter than that, any MIR variability is likely to be heavily suppressed into the noise of the WISE light curves. [Mo et al. \(2025\)](#) have recently built a image-subtraction pipeline for WISE / NEOWISE data, which could be used to solve this problem.

CL-AGN show variability on similar timescales as those that I have found here ([Ricci & Trakhtenbrot, 2023](#)). However, CL-AGN are a subset of AGN and the galaxies hosting my transients are not AGN according to the WISE criterion from [Hviding et al. \(2022\)](#). My objects are also much lower luminosity than seen for CL-AGN, which also makes this classification unlikely.

4.4.1.3 CSM interaction

In chapter 3, I estimated an intrinsic rate, through simulations of the detection efficiency for three out of 3628 SNe Ia with potential late-time CSM signals, of $< 0.5\%$ of normal SNe Ia displaying late-time (> 100 days after the peak) CSM signatures. [Dubay et al. \(2022\)](#) estimated a rate of late-time interaction in $< 5\%$ of SNe Ia. In this work, I identified three SNe Ia whose signals in ZTF cannot be ruled out as being due to late-time CSM interaction, giving a raw rate of late-onset CSM interaction of $\leq 0.1\%$ for the 4991 SNe Ia in this sample (see Table 4.6). No potential late-time signature from CSM interaction was identified in any other SN type.

SN 2017frh is my best candidate for late-onset CSM interaction. It has the earliest period of late-time detections (~ 250 days after the first detection) and a similar estimate for M_{CSM} as [Graham et al. \(2019b\)](#) found for SN 2015cp. If I assume a nova-like shell at $T = 10^7$ K, I get a reasonable CSM mass of $M_{\text{CSM}} \approx 4 M_{\odot}$ (although see the discussion of the limitations of my CSM estimates in section 4.4.2). My two other CSM candidates, SN 2017fby and SN 2016cob, have late-time signals five years after the SN was first detected. When I assume these objects to have shells of similar thickness as the shell in SN 2017frh, a CSM mass estimate of $\sim 5 M_{\odot}$ can be found for these objects as well. These estimates are lower limits, as no host extinction is assumed for objects in regions of galaxies where significant amounts of extinction can be expected.

I have not performed detailed rate calculations for the events discovered in this study. However, there are ~ 600 objects in my sample that were first detected less than 300 days before the start of ZTF, and I have one candidate with a late-time signal detected around 300 days after discovery. This gives a raw rate of 0.2% when not taking into account any bias. Two additional candidates have their detections in the fifth year after the explosion. About half of my sample exploded less than five years before the start of ZTF, and have their fifth year observed the survey. Two of these are candidates for having late-time CSM interaction at these phases, giving a rate of 0.05% when not taking any bias into account. My rough estimates of the rates suggest similar values to those of chapter 3 within the uncertainty of the small numbers of events detected.

The three events in chapter 3 that were identified as consistent with CSM interaction were also close to their host nuclei. As discussed in chapter 3 this could suggest a preference for CSM interaction to occur in SNe Ia in these environments or an alternative explanation is that they are caused by nuclear activity. However, as discussed in section 4.4.1.2 and shown in figure 4.9, these late-time signals are similar but not entirely consistent with nuclear activity/transients.

4.4.2 Limitations

I found a large group of SNe whose tail was detected by ZTF but never properly picked up by the survey – all of them required more thorough baseline corrections as the transient was present in the reference images. Several of these were exceptionally long-lived, with detections in the binned data being recovered over a year after the start of ZTF. While some of these can be explained by CSM interaction slowing the decline of the light curve, the gaps between the pre-ZTF and ZTF detections can make it difficult to say whether the late-time signal is due to new interaction or if the transient has been steadily fading between the observations. Some objects were already known to be interacting before the start of ZTF but I still grouped them with the tail detections as their interaction did not start at late times. I also found some periodic variable sources that have this behaviour for at least the length of the survey, making it impossible to generate proper reference images or apply a good baseline correction. Both groups can leave behind a ghost, which is an imprint at the transient location when the object is fainter than it was in the references. These events can be recovered from the data, particularly if different time frames for estimating the baseline correction are used, but they are often hidden if a baseline correction is not applied.

As in chapter 3, the magnitude limit of the binning technique is set by the magnitude limit of the references, as their uncertainty starts to dominate when more epochs of observation are binned together. One way to improve the sensitivity, though computationally expensive, is to find a region where I am certain there is no excess at the SN location and use it to generate custom references to use for difference imaging. This would immediately lessen the need for a baseline correction significantly, as an adequately chosen reference region already has this built in automatically.

Any search for late-time signals from SNe will also be dependent on the constraints placed on the sample. I required a first detection between 2008 and 2018 and a classification or candidate SN status. My sample contained several false positives from objects that had spurious detections within the initial sample definition, but a SN (which was

classified) exploded years later and was observed by ZTF. Siblings, whether they are true SN siblings or false due to spurious detections or projection effects, can be confused if they occur at (nearly) the same position in the sky. One way to remove false siblings such as AT 2017keg and SN 2013ld is to require a spectroscopic classification in the same time period as the first detection, though this would also remove objects like AT 2017gcd as in that case only the second sibling was classified, and candidate objects like AT 2017fwf and AT 2017gpv that were never spectroscopically classified.

I attempted to estimate the CSM mass required to produce late-time detections as bright as I observed them, assuming Bremsstrahlung to be the main mechanism producing light. In most cases when assuming a thin nova-like shell the required CSM mass is above $\sim 10 M_{\odot}$, with the only exception being SN 2017frh due to its relatively short delay time between the SN explosion and late-time detection. For SN 2016cob and SN 2017fby, the only way to get realistic CSM mass estimates is by assuming a colder CSM and a lower F_R , which could be explained by the shell size that is needed for interaction to start over five years after the SN (see the end of Sect. 4.3.4.1). [Graham et al. \(2019b\)](#) argue that the simplified assumption of Bremsstrahlung being the main mechanism is largely responsible for the high CSM mass estimates. They find an upper limit of $M_{\text{CSM}} \sim 7 M_{\odot}$ for SN 2015cp with this method, while [Harris et al. \(2018\)](#) find $M_{\text{CSM}} \lesssim 0.5 M_{\odot}$ using radio non-detections. They argue that line emission from elements such as H, Mg, and Ca are the main emission contributors instead. However, as line emission requires much more assumptions on the state of the emitting gas, it is much more complicated to derive a good estimation for M_{CSM} for it.

Even if I assume that Eq. 4.4 gives an upper limit and the better estimates from modeling of radio observations to be ~ 14 times smaller (based on the two estimates for SN 2015cp), the CSM interaction scenario for late-time signatures at 2 000 – 4 000 days after the initial SN discovery is still unlikely for five of my events that are also potentially due to nuclear activity. I have also not considered the impact of extinction in the host galaxy, which is likely non-negligible given the nuclear location of the transients considered. The CSM mass can be reduced by assuming a smaller F_R , but

the effect is limited unless I consider shells that are unphysically thin. However, thin shells are not ruled out by a several hundred-day long period of late-time detections, as the time delay due to the large radius of the shell can smear even a 1 s interaction signal out over several months. This shows that it is very difficult to properly constrain the CSM mass from photometry alone without knowing the mechanism powering it.

4.5 Conclusions

In this work I have presented a search for late-time signals observed in ZTF from transients that were first detected up to ten years before the start of the survey. By binning the ZTF data I was able to go nearly one magnitude below the noise limit of unbinned data points. My sample consists of 7718 unique objects that are in the ZTF footprint and could be searched for late-time signatures. By careful inspection of the ZTF images and comparison to known classes of transients, I determined the most likely source of the late-time signatures I identified at the positions of 98 transients discovered prior to the start of ZTF. My main conclusions are:

1. I identified several very long-lived transients that began pre-ZTF and continued for many hundreds to thousands of days. These include the SLSN II_n, SN 2015da, with data extending to eight years after discovery, and the spectroscopically unclassified AT 2017gpv, whose bumpy long-lived light curve (detected up to nearly four years after the discovery) is reminiscent of iPTF14hls ([Arcavi et al., 2017](#); [Sollerman et al., 2019](#)) and SN 2020faa ([Yang et al., 2021](#); [Salmaso et al., 2023](#)).
2. I found ten confirmed pairs of sibling transients where the pre-ZTF and ZTF siblings are at nearly the same sky position. One additional unconfirmed but likely sibling transient was identified when applying my binning technique at the positions of SN 2017ige (it was detected by Pan-STARRS with the internal name PS21aos). In most cases, one of the siblings was never spectroscopically classified. I also found two cases where the ZTF transient was mistakenly reported as a pre-ZTF transient due to spurious detections from the nearby host.

3. I found six flux excesses at the position of pre-ZTF events classified as Type II SNe (three Type II, one IIn and two SLSN-II) consistent with AGN activity in known AGN host galaxies. I speculate that some of these were likely misclassified at the time of the pre-ZTF detection.
4. I found five flux excesses that were consistent with activity of the host nucleus close to the location of the original transient. The brightness, light curve shape, and duration are broadly consistent with (ambiguous) nuclear transients, showing that there may be a previously unknown population of faint nuclear transients that requires deeper (binned) observations to be discovered.
5. I found three flux excesses that were consistent with late-time interaction with a $4 - 5 M_{\odot}$ CSM shell, assuming Bremsstrahlung as the main emission mechanism. In two of these, the signal is ~ 5 years after the SN and could also be explained by a nuclear transient unrelated to the SN. My best candidate is SN 2017frh with a signal ~ 250 days after the SN, consistent with a $4 M_{\odot}$ nova-like CSM shell at 10^7 K. However, a nuclear transient cannot be completely ruled out.

The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST; [Ivezic et al., 2019](#)) will observe the transient universe several magnitudes deeper than ZTF. This is crucial for finding faint interaction signals like those I have searched for in this chapter, as I have shown these to be rare. Increasing the volume that can be probed is the best way to find new events and follow them up. Current samples of objects will be observed at late times by LSST, making it possible to do a late-time signal search in real time. By using a well-defined and complete sample such as the ZTF SN Ia DR2 ([Rigault et al., 2025a](#), , [Smith et al., in prep.](#)), Bright Transient Survey (BTS; [Fremming et al., 2020](#); [Perley et al., 2020](#)), or a similar sample for ZTF SNe first detected in 2021+ as the basis for a similar search of late-time signals observed with LSST, it will also be possible to properly estimate their rates. Not only can these methods be used to search for late-time CSM interaction, but as I have shown, they are also suitable for finding other weak signals, such as faint nuclear transients.

A real-time search for Type Ia SNe with late-time CSM interaction in ZTF

In chapters 3 and 4 I searched for late-time signals in old transients, with most results being found after the period of activity. Follow-up was only attempted for SN 2019ldf (photometric) and SN 2020alm (spectroscopic) in section 3.3.2, but these provided no further signs of late-time CSM interaction. In order to understand more about the late-time signals that are recovered by binning the ZTF light curves, it is crucial to follow promising objects up using larger telescopes.

For this reason, I automate my pipeline to search a sample of SNe Ia observed by ZTF for currently active late-time signals using the latest observations. By doing this regularly, it is possible to find promising objects early and follow them up. My position as a student support astronomer at the Isaac Newton Group of Telescopes (ING) and Nordic Optical Telescope (NOT) proved to be extremely advantageous in obtaining time to observe targets using a 2.5 m telescope, even though this time has very low priority. It proved to be crucial for quick follow-up of several objects that are presented in this chapter, especially SN 2020qzx (see section 5.3.3).

This chapter is based on Terwel et al., (in prep.), which is currently being written and will be submitted to A&A. The sample selection, data reduction, analysis and interpretation were performed by me. Where relevant, I have noted the specific parts that were carried out by collaborators. In section 5.1 I present the sample used in the real-time search and explain where the latest observations come from. Section 5.2 describes the automated pipeline, details the telescopes used in the follow-up campaign, and gives a short overview of the reduction. In section 5.3 I describe the results for the promising objects that were found during the run of the real-time search. These are discussed in section 5.4, and I conclude in section 5.5.

5.1 Data

To create a SN Ia sample to search for late-time CSM interaction in real time, I collect all spectroscopically classified SNe Ia in ZTF using the Fritz broker ([van der Walt et al., 2019](#); [Coughlin et al., 2023](#)) and selected those objects that were first detected before 8 July 2023 (mjd 60133). Using this date cutoff ensures that all targets are over 100 days old at the start of the real-time search, and any new observations can be considered to be late-time observations of the SNe. This gives me a sample of 6 914 SNe Ia. Using the FPBOT package ([Reusch, 2023](#))¹ I construct the most up-to-date ZTF light curves for every object in my sample. ZTF uses difference imaging to remove constant sources from the science images and isolate the transient. Then, FPBOT forces a PSF fit at a given location in the difference images to construct a light curve at that position spanning the entire survey.

As this sample contains all ZTF SNe Ia up until the first half of 2023, this naturally includes the entire ZTF SN Ia DR2 ([Rigault et al., 2025a](#), Smith et al., in prep.), which have previously been examined for late-time signals in chapter 3. 69 objects that were not available on Fritz when I made my sample have been excluded. However, as the latest data point used in chapter 3 is on 26 June 2022 (mjd 59756) I probe their SNe at later phases, and they could (have) become active after the time period examined in chapter 3. Therefore, checking the same sample again could identify new late-time detections.

I aim to find late-time excesses while they are active, so the light curves have to be regularly updated with the latest observations. The ZTF survey is split into three parts: 40% is a private partnership survey whose data I can access immediately after the observation and becomes publicly available after 18 months. 40% is a public survey whose data is released every three months after a three-month proprietary, resulting in a three to six-month delay. The last 20% is private Caltech time, whose data is released publicly after 18 months.

¹<https://github.com/simeonreusch/fpbot>

I use the binning method from chapter 3 to search for faint signals, binning observations over 100 days after the SN peak together in bins of 25, 50, 75, and 100 days. (see section 5.2). The light curves do not have to be updated daily, as it will take several observations to push the detection limit deep enough to detect a faint signal in the latest bin. As the smallest bin size in the pipeline is 25 days, I aim to be able to create one additional bin each time an object is updated. I therefore split the sample into 28 lists and update the objects in one of these each day. This results in every object in the sample being checked and updated once every four weeks.

5.2 Analysis

Each time the light curves of the SNe Ia in my sample are queried and updated, they are run through the analysis pipeline from chapter 3. This pipeline takes all observations at a given sky position and after pre-processing (including applying a baseline correction to correct for systematic offsets across the light curve, see e.g. Yao et al. 2019; Miller et al. 2020) and finding the SN peak, it will bin all observations in each band that are over 100 days after the SN peak in bins that have a size of 25, 50, 75, or 100 days. By doing so, the noise in the data is reduced, allowing to probe objects that are up to nearly a magnitude below the noise limit of single observations ($m \approx 20.5$ mag) at the price of a reduction in the time sensitivity. This method is able to push the detection limit until the magnitude limit of the used reference images become the dominant source of uncertainty, which is around $\bar{m} \approx 21.8$ mag for ZTF (Strotjohann et al., 2021).

When the pipeline detects a possible late-time signal, whether currently or previously active, the object is rechecked manually to determine the possible origin of the late-time time signal and if it should be followed up. Possible explanations for the signal include, e.g., false positives due to a baseline offset, contamination from a possibly active host nucleus, known transients that are detectable over 100 days after the SN peak, or sibling transients at nearly the same sky position (see sections 3.3 and

4.3 for more details on these groups of objects). If the late-time signal is deemed real but has clearly faded again before it was found, all I can do is an analysis akin to that of chapters 3 and 4. However, if the signal is still active, I attempt to follow it up as quickly as possible using bigger telescopes to gain deep photometry (to confirm the late-time signal in a single observation) and spectroscopy (to search for spectral signature of the signal found in the binned ZTF observations. After all pre-existing light curves were generated, the real-time monitoring program started on 13 November 2023 (mjd 60261) and ran until 9 July 2024 (mjd 60500).

5.2.1 Follow-up observations and reduction

Follow-up observations of promising late-time detections were carried out with the Nordic Optical Telescope (NOT) and the Gran Telescopio CANARIAS (GTC) at Observatorio Roque de Los Muchachos, La Palma. I obtained the NOT data through staff time, except for the SN 2020qxx host spectra which I obtained through fast-track proposal p69-401 of which I am the PI. The GTC data was obtained and reduced by Dr. Lluís Galbany. General information on these telescopes is given in sections 2.2.2 and 2.2.3, respectively. An overview of the reduction process is given in section 2.4. The NOT photometry is reduced using a custom pipeline in PYTHON using the ASTROPY package. The GTC photometry is reduced with IRAF. After the basic photometry reduction, difference imaging is done using AUTOPHOT (Brennan & Fraser, 2022)¹ by Dr. Seán Brennan, using observations from Pan-STARRS (Chambers et al., 2016) or the Sloan Digital Sky Survey (SDSS, York et al., 2000; Adelman-McCarthy et al., 2006; Gunn et al., 2006) as reference images.

GTC+OSIRIS+ spectra were obtained at the location of SN 2019zbq at late times using the R1000R and a 1.0'' slit. They were reduced using a custom pipeline based on PYPEIT (Prochaska et al., 2020; Prochaska et al., 2020).

The NOT+ALFOSC spectra obtained at the position of SN 2020qxx were obtained with grism 4 with a 1.3'' slit on the 20, 30, and 31 March 2024 and 13, 14, and 15 April

¹<https://github.com/Astro-Sean/autophot>

2024. Spectra of the host galaxy of SN 2020qxx were obtained with the NOT+ALFOSC on 3 May and 19 June 2024 using a $1.0''$ slit. For these spectra I used the WG345 order-blocking filter to remove second-order contamination of blue light in the red part of the spectrum. This was not used in the spectra of the transient itself as it was already very faint, and removing an extra $\sim 10\%$ of the light that hits the detector would decrease the SNR of the spectra. These spectra were reduced using a pipeline made by Dr. Steve Schulze based on PYPEIT¹, taking care to separately extract the different sources lying close together.

Since a signal due to CSM interaction is not expected to evolve on the timescale of a few days (e.g., [Graham et al., 2019b](#)), I combined the spectra obtained in March 2024 to gain a combined exposure time of 3 hours on the target. I also combined the spectra taken in April 2024 together to obtain a second spectrum with 2 hours and 40 minutes on the target. These will be called the ‘March’ and ‘April’ spectra of SN 2020qxx below (see section [5.3.3](#)).

5.3 Results

The aim of this analysis is to identify in real-time SNe Ia with potential signatures of late-time interaction. My pipeline identified 349 events that had a late-time detection, including ones that have been analysed in chapter [3](#). 236 of these were ruled out as due to data issues (baseline offsets, etc.) and 101 due to known transients (tails, siblings, known Ia-CSM), leaving 12 objects.

Compared to chapter [3](#) and chapter [4](#), the pipeline returns a slightly larger fraction of objects from my sample. In both chapter [3](#) and my sample in this chapter, around two-thirds of the returned objects are false positives and around one-third are known objects, leaving the final 3-5% to be investigated further. In chapter [4](#), a little over 75% of the returned objects are false positives, and around 22% are known objects, leaving a comparable 2% of objects for deeper investigation. The main reason for this

¹https://gitlab.com/steveschulze/pypeit_alfosc_env

Table 5.1: Objects with a detected late-time signal in the binned ZTF light curves. The first two columns give the ZTF and IAU names of the object, followed by the position, classified type, and redshift. The classifications are taken from the ZTF DR2 release (Burgaz et al., 2025a; Dimitriadis et al., 2025), except for SN 2022uej, whose classification comes from TNS. Column 7 gives the peak date of the SN, and column 8 gives the discovery date of the late-time signal. Column 9 gives the start of the first bin that is part of the late-time signal with the bin size that best shows the signal. Column 10 gives the duration of the late-time signal in rest frame days. The last column is a comment on the nature of the late-time signal. The table is divided into two parts, with the objects that were not followed up above and objects that were followed up below the division line.

ZTF name	IAU name	RA	Dec	Type	z	MJD _{peak-SN}	MJD _{disc-lt}	MJD _{start-lt}	Δt (day)	Comment
ZTF18abtgys	SN 2018gtr	00:04:36.30	51:59:37.63	Ia-norm	0.042 ± 0.003	58372	60279	59712	162	chapter 3
ZTF19ablekwo	SN 2019mse	00:15:21.36	46:44:08.64	Ia-norm	0.088 ± 0.004	58715	60283	59167	1170 ⁽²⁾	chapter 3
ZTF20abjftv	SN 2020tfc	22:17:00.81	30:39:20.97	Ia-norm	0.031 ± 0.001	59116	60269	59671	804 ⁽³⁾	chapter 3
ZTF20aafytx	SN 2020alm	16:52:01.49	23:32:22.77	Ia-norm	0.06001 ± 0.00001	58873	60271	59624	826 ⁽³⁾	chapter 3
ZTF22abgbbez	SN 2022uej	07:19:57.72	46:00:59.00	Ia	0.03272 ± 0.00008	59839	60454	60168	212	r detected only
ZTF19acgonvy	SN 2019spg	08:21:48.09	-14:04:12.00	Ia	0.073 ± 0.001	58790	60353	58891	110	Nuclear
ZTF19abvypoh	SN 2019pqn	16:04:03.86	18:28:18.46	Ia-norm	0.03720 ± 0.00001	58745	60458	60436	62 ⁽³⁾	Nuclear
ZTF20acpbhbf	SN 2020yvs	10:24:22.78	41:42:25.97	Ia-norm	0.04453 ± 0.00001	59160	60268	59260	1127	False positive
ZTF20aazhzna	SN 2020jsa	14:24:23.94	26:41:23.00	Ia-norm	0.03657 ± 0.00001	58992	60431	60315	72	Nuclear
ZTF18acauscw	SN 2021nbt	22:42:29.54	01:49:53.84	Ia	0.07 ± 0.01	59359	60283	59718	730 ⁽³⁾	Nuclear / ANT
ZTF19adedgea	SN 2019zbq	09:08:53.91	55:19:08.86	Ia	0.04678 ± 0.00001	58859	60261	59958	518 ⁽³⁾	Nuclear / ANT
ZTF20abqkbfx	SN 2020qgx	18:04:00.24	74:00:50.07	Ia-CSM	0.0964 ± 0.0004 ⁽¹⁾	59094	60388	60332	46	Late-time line emission

- ⁽¹⁾ z value as derived from the classification spectrum, see section 5.3.3.2 for more details.
⁽²⁾ Two periods of detections were identified.
⁽³⁾ Active at the end of the monitoring period.

Table 5.2: Follow-up observations of transients with late-time signals. The exposure time is noted as the number of exposures x exposure time of a single exposure. Photometric observations taken in multiple bands with the same exposures are mentioned in the same row.

Name	Tel. / Inst.	Mode	Setup	MJD	Exp. time (s)	Seeing (")	Comment
SN 2020yvs	NOT / ALFOSC	phot	<i>gr<i>i</i></i>	60388	10 x 10	0.9	$g > 22.2$, $r > 21.8$, $i > 21.6$ mag
SN 2020jsa	NOT / ALFOSC	phot	<i>gr<i>i</i></i>	60432	10 x 30	1.0	$g > 22.76$, $r > 22.49$, $i > 23.05$ mag, host residual
SN 2021nbt	NOT / ALFOSC	phot	<i>ri</i>	60482	4 x 30	0.9	$r = 19.458 \pm 0.040$, $i = 19.123 \pm 0.062$ mag
SN 2019zbq	GTC / OSIRIS+	phot	<i>r</i>	60271	1 x 10	1.2	$r = 20.7 \pm 0.1$ mag
		spec	Grism R1000R, 1.0" slit	60282	3 x 1200	1.2	Contains host & excess spectra
		phot	<i>i</i>	60389	1 x 600	1.7	$i > 21.5$ mag
SN 2020qxz	NOT / ALFOSC	spec	Grism #4, 1.3" slit	60389	3 x 1200	1.7	Part of March spectrum
		phot	<i>i</i>	60399	1 x 600	1.4	$i > 21.5$ mag
		spec	Grism #4, 1.3" slit	60399	3 x 1200	1.4	Part of March spectrum
		spec	Grism #4, 1.3" slit	60400	3 x 1200	0.8	Part of March spectrum
		spec	Grism #4, 1.3" slit	60413	3 x 1200	1.5	Part of April spectrum
		spec	Grism #4, 1.3" slit	60414	3 x 1200	1.7	Part of April spectrum
		spec	Grism #4, 1.3" slit	60415	2 x 1200	1.3	Part of April spectrum
		spec	Grism #4, 1.0" slit (*)	60433	3 x 1200	1.4	Part of host spectrum
		spec	Grism #4, 1.0" slit (*)	60480	3 x 1200	0.7	Part of host spectrum

(*) A WG345 order blocking filter was used to remove second order contamination in the red part of the spectrum.

is because the sample in chapter 4 consists of objects first detected before the start of ZTF, leaving less opportunity to recover the declining tails of bright SNe, which is the group of known objects returned by the pipeline.

Table 5.1 shows the 12 objects whose signal was detected in the binned ZTF light curves at ≥ 100 days after the SN peak. Several of the objects listed have been found in chapter 3 as their sample is a subset of mine. Even though their sample has been checked before, some objects might have become active after their search, but a visual inspection showed that this is not the case for the recovered objects. As their previous detections are already described in detail in chapter 3, I will not consider them further here. Table 5.2 shows the follow-up observations for those objects that were found while the signal was still present.

First, I will discuss the three new objects that have been found and were not previously identified but I could not follow up on (section 5.3.1). After this I will discuss the five late-time detections that were followed up on and the interpretation of their data (section 5.3.2).

5.3.1 Faded objects

I found three new objects (SN 2022uej, SN 2020spg, SN 2019pqn) where a late-time detection was identified but I could not follow up on. In two cases, the late-time signal was recovered too late for follow-up observations, and in one case, I had no access to suitable telescope time to obtain constraining follow-up observations. The light curves of these objects are shown in figure 5.1.

SN 2022uej. A late-time r -band signal in SN 2022uej was first picked up by my pipeline on 24 May 2024. It appears to have started 276 rest frame days before this and lasted over 200 days. There were observations in the g - and i -bands at similar epochs, but no detections at $> 5\sigma$ were found. The date on which the r -band detection was identified is a few days after a ZTF data release, which might have included some images from the public survey needed for a detectable signal. SNAP shows that the




Figure 5.1: **Left:** Binned light curves in flux space in the SN rest frame of the three objects with a late time signal that I could not follow up on. The three ZTF bands *g*, *r*, and *i* are shown in green, orange, and red, respectively. Individual observations are shown before the binning starts. Bins are shown as coloured blocks showing the size of the bin, mean value, and 1σ uncertainty of the binned observations, with the point showing the mean observation date of the bin. Bins that are $> 5\sigma$ above zero flux are filled, and bins $\leq 5\sigma$ from zero flux are dashed. **Right:** Same plot as on the left but in absolute magnitude, corrected for MW extinction. 5σ detections in individual observations are shown before the binning starts, and the 5σ binned detections are shown in the same way as in the left plots. 5σ upper limits are shown with a downward arrow for the binned non-detections. The grey vertical line shows when the excess was first discovered.

object is close to the imperfectly subtracted host nucleus, which causes some extra noise in the forced photometry light curve. However, the resulting dipole at the host location does not overlap with the SN location as it is oriented perpendicular to the direction of the SN, while there is a slight extra excess at the SN location $2.6''$ from the host. The identified signal is real in the data, but I was unable to follow up on the object because it had already faded by the time it was discovered.

SN 2019spg. The location of SN 2019spg has a relatively short, sparse baseline before the SN explosion. Immediately after the SN fades, the flux seems to settle at a plateau significantly above the pre-SN baseline. Due to it staying at a nearly constant flux level for much longer than the baseline, this could suggest that the baseline

correction was incorrect and the actual baseline should be at the plateau level. With a separation of $0.5''$, this object is also very close to the host galaxy nucleus. According to the AGN criterion presented in [Hviding et al. \(2022\)](#), which uses data gathered with the *Wide-field Infrared Survey Explorer* (WISE, [Wright et al., 2010](#)), the host galaxy is not an AGN. This does, however, not mean that the host cannot show small variability, as has been shown in chapter 3 and chapter 4. Because of these reasons, I decided not to follow up on this late-time signal. After around 1 000 days, the flux level dropped again, nearly to the pre-SN baseline level. This drop shows that the previous plateau was a real excess after all, although it was likely related to variability of the host nucleus, not the SN. Due to a combination of unfortunate timing of the excess and unfortunate sampling of the light curve, specifically the pre-SN light curve, these late-time detections of likely nuclear activity looked like a false positive until the signal faded and the original baseline was recovered.

Interestingly, a baseline correction was needed to recover the signal. SNAP shows no excess in the difference images, but it does show a ghost in the pre-SN region and after the binned signal has faded, which suggests that the excess is in the reference images and that the host dimmed between the time when the reference images were made and the SN explosion.

SN 2019pqn. A late-time signal at the position of SN 2019pqn was identified on 28 May 2024, only showing faint g -band detections that had started 21 days earlier in the rest frame of the SN. This lasted at least 64 days and was active at the end of the monitoring period. Only upper limits in r - and i -bands were found at similar epochs. SNAP shows that the late-time signal lies between the SN and host nucleus locations (which are $1.9''$ apart) but seems to slightly favour the host nucleus location. When this excess was found, I did not have the opportunity to follow it up as I only had very low priority time at the NOT. Therefore, I only have the short period of binned photometry detections. Given that it was only detected in the g -band, combined with the excess being slightly offset from the SN location, a nuclear transient is a good explanation

of this late-time signal, even though the host nucleus is not an AGN according to its WISE colours.




Figure 5.2: Same as figure 5.1, but for the objects for which follow-up observations were made. The grey vertical line shows when the excess was first discovered, and the grey dashed and dotted lines show the photometric and spectroscopic follow-up, respectively. The coloured lines points at the follow-up dates show the follow-up detections and the crosses show the follow-up 5σ upper limits. Note that for SN 2020qxz the time period between the SN tail and late-time signal has been cut out to better show the detections and follow-up campaign.

5.3.2 Follow-up results

I found five objects with an active late-time signal for which I could obtain follow-up observations. SN 2019zbq was followed up with the GTC, while the others were followed up with the NOT. Below, I discuss each object individually, starting with the

three for which only photometry was taken and then SN 2019zbq, for which I obtained photometry and spectroscopy. SN 2020qxz, for which more extensive follow-up was obtained, is discussed separately in section 5.3.3. The light curves of these objects are shown in figure 5.2. The details of their follow-up observations are presented in Table 5.2.

SN 2020yvs. On 20 November 2023, SN 2020yvs was found to have elevated flux levels in all three ZTF bands that had started nearly 1 000 d before, with the effect being stronger in the redder bands. The detected excess is clearly visible with SNAP, though at times it is merged with a residual dipole from the host. As the SN is very close to the bright nucleus of the host galaxy ($1.3''$, host SDSS r mag = 13.56), I decided to keep a close eye on the development of the light curve to determine if it was more likely to be associated with the SN or not.

As the late-time detections grew stronger over time, I obtained confirmation through photometry with NOT/ALFOSC in the *gri*-bands. These observations were obtained on 19 March 2024, with ten exposures of 10 s each using the SDSS *gri* filters that were stacked together for each band. The short exposures are needed to avoid saturating the host galaxy as image subtraction has to be performed to search for an ~ 20 mag excess in an ~ 13.5 mag environment. The image subtraction was done through AUTOPHOT (Brennan & Fraser, 2022) using SDSS images as templates and was attempted using Pan-STARRS images as templates, though these proved to have a saturated host galaxy, making them unusable as templates.

After image subtraction, a small residual at the host nucleus was left, but no excess was found at the SN location. The host residual is the brightest source close to the location, having at most a 2σ detection at 21.6 mag in the *i*-band. This is significantly below the detections found in the late-time ZTF light curve (brighter than mag 20 in both *r* and *i*). Therefore, I conclude that the binned detections are not real but likely due to artefacts in the ZTF images after subtracting the bright host galaxy.

SN 2020jsa. The late-time signal in SN 2020jsa was found on 1 May 2024, showing a 21.2 ± 0.1 mag r -band excess that had begun 112 days earlier and was already fading at the time of discovery. The SN is located $3.0''$ from the host galaxy nucleus, and a previous faint detection in the binned late-time light curve can be seen around a year after the SN peak. While the host nucleus is relatively far from the SN, the difference images show a dipole with the positive side towards the SN after host subtraction. This would suggest that there are subtraction issues or that the host shows slight variability. However, according to the WISE criterion, the host galaxy is not an AGN. The object was observed with the NOT in gri the next night, using short exposures to avoid saturation of the bright host and nearby stars. The resulting difference images found no excess at the SN location, providing a deeper limit than the binned ZTF r -band light curve. However, a faint residual was found at the host nucleus location, showing that the excess is real but unlikely to be related to the SN.

SN 2021nbt. The late-time signal in SN 2021nbt was first detected on 5 December 2023. The object is only $0.6''$ from the host nucleus, and its ZTF name is from 2018, which suggests that some host variability caused a first detection long before the SN exploded. This immediately raised suspicion that the detected late-time signal could be related to the host instead. Because of this, it took until 21 June 2024 before I decided to follow it up with the NOT. The signal started relatively shortly after the SN and persisted for several months, which could suggest it is connected to the SN in some way.

After difference imaging was performed on the ri -band photometry obtained with the NOT (see Table 5.2), a residual was recovered at $r = 19.458 \pm 0.040$ mag and $i = 19.123 \pm 0.062$ mag. This is somewhat brighter than the binned detections using ZTF observations taken around the same time, and a signal this bright would have been detectable in the individual ZTF observations. If the host nucleus varies slightly over time, however, this could cause different reference images to show the host at different brightness, which would result in different residuals after difference imaging.

Unfortunately, there is no spectrum of the host galaxy, and since the signal is superimposed on the much brighter host, it is impossible to extract a spectrum of the excess without it being completely dominated by the host. The host is too small to attempt extracting a pure host spectrum away from the excess location, and without a host spectrum, the excess contribution cannot be isolated.

At the distance of the host, the magnitude of the excess observed by the NOT translates to an absolute magnitude of $M \sim -18$ mag. As can be seen in figure 5.2, the excess appears after a gap in the observations and persists for at least 780 days at a fairly constant magnitude. This is at the bright end of the potential ambiguous nuclear transient (Hinkle et al., 2022) population identified in chapter 4, but still consistent with the brightness of the comparison object they use. Due to its shape, brightness, and duration, an ANT that started somewhere between 200 and 400 days after the SN is a good explanation for this excess. However, I do not have spectroscopy to confirm since it was so faint relative to the host.

SN 2019zbq. On 13 November 2023, SN 2019zbq was noted to have an active late-time signal in the binned r -band at $m = 21.5 \pm 0.2$ mag, with the signal already being present for nearly 300 days before and being recovered regardless of bin size or placement. The binned g -band did not rise from the baseline during this time, but the i -band did, though the larger uncertainties in this band prevented the detections to reach a 5σ confidence level. The SN is close to the centre of its host galaxy, with an offset of $1.09''$.

Checking the r -band difference images with SNAP ruled out image defects, reduction or subtraction errors, or a sibling transient with a small sky separation from SN 2019zbq (within the uncertainty of the ZTF resolution) as potential explanations for the detections. The host galaxy has $W1 - W2 = 0.02 \pm 0.04$ and $W2 - W3 = 0.9 \pm 0.5$, which is not an AGN according to the AGN criterion presented in Hviding et al. (2022).

On 23 November 2023 an r -band image was obtained at the GTC with OSIRIS+. Using AUTOPHOT and reference images from Pan-STARRS, a $m = 20.7 \pm 0.1$ mag




Figure 5.3: **Left:** r -band image of the location and host galaxy of SN 2019zbq, taken on 23 November 2023 with the GTC using OSIRIS+. The SN location is marked in red and the purple dot is the galaxy nucleus location. **Right:** Difference image of the left region after subtracting a Pan-STARRS template image. There is a residual visible at the host galaxy location. Note that the colour scaling is different for the two images to highlight the important sources.

source at the host nucleus location was obtained, which is clearly offset from the SN location, confirming the late-time excess is real but not related to the SN. The GTC image and the resulting difference image are shown in figure 5.3.

On 4 December 2023, a spectrum was obtained with the GTC and OSIRIS+. Since the transient is located on top of the host, this spectrum contains light coming from both the host and the transient. The host is 4.5 mag brighter and dominates the spectrum heavily, meaning the host must be subtracted to recover any potential transient spectrum. Since a slit was used in obtaining the spectrum and the host galaxy is an extended object, a pure host galaxy spectrum is extracted away from the transient location, and I subtract it from the spectrum at the brightest part of the trace containing the transient. Assuming that the spectral features of the host are similar enough at these two locations, the main difference between the two, apart from total flux observed, will be the excess I am interested in.

Figure 5.4 shows the extracted host and excess spectra. Since the red edge of the detector has very low sensitivity, this can be used to scale the pure host spectrum to the combined host and excess spectrum, allowing me to take the difference and




Figure 5.4: GTC spectrum of the transient in the SN 2019zbq host galaxy, taken on 4 December 2019. Vertical dotted lines give the locations of sky lines and tellurics. **Top:** Extracted spectra at the centre of the trace (host + excess) and at the side of the extended host galaxy trace (host). **Bottom:** The difference of the two spectra in the top panel, showing the isolated spectrum of the excess. The host spectrum has been re-scaled to ensure that the red edge is the same for both spectra.

extract the pure excess spectrum. The pure excess spectrum is shown in the bottom panel. Integrating this spectrum over the r -band shows this excess to be ~ 21 mag, which agrees with the photometry values. A few small spikes can be seen in the excess spectrum, but these do not match up with any known emission lines at the host redshift, even when assuming a consistent velocity offset of the excess from the host. Overall, the spectrum looks red, which matches the binned ZTF data, as the g -band does not rise from zero flux (and the i -band rises the most but with a large scatter).

The excess has an absolute magnitude $M = -16.0$ mag. This, together with the nuclear location and the sudden appearance of the transient followed by a relative stable plateau for over 500 days, is consistent with the faint ANT-like nuclear transients found in chapter 4. Therefore, I conclude that a faint ANT-like transient provides a good explanation for the signal, although I caution that extracting the transient spectrum and removing the host contribution is not straightforward.




Figure 5.5: Pan-STARRS DR1 colour image of the location and host galaxy of SN 2020qzx. The SN location is marked in red, and the two host galaxy lobes are marked in white.

5.3.3 SN 2020qzx

SN 2020qzx was classified as a SN Ia-CSM at early times, with a prominent $H\alpha$ line visible in the classification spectrum taken around the SN peak which was used to obtain a redshift of $z = 0.0964$ for the object. The slowly decaying tail is visible in the ZTF observations up to 1.2 years after the peak and an additional three months in the binned light curve before fading below the noise limit, resulting in its recovery in chapter 3.

On 20 March 2024 (mjd 60389), at nearly 1180 d after the SN peak in the rest frame, this object was flagged as currently active due to an elevated flux level in the i -band in the last four weeks, with the detection being at $m = 20.9 \pm 0.2$ mag (see

bottom row of figure 5.2). As can be seen in figure 5.5, the SN location is visually offset from the host galaxy, so I obtained follow-up observations with the NOT+ALFOSC on the same night. The resulting spectrum showed a prominent emission line around $\lambda = 9210 \text{ \AA}$ in the observer frame ($\lambda = 8400 \text{ \AA}$ in the SN rest frame). After this promising detection, I launched a campaign to take more observations (photometry and spectroscopy) with the NOT+ALFOSC over the following weeks.

5.3.3.1 Late-time photometry

Two epochs of *i*-band photometry were taken with the NOT+ALFOSC on 20 and 30 March 2024. Both were single 600 s exposures, which were reduced using the method described in section 5.2.1. No residual flux above 5σ was found in either image, resulting in $m = 21.5$ mag upper limits. As can be seen in figure 5.2, all observations taken with the NOT were after the detection in the binned ZTF data. However, likely transient flux was detected in the spectra taken at the same epochs as this follow-up photometry. The most likely explanation for the photometric non-detections at the same epochs is that the emission lines seen in the spectra are no longer strong enough to cause a detectable increase in the broadband flux.

5.3.3.2 Host spectroscopy

The host galaxy of SN 2020qxx has two lobes, a big one on the South-East side and a smaller one on the North-West side (see figure 5.5). By orienting the slit along the axis of the galaxy, I obtained spectra of both lobes on the nights of 3 May and 19 June 2024 with the NOT+ALFOSC, resulting in a two-hour spectrum on the big (South-East) lobe and a one hour and 40 minutes spectrum on the small (North-West) lobe as one of the 20-minute exposures on 3 May 2024 was contaminated at the location of the small lobe and could not be used.

The spectra of the big and small lobes shown at the bottom of figure 5.6 look nearly identical, and the galaxy emission and absorption lines discussed below are at the same position for both lobes. This shows that the host is more likely a single object at a

single z rather than two distinct galaxies overlapping in the line of sight.




Figure 5.6: Combined late-time spectra of SN 2020qzx and its host in the SN rest frame. The transient emission lines are marked with grey vertical lines, and the best-fit Gaussians are overlaid on top of the combined March spectrum. The dashed vertical lines show the location of host galaxy lines (as well as $H\beta$ to compare its location to the transient emission line) at the host redshift of $z = 0.0975$, and the dotted vertical lines show the location of sky lines. I also show the classification spectrum of SN 2020qzx taken around peak magnitude in light blue and the late-time spectrum obtained for SN 2015cp in magenta, scaled down by a factor of 2 for readability.

Several lines can be seen in the host spectra. I identify weak lines which I associate with $H\alpha$ emission, $[S II]$ emission, Ca II H&K, and NIR triplet absorption, and Na ID and Mg I absorption. These galaxy lines match the spectrum at $z = 0.0975$, which is slightly larger than was found for the $H\alpha$ emission seen in the maximum-light classi-

5. A REAL-TIME SEARCH FOR TYPE IA SNE WITH LATE-TIME CSM INTERACTION IN ZTF

fication spectrum of SN 2020qxx as Ia-CSM (top spectrum in figure 5.6). Assuming that the host galaxy is indeed at this larger z and that the SN is associated with the host, the blueshifted $H\alpha$ line in the classification spectrum of SN 2020qxx translates to the SN moving towards us with $\sim 200 \text{ km s}^{-1}$ with respect to the host galaxy. The different z value would also slightly change the distance modulus, resulting in all absolute magnitudes being increased by $\sim 0.03 \text{ mag}$. Unless specified otherwise, I use the SN rest frame at $z = 0.0964$ while further discussing this object.




Figure 5.7: Same as figure 5.6 but zoomed in on the regions containing the transient emission lines and the $H\alpha$ region. The 1σ uncertainties of the Gaussian fits are shown in grey.

5.3.3.3 Late-time spectroscopy at the SN location

I obtained spectra at several epochs at the position of SN 2020qxx in March and April 2024, which were combined into a ‘March’ spectrum and ‘April’ spectrum (see section 5.2.1 for further details). In the March spectrum, I identify four prominent emission lines and zero absorption lines that are not associated with the host galaxy. Both the March and April spectra show some signs of slight host contamination, but the four emission lines are only seen in the March data, not in the April spectrum nor the host spectra obtained in May and June 2024. This suggests that these features are transient in nature, and are decaying between the two periods of observations. These emission lines are seen at $4\,766.0 \pm 1.1$ Å, $8\,395.8 \pm 0.6$ Å, $8\,423.6 \pm 1.4$ Å, and $8\,615.8 \pm 1.3$ Å in the rest frame of the SN and are marked in figure 5.6 and figure 5.7. figure 5.8 and figure 5.9 show the individual spectra that were combined into the March and April spectra.

I have fitted these four emission lines with Gaussian profiles to measure their strengths and widths, the results of which are given in Table 5.3. The fits to the four lines are shown in figure 5.6 and figure 5.7. Three of the lines have a similar dispersion, suggesting that they are from the same source. However, the bluest line has a much higher dispersion that is incompatible with the other three.

Identification of the origin of these emission lines is not straightforward. SN 2020qxx is a known SN Ia-CSM that showed H α emission near peak brightness, so it is natural to assume that at least one of the four lines in the March 2024 spectrum is due to H emission. Indeed, the line at $4\,776$ Å can be interpreted as H β emission, though this requires it to be blueshifted by $5\,920 \pm 70$ km s⁻¹. Offset emission lines have been found before at up to $10\,000$ km s⁻¹, see e.g., [Chen et al. \(2020\)](#) and [Prentice et al. \(2020\)](#). However, the significant difference in the hydrogen velocity between the early and late-time detections suggests that this material has been accelerated in the time between detections or is coming from a different location.




Figure 5.8: Individual late-time spectra of SN 2020qxx taken with the NOT. The three individual epochs are shown in red, green, and blue, and the stacked spectrum is shown in black. Each individual spectrum has an exposure time of 3×20 minutes, except for the spectrum taken on 15-04-2024 which is 2×20 minutes. The transient emission lines are marked with grey vertical lines, the dashed vertical lines show the location of host galaxy lines (as well as $H\beta$ to compare its location to the transient emission line) at the host redshift of $z = 0.0975$, and the dotted vertical lines show the location of sky lines.

SN ejecta sweep up the CSM they are interacting with, so some acceleration can be expected. This has been seen before in some SNe IIn, where aspherical, possibly clumpy CSM has been accelerated to several thousand km s^{-1} resulting in wide boxy or multi-peaked emission lines (Leonard et al., 2000; Pozzo et al., 2004; Smith et al., 2015). The other issue is that where there is $H\beta$ emission, one would expect $H\alpha$ emission as well. However, there is no sign of any emission line at the position of $H\alpha$

in the March or April spectra. There is a small bump that is consistent with being due to host $H\alpha$ contamination, but nothing is seen with the same offset as $H\beta$.




Figure 5.9: Same as figure 5.8 but zoomed in on the regions containing the transient emission lines and the $H\alpha$ region.

The other three transient emission lines at 8395.8 Å, 8423.6 Å, and 8615.8 Å in the rest frame of the SN have their own identification challenges. First of all, these are in the same region as the host galaxy Ca II NIR absorption. For the transient emission lines to be visible, this dip in the background spectrum has to be overcome first. Three emission lines are found, but their relative positions are different from the Ca II NIR triplet, meaning that not all lines can be explained by them. If I however shift the

Table 5.3: Fits of the emission lines identified in the March 2024 spectrum at the position of SN 2020qzx. The first three columns give the central wavelength and width (in Å and km s⁻¹) of the fitted lines. The last three columns give the most likely matched elements and emission lines, assuming some amount of blueshift velocity offset (given in the last column). The redshift uncertainty contribution to the velocity offset (120 km s⁻¹) has not been included. All fits are done in the SN rest frame.

Central wavelength (Å)	σ (Å)	σ (km s ⁻¹)	Element	Rest wavelength (Å)	Velocity offset (km s ⁻¹)
4 766.0 ± 1.1	12.7 ± 1.4	800 ± 88	H β	4 861.00	5 920 ± 70
8 395.8 ± 0.6	9.3 ± 0.7	332 ± 25	Ca II ⁽¹⁾	8 542.00	5 180 ± 20
8 423.6 ± 1.4	5.6 ± 1.5	199 ± 53	NI	8 567.74	5 090 ± 50
8 615.8 ± 1.3	5.8 ± 1.5	202 ± 52	K I ⁽²⁾	8 763.96, 8 767.05	5 170 ± 50

⁽¹⁾ Only the middle line of the Ca II NIR triplet is visible, the other two lines are not present due to atmospheric and host contamination for the blue and red component, respectively.

⁽²⁾ K I is a doublet for which I use the mean to estimate the velocity offset.

middle line of this triplet to match with the bluest transient line, the blue line in the Ca II NIR triplet overlaps with a skyline while the red line is in the middle of the dip mentioned before. This could explain why only one of the three components is visible. The required blueshift is $5\,180 \pm 20$ km s⁻¹, significantly lower than the shift required for H β of $5\,920$ km s⁻¹. A shift of $\sim 5\,150$ km s⁻¹ allows the remaining two transient lines to be matched to a NI line and K I doublet with the same velocity, resulting in a possible identification of three offset emission lines with similar widths. For the emission line matching with the K I doublet I use the mean wavelength.

Neither of these elements are expected to have a strong presence in the ejecta of SNe Ia, nor have these lines been seen in other SNe with CSM. [Graham et al. \(2015b\)](#) did, however, find time-varying K I absorption lines in SN 2014J around its peak, which they associate with photoionization of distant CSM, although the inferred distance is larger than what the ejecta could travel in a couple of years. It might also be possible to obtain a nitrogen-enhanced CSM if the donor star is an evolved star massive enough to burn hydrogen through the CNO cycle but small enough for the primary to have ended up as a WD. Dredge-ups could bring fusion products to the surface, including nitrogen, which is then lost from the system to form the CSM. [Barmantloo et al. \(2024\)](#) make a similar argument for the production of [N II] lines in stripped envelope CCSNe.

The late-time CSM in SN 2015cp was primarily $H\alpha$ dominated (Graham et al., 2019b), which is absent in SN 2020qzx. At the same time, three of the four emission lines found in SN 2020qzx are absent in SN 2015cp. There is a tentative detection of the Ca II NIR triplet in SN 2015cp, but unlike in SN 2020qzx, all three components are observed. They only require a velocity offset of $\sim 1\,400\text{ km s}^{-1}$, though with a FWHM of $\sim 3\,000\text{ km s}^{-1}$ the Ca II lines in SN 2015cp are much broader than any of the transient emission lines found in SN 2020qzx.

5.4 Discussion

In this study I monitored a sample of 6 914 SNe Ia discovered by ZTF for signs of late-time rebrightening due to potential CSM interaction. I created forced photometry light curves at the location of each SN for the full duration of the survey and updated these every four weeks with the latest observations in an attempt to find late-time signals in real time and follow them up with further observations. I found 12 objects with late-time rebrightening that could not be explained by image defects, known contamination of host nuclei, or known sibling transients. Four of these were known from chapter 3 and discussed there. Their rediscovery here did not lead to any new insights. In three objects, the signal was found when follow-up was not possible. I was able to follow the signals in the other five objects up, using the NOT for four objects and the GTC for one. Out of these, there was one false positive, three likely nuclear transients unrelated to the SN, and a confirmed late-time signal in SN 2020qzx.

5.4.1 Detection challenges

As was shown in chapter 3 and chapter 4, late-time signals in SNe Ia are rare. Out of nearly 7 000 objects, I found eight objects with a late-time signal detected somewhere in the 240 days as I was actively monitoring them and an additional four objects whose late-time signals were already discussed in chapter 3. During this time many other objects were flagged but turned out to be false positives, bright long-lived SNe,

or sibling transients. SNe near their host nucleus are especially prone to give false detections.

Most objects for which a late-time signal was flagged in Table 5.1, chapter 3, and chapter 4 are close to the host nucleus. These regions are difficult to handle with difference imaging, often leading to residuals that can be mistaken for a weak signal by forced photometry. Even if the excess is real, the busy centres of galaxies leave plenty of opportunity for alternative explanations, including ones involving the central supermassive black hole such as AGN, tidal disruption events and other sources of nuclear variability, e.g. ANTs (Hinkle et al., 2022). One could make a cut on the host separation to ensure that contamination from nuclear events is removed, but I did not do this here as it would introduce a bias in the SN environments in my sample. Such a cut is also dependent on the pixel size of the used observations, which means that more distant SNe require larger host separations in order to not be removed by the cut. I decided to instead be as inclusive as possible with the knowledge of the contamination this leads to in my results.

Even when searching for faint late-time signals in real time, the fact that my method requires several observations to bin together in order to reach a deeper detection limit results in my method still being slightly behind current events. This makes it crucial to follow up quickly on newly found signals with bigger telescopes and better instruments that are able to measure such faint signals, especially as sometimes the late-time signals have been relatively short-lived and were already fading again by the time they were found. Without my ability to follow up quickly, the transient emission lines in SN 2020qxx would likely have been missed.

5.4.2 Follow-up campaigns

For five objects that showed late-time signals I was able to get follow-up photometry, on which difference imaging and forced photometry was performed using AUTOPHOT. In one case (close to the position of SN 2020yvs), the follow-up observations provided deeper non-detections than the late-time signal that was found in the binned ZTF

observations. The host galaxy of SN 2020yvs is very bright, which likely results in small artefacts in the difference images that are misidentified as the late-time signal. In three other cases (close to the positions of SNe 2020jsa, 2021nbt and 2019zbq), a residual was detected at the host nucleus location, showing that for these objects the recovered late-time signal was unrelated to the SN.

The host galaxy of SN 2019zbq was confirmed to have an active nuclear transient and was followed up with spectroscopy. However, since the host galaxy is 4.5 mag brighter than the transient, it dominates the spectrum heavily, and the only way to extract the transient spectrum is by subtracting the host contribution. This requires the existence of a pre-transient host spectrum (from e.g. SDSS, as was done for the late-time signal in SN 2020alm in chapter 3) or a host spectrum without contribution from the transient extracted from the same observation as the spectrum containing the transient, as was done for SN 2019zbq in section 5.3.2. This is a tricky process and can introduce a lot of uncertainty. For this reason, I did not attempt to get a spectrum of the confirmed excess in the host of SN 2021nbt, as there was already enough evidence of the late-time excess being unrelated to the SN. Integrating the flux of the SN 2019zbq host excess spectrum over the r -band matches the measured r -band brightness, showing that the excess extraction has worked.

The excess spectrum in the SN 2019zbq host looks quite red but has no distinctive emission or absorption features. This is similar to the ANT AT 2020ohl (Hinkle et al., 2022), although AT 2020ohl is much bluer and brighter. The ANTs presented in Wiseman et al. (2025) are much closer to the excess in colour but again much brighter, and they all show Balmer emission lines. The small peaks that do stand out in the excess spectrum do not match any set of elements at a single velocity offset. The best I can say about it is that this is indeed some sort of ambiguous nuclear transient.

5.4.3 The late-time signal in SN 2020qxz

SN 2020qxz stands out from the other objects in the sample for several reasons. First of all, it is a SN Ia-CSM with detected interaction around its peak. It is also visually

separated from its host nucleus, which makes it easier for the excess to be interpreted as being related to the SN instead of nuclear activity, and it is also easier to get a spectrum of the excess without much contribution from the host galaxy. I followed this excess up over three months and obtained nearly six hours of spectroscopic data of the transient, which I combined into a March spectrum and an April spectrum. We also observed the host galaxy for one hour and 40 minutes. I also got a spectrum of the two lobes of the host galaxy to confirm that they are not two galaxies overlapping in the line of sight and finding a host redshift of $z = 0.0975$. This is slightly larger than the estimated $z = 0.0964$ from the $H\alpha$ line in the SN peak spectrum, and gives the SN a velocity of $\sim 200 \text{ km s}^{-1}$ with respect to the host galaxy. The SN is $3.4''$ from the centre of the South-East lobe, which translates to $\sim 7 \text{ kpc}$ at the host redshift. This means that the SN velocity can be interpreted as its rotation velocity around the host.

The March spectrum shows four emission lines that are not in the host spectrum, and these had disappeared a few weeks later in the April spectrum. The transient nature of these lines shows that they are associated with the i -band detection in the binned ZTF light curve. As the follow-up was performed when the late-time signal was already fading again, both the photometric follow-up and the binned observations resulted in non-detection during the follow-up campaign. The remaining transient emission lines are too faint to be picked up in the broadband photometry and are only found during the long spectroscopic observations. This suggests that I only managed to obtain the tail end of the identified signal with spectroscopy.

The only way to match the transient emission lines to elements that can be expected in such an environment is to allow for a large velocity offset. The three lines around $8500 \text{ \AA}$ all require a velocity of $\sim 5150 \text{ km s}^{-1}$, matching them with one of the lines of the Ca II NIR triplet (the other two are likely contaminated by the host lines, see section 5.3.3.2), NI, and KI emission lines. Fitting the lines with Gaussians shows that they have similar widths, suggesting that the elements producing these lines are in the same location. The bluest transient line, however, is best matched by $H\beta$ with an offset of $5920 \pm 70 \text{ km s}^{-1}$. It is also broader than the other three lines. This

suggests that, if this line identification is correct, the H-rich material is in a different and faster moving cloud than the N, K, and Ca. It is important to note that although I did not use an order-blocking filter for these observations, none of the emission lines match a second-order diffraction of another emission line, showing that they are four different emission lines.

One problem with the interpretation of the emission line as being $H\beta$ is that in most cases, when $H\beta$ is detected, one would also expect to detect $H\alpha$ with at least the same strength. This is not the case here, nor is there any sign of an absorption feature at the expected location of $H\alpha$ that could completely wipe out a strong, relatively broad emission line like this. However, identifying this line as another element requires a very different velocity and leaves the open question of its strength and width, making $H\beta$ the best match despite its shortcomings.

The classification of SN 2020qxx as a SN Ia-CSM is due to the $H\alpha$ signal around the SN peak from interaction with CSM relatively close to the progenitor system. The late-time signal suggests a second region of CSM further from the progenitor system. Assuming that the inner CSM has a shell or torus-like structure, the interaction signature has the same total redshift as the SN. The late-time transient lines require a significant velocity offset, and their Gaussian profiles were much wider than that of the $H\alpha$ line. This shows that the outer CSM is moving much faster and has much more of a patchy structure. Velocity offsets of a few thousand km s^{-1} have been seen before (Maeda et al., 2010; Maguire et al., 2018). While $> 5\,000 \text{ km s}^{-1}$ is high, it is not unprecedented, as the Ca II lines in SN 2019bkc have been measured to be up to $10\,000 \text{ km s}^{-1}$ (Chen et al., 2020; Prentice et al., 2020).

CSM with velocity offsets in this regime has been seen before in some SNe IIn. It has been interpreted as material that has been accelerated as it has been swept up. If the CSM is clumpy or disk-like, this can result in a multi-peaked emission line where the redshifted peak is weakened as it comes from the receding end of the SN and has had to travel through dust-forming SN material (Leonard et al., 2000; Pozzo et al., 2004; Smith et al., 2015). In the case of SN 2020qxx, the late-time signal could

be interpreted as the ejecta running into a shell of H-rich CSM around peak light, resulting in the $H\alpha$ line in the classification spectrum. As the CSM is swept up, it is accelerated to at least $\sim 6\,000\text{ km s}^{-1}$ before a part of the accelerated ejecta crashes into a cloud of CSM, resulting in the interaction signal that was observed in 2024. Only the expanding part of the ejecta with accelerated CSM that interacts with this CSM cloud would be emitting light, explaining the offset. However, the viability of such a scenario needs proper modelling.

5.5 Conclusions

In this work I have presented a real-time search of late-time signals in 6914 Type Ia SNe in ZTF, using the latest observations to create the most up-to-date light curves for the objects in my sample. By binning the latest observations I was able to go up to nearly one magnitude beyond the single observation limit, which allowed me to find eight new faint signals during the 240-day monitoring effort. Using the NOT and GTC, three of these objects were subsequently followed up with photometry, and two objects were followed up with photometry and spectroscopy. My main conclusions are:

1. The requirement to bin observations to get deeper magnitude limits makes even the most real-time strategy lag slightly behind current events. In some cases this resulted in the discovery of a (short-lived) late-time signal while it was already declining again. Quick follow-up is required, but only the tail end of the signal can likely be observed.
2. Late-time signals are rare, and detections of late-time signals that are associated with the SN are contaminated with detections of associated objects in the same line of sight. This is more prevalent for SNe in or close to the centres of their host galaxies.
3. In five objects, the late-time detections have been determined to be due to nuclear activity. One of these was not followed up on as it was initially dismissed as a

false positive until it faded. This shows the importance of a good baseline when searching for faint signals. The three objects that were followed up showed an excess at the location of the host nucleus, clearly separated from the SN position due to the better resolution.

4. The excess in the SN 2019zbq host was followed up spectroscopically and showed both similarities and differences to previously studied ANTs, with the main difference being the absolute magnitude of the signal. Since the excess is 4.5 mag fainter than the host nucleus, extracting its spectrum required careful subtraction of the host signal, which complicates studying faint ANTs like these.
5. The late-time signal in SN 2020qzx stood out from the others, as this is a SN Ia-CSM at ~ 7 kpc from the host nucleus, making it significantly easier to follow up. While the signal was short-lived, I was able to follow it up spectroscopically, finding four transient emission lines that are best matched by four different elements at two different but highly blueshifted velocities ($5\,150 - 5\,920$ km s $^{-1}$). This suggests that the progenitor system created at least two regions of CSM: one nearby that resulted in the initial Ia-CSM classification and one further away that resulted in the late-time signal.

Faint late-time signals in SNe Ia due to interaction with distant CSM are rare, and even with techniques like binning observations to reach the deepest detection limits, only the brightest and strongest signals can be detected with current surveys. The upcoming Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST; [Ivezić et al., 2019](#)) will be able to detect objects several magnitudes deeper than ZTF, allowing the detection of fainter objects in a larger volume. Monitoring existing SN samples like those found by ZTF with LSST will be the best way to find faint late-time signals, which could then be followed up spectroscopically to find the chemical composition of the source of the signal. Constraining the properties of these rare events will allow us to probe the physics that govern SN Ia explosions more thoroughly and gain new insights to help explain them.

Conclusions and future work

In this thesis I have presented my work on the search for signs of interaction between the ejecta of SNe Ia and distant CSM. Due to the distance of the proposed CSM the interaction would start at late times, which I defined to be ≥ 100 days after the observed SN peak in the observer frame, and cause an extreme flattening or rebrightening in the SN light curve. I searched through three different samples of objects, these being the ZTF SN Ia DR2 (section 3), transients that were first detected between 2008 and 2018, including objects that were not SNe Ia (section 4), and a real-time search using all SNe Ia discovered by ZTF up to July 2023 in order to follow candidates up using the NOT and GTC (section 5).

To conduct this search in a systematic manner and to be able to recover the faintest signals hidden within the data, I built a custom analysis pipeline, which is presented in section 3.2.1 and tested through a simulated observing campaign in section 3.2.4. After an initial cleaning and calibration of the data, the late-time observations were binned together in bins with a size between 25 and 100 days. By doing so, I was able to push the detection limit beyond that of the individual observations, near the noise limit of the reference images used by ZTF in their difference imaging pipeline. This was done separately for the different ZTF *gri*-bands.

6.1 Conclusions

Searching through these three samples made one thing abundantly clear: late-time CSM interaction in SNe Ia is very rare, and the only way to systematically find them is by searching through large data sets, preferably catching them while they are active so they can be followed up on. By doing so, it will be possible to study properties of the CSM, such as its size and composition, gain new insights in SN Ia explosions and the physics that governs them, and place constraints on their progenitor systems. Based on the work I presented, my main conclusions are:

1. Across the three samples I analysed, the late-time data consisted of 11 973 unique spectroscopically confirmed Type Ia SNe, and an additional 2 728 transients in the pre-ZTF sample that were not classified as SN Ia or one of its subclasses. For some of these transients the data extended up to 15 years after they were first discovered. Out of these, I identified seven candidate objects showing signs of possible late-time CSM interaction, with the first detection of their late-time signal being between 250 and 2 200 days after discovery in the SN rest frame. Additionally, SN 2020qxz was spectroscopically confirmed to have a period of rebrightening around 1 150 days after peak light in the SN rest frame.
2. The largest fraction of objects that are flagged by the pipeline are false positives that show spurious detections in their binned light curves. This is a result of the wide net it casts in an attempt to minimize the chance of missing a faint signal. The visual inspection of these light curves, as well as inspecting the difference images using SNAP, showed that in most cases the spurious detections are either due to image issues (e.g., in a difficult environment) or due to software issues (e.g., imperfect image subtractions).
3. Another large group of objects that were flagged by the pipeline are bright nearby SNe, whose declining tail could still be recovered in the bins. Despite attempting to filter for these cases, some slipped through due to gaps in the observations or a tail that did not follow the simple decline model used in the test. These included most known SNe Ia-CSM in the sample and several very nearby objects whose decline displayed a kink around a year after the peak due to a change in their opacity.
4. I also recovered over a dozen pairs of sibling transients, two transients at (nearly) the exact same location. In one case, this is a line-of-sight effect, with AT 2018iml likely being a CV in the foreground of SN 2009hz. The other cases were sibling SNe occurring in the same galaxy. Some of these were already known, with both siblings already having been classified. In other cases, a possible but unconfirmed

(extinct) sibling transient fits the late-time signal. While such sibling transients are not a very common occurrence, modern surveys discover enough new transients that these rare events can be found regularly.

5. Most objects with a late-time signal that was not easily recognized as a false positive, SN tail, or sibling transient, were close to or on top of the nucleus of their host galaxy. These are difficult and busy environments, and even if the nucleus is not active, there is a lot of room to interpret a late-time signal as something related to the host nucleus instead. One class of possible contaminants I considered are ANTs, ambiguous nuclear transients that do not really fit in the AGN or TDE classes. The objects I recovered are broadly consistent but generally fainter than most known ANTs, suggesting the existence of a previously unknown population of faint ANTs.
6. Transients that were first detected before the start of ZTF but still visible after the survey started were never properly detected by ZTF, as they are in the reference images that were made shortly before the survey started in March 2018. As a result of this, difference imaging will leave a ghost, a negative imprint that is easily identified using SNAP (see section 3.2.3). Forced photometry still works, but the measured flux will be too small or even negative. Through a proper baseline correction, as is done at the start of my pipeline, these objects can still be recovered. AGN and variable stars can be recovered in the same way, though a proper baseline correction is more difficult.
7. Between the three samples, I recovered seven candidate objects, with their late-time signals starting between 0.7 and 5.9 years after the SN, lasting between ~ 100 days to over a year, and having an absolute magnitude of $M \sim -16$ mag or brighter in the ZTF bands. In section 4.3.4.1, I showed a first-order estimate for CSM mass based on broadband magnitude, assuming that the main emission mechanism is bremsstrahlung. Through this, I showed that it is possible to have interaction years after the SN that is this bright, but it requires a large CSM

mass in a thin shell. Due to the size of the shell and the finite speed of light, even a short interaction will be smeared out over a long period of time.

8. Through a quick detection and fast follow-up of SN 2020qxx, I was able to spectroscopically confirm the late-time signal through four emission lines that had disappeared a month later. This object stands out from the candidates for several reasons, as it is a known Ia-CSM that showed signs of interaction near its peak, had a very short period of late-time interaction that was detectable with ZTF, and was visually offset from its host galaxy (making spectroscopic follow-up much easier). The emission lines appear very blueshifted ($5\,150 - 5\,920 \text{ km s}^{-1}$), and include $\text{H}\beta$ emission. This could suggest that the second episode of CSM interaction is between accelerated material from the first episode of CSM interaction and a cloud of CSM on our side of the SN.
9. With the method that I used in this thesis to find faint late-time signals, even a real-time search is not completely in real time, but lags a few days to weeks behind. This is because several observations are needed before a newly appearing late-time signal can cause its bin to rise significantly above the noise. The need for several new observations can, however, be used to limit the amount of processing that has to be done on a daily basis. There will always be a trade-off between catching things quickly and catching them while they're faint, which can result in situations like that of SN 2020qxx, where follow-up was only possible for the fading part of the signal.
10. In section 3.3.3 I estimated that the rate of strong late-time interaction in SNe Ia is $< 0.5\%$ of SNe Ia, which translates to an absolute rate between 8_{-4}^{+20} and $54_{-26}^{+91} \text{ Gpc}^{-3} \text{ yr}^{-1}$ when assuming a constant SN Ia rate of $2.4 \times 10^{-5} \text{ Mpc}^{-3} \text{ yr}^{-1}$ for $z \leq 0.1$. While the pre-ZTF and real-time samples did not allow me to make separate robust estimates of this rate, the low number of recovered (candidate) objects confirms the rarity of late-time CSM interactions in SNe Ia.

6.2 Future work

6.2.1 Improvements to the pipeline

One thing that has become very clear, and is commented on in the conclusions of chapters 3, 4, and 5 is the limitations that the depth of ZTF places on my ability to recover such faint signals, even when using techniques like light curve binning. The upcoming Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST; Ivezić et al., 2019) will be able to detect objects several magnitudes deeper than ZTF, allowing the detection of fainter objects in a larger volume. With this in mind, there are several improvements that can be made to the pipeline in preparation for a similar search using LSST. This has the advantage that if this is done before the survey starts, everything can be done in real time with possible follow-up opportunities. Possible samples to monitor from the start of LSST include the ZTF SN Ia DR2 or a similar well-understood sample from ZTF that includes other types of transients and transients discovered after the start of 2021. Possible improvements to the pipeline include:

- Integrating the AGN check based on WISE data with the binning program, automatically generating SNAP movies, and other adaptations to streamline and automate the analysis process, making it easier for a real-time search.
- Lessen the amount of false positives flagged by the pipeline, and improve the methods for detecting SN peaks and removing bright SN tails.
- Generalize the procedure to work for other types of transients. One of the reasons why only the pre-ZTF sample included non-SN Ia transients was that the tail removal procedure was not built to handle those objects and could not be used.
- Machine learning. Different classes of objects recovered with the pipeline have different kinds of binned signals. Machine learning can streamline, automate, and speed up the classification of signal types and possibly identify new classes.

6.2.2 Alternative uses for the pipeline

As has been shown repeatedly throughout this thesis, the binning procedure is a catch-all method able to find any kind of faint signal. This is one of the reasons why signals coming from, e.g., declining SN tails, siblings, and nuclear transients are recovered even though the pipeline is not optimized for them. This means that the method and pipeline can easily be adapted for other types of transients or research goals, such as:

- Declining SNe to follow their evolution out to later epochs.
- Sibling transients to learn more about the environmental effects on SNe. The current pipeline only finds sibling transients when the first one is the brightest. To find sibling transients that are spatially far enough apart not to have one leave an imprint in the forced photometry light curve of the other transient, a different pipeline will be needed (e.g., SN 2024nqr and SN 2024pgd, see Fig. 2.1).
- (faint) ANTs, other nuclear transients, and estimating the containment rate for late-time signals in SNe in nuclear environments. This could use the pipeline version that was used for the pre-ZTF sample in chapter 4, utilizing the choice of different baseline regions. All that is needed is to define a new sample of positions to generate forced photometry light curves at.
- Precursor events. Everything that I have done and everything that can be done to look at what happens after the SN can be reversed to check what happened before the SN. This can, for instance, be done for LSST transients to see if ZTF observed any signs of precursor activity.
- Use in other surveys. There are many other surveys, both active and complete, whose data could be processed with my pipeline, as long as the data structure is transformed to what the code expects. ASASSN, ATLAS, GOTO (Gravitational-wave Optical Transient Observer, Steeghs et al., 2022; Dyer et al., 2024), and (i)PTF are just a few examples. This could be especially powerful if data from multiple surveys is used together to create long and/or well-observed light curves.

6.2.3 Other future work

Besides future work that involves using an adapted version of the binning procedure or pipeline, there are more questions to be answered based on the results I presented in this thesis. Some examples of these are:

- A proper CSM mass estimate using broadband observations. While the best estimates will rely on the emission lines that can be spectroscopically confirmed, if it is possible to get a decent CSM mass estimate given a certain assumption for the composition and state of the CSM, this should be useful for getting a first estimate or in cases where no spectroscopic follow-up was possible. The attempt I made in section 4.3.4.1 is a good step, but it can be improved upon.
- In section 5.4.3, I propose a scenario that could explain the large velocity offsets that I found in the late-time transient emission lines of SN 2020qxx. Whether or not such a scenario is actually viable, and how common the right conditions would be to generate such offset emission lines, will require detailed modelling and simulations.

Modern large-scale surveys that observe the night sky and monitor its changes on a regular basis generate a wealth of data that can be mined for years. By combining observations that have been taken over timespans of years, it is possible to follow the evolution of transients for much longer timescales than they are usually studied. Binning observations together is a powerful way to detect the faintest signals and can be used to find new and rare events that may or may not be connected to brighter types of events that are already known. Through this, it will be possible to find new clues to help us understand some of the most energetic events in the universe and the physical processes that govern them.

Acknowledgements

Based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under Grants No. AST-1440341 and AST-2034437 and a collaboration including current partners Caltech, IPAC, the Oskar Klein Center at Stockholm University, the University of Maryland, University of California, Berkeley, the University of Wisconsin at Milwaukee, University of Warwick, Ruhr University, Cornell University, Northwestern University and Drexel University. Operations are conducted by COO, IPAC, and UW.

This work was supported by the GROWTH project ([Kasliwal et al., 2019](#)) funded by the National Science Foundation under Grant No 1545949.

The Gordon and Betty Moore Foundation, through both the Data-Driven Investigator Program and a dedicated grant, provided critical funding for SkyPortal.

The ztfquery code was funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n°759194 - USNAC, PI: Rigault).

The Legacy Surveys consist of three individual and complementary projects: the Dark Energy Camera Legacy Survey (DECaLS; Proposal ID #2014B-0404; PIs: David Schlegel and Arjun Dey), the Beijing-Arizona Sky Survey (BASS; NOAO Prop. ID #2015A-0801; PIs: Zhou Xu and Xiaohui Fan), and the Mayall z-band Legacy Survey (MzLS; Prop. ID #2016A-0453; PI: Arjun Dey). DECaLS, BASS and MzLS together include data obtained, respectively, at the Blanco telescope, Cerro Tololo Inter-American Observatory, NSF’s NOIRLab; the Bok telescope, Steward Observatory, University of Arizona; and the Mayall telescope, Kitt Peak National Observatory, NOIRLab. Pipeline processing and analyses of the data were supported by NOIRLab and the Lawrence Berkeley National Laboratory (LBNL). The Legacy Surveys project is honored to be permitted to conduct astronomical research on Iolkam Du’ag (Kitt Peak), a mountain with particular significance to the Tohono O’odham Nation.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American

Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

The Pan-STARRS1 Surveys (PS1) and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation Grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

The intermediate Palomar Transient Factory project is a scientific collaboration among the California Institute of Technology, Los Alamos National Laboratory, the University of Wisconsin (Milwaukee), the Oskar Klein Centre, the Weizmann Institute of Science, the TANGO Program of the University System of Taiwan, and the Kavli Institute for the Physics and Mathematics of the Universe.

ASAS-SN is funded by Gordon and Betty Moore Foundation grants GBMF5490 and GBMF10501 and the Alfred P. Sloan Foundation grant G-2021-14192.

Based on observations made with the Nordic Optical Telescope, owned in collaboration by the University of Turku and Aarhus University, and operated jointly by Aarhus University, the University of Turku and the University of Oslo, representing Denmark, Finland and Norway, the University of Iceland and Stockholm University at the Observatorio del Roque de los Muchachos, La Palma, Spain, of the Instituto de Astrofísica de Canarias.

The data presented here were obtained [in part] with ALFOSC, which is provided by the Instituto de Astrofísica de Andalucía (IAA) under a joint agreement with the University of Copenhagen and NOT.

Based on observations made with the Gran Telescopio Canarias (GTC), installed at the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias, on the island of La Palma. This work is (partly) based on data obtained with the instrument OSIRIS, built by a Consortium led by the Instituto de Astrofísica de Canarias in collaboration with the Instituto de Astronomía of the Universidad Autónoma de México. OSIRIS was funded by GRANTECAN and the National Plan of Astronomy and Astrophysics of the Spanish Government.

This research made use of PYPEIT,¹ a Python package for semi-automated reduction of astronomical slit-based spectroscopy (Prochaska et al., 2020; Prochaska et al., 2020).

This work made use of Astropy:² a community-developed core Python package and an ecosystem of tools and resources for astronomy (Astropy Collaboration et al., 2013, 2018, 2022).

Many people have contributed to lmfit. The attribution of credit in a project such as this is difficult to get perfect, and there are no doubt important contributions that are missing or under-represented here. Please consider this file as part of the code and documentation that may have bugs that need fixing.

Some of the largest and most important contributions (in approximate order of size of the contribution to the existing code) are from: Matthew Newville wrote the original version and maintains the project.

Renee Otten wrote the brute force method, implemented the basin-hopping and AMPGO global solvers, implemented uncertainty calculations for scalar minimizers and has greatly improved the code, testing, and documentation and overall project.

Till Stensitzki wrote the improved estimates of confidence intervals, and contributed many tests, bug fixes, and documentation.

A. R. J. Nelson added differential_evolution, emcee, and greatly improved the code, docstrings, and overall project.

Antonino Ingargiola wrote much of the high level Model code and has provided many bug fixes and improvements.

Daniel B. Allan wrote much of the original version of the high level Model code, and many improvements to the testing and documentation.

Austen Fox fixed many of the built-in model functions and improved the testing and documentation of these. Michal Rawlik added plotting capabilities for Models.

¹<https://pypeit.readthedocs.io/en/latest/>

²<http://www.astropy.org>

The method used for placing bounds on parameters was derived from the clear description in the MINUIT documentation, and adapted from J. J. Helmus’s Python implementation in `leastsqbounds.py`.

E. O. Le Bigot wrote the `uncertainties` package, a version of which was used by `lmfit` for many years, and is now an external dependency.

The original AMPGO code came from Andrea Gavana and was adopted for `lmfit`.

The propagation of parameter uncertainties to uncertainties in a Model was adapted from the excellent description at ¹, which references the original work of: J. Wolberg, *Data Analysis Using the Method of Least Squares*, 2006, Springer.

Additional patches, bug fixes, and suggestions have come from Faustin Carter, Christoph Deil, Francois Boulogne, Thomas Caswell, Colin Brosseau, nmearl, Gustavo Pasquevich, Clemens Prescher, LiCode, Ben Gamari, Yoav Roam, Alexander Stark, Alexandre Beelen, Andrey Aristov, Nicholas Zobrist, Ethan Welty, Julius Zimmermann, Mark Dean, Arun Persaud, Ray Osborn, @lneuhaus, Marcel Stimberg, Yoshiera Huang, Leon Foks, Sebastian Weigand, Florian LB, Michael Hudson-Doyle, Ruben Verweij, @jedzill4, @spalato, Jens Hedegaard Nielsen, Martin Majli, Kristian Meyer, @azelcer, Ivan Usov, and many others.

The `lmfit` code obviously depends on, and owes a very large debt to the code in `scipy.optimize`. Several discussions on the SciPy-user and `lmfit` mailing lists have also led to improvements in this code.

Other software: `numpy` (Harris et al., 2020), `pandas` (pandas development team, 2020; Wes McKinney, 2010), `matplotlib` (Hunter, 2007), `sncosmo` (Barbary et al., 2021), `simsurvey` (Feindt et al., 2019a)

¹<https://www.astro.rug.nl/software/kapteyn/kmpfittutorial.html#confidence-and-prediction-intervals>

References

- Aamer, A., Nicholl, M., Jerkstrand, A., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 11970 (Cited on page [vii.](#))
- Adelman-McCarthy, J. K., Agüeros, M. A., Allam, S. S., et al. 2006, *Astrophysical Journal Supplemental Series*, 162, 38 (Cited on pages [36](#), [90](#) and [158.](#))
- Ailawadhi, B., Dastidar, R., Misra, K., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 248 (Cited on pages [87](#) and [137.](#))
- Alard, C. & Lupton, R. H. 1998, *Astrophysical Journal*, 503, 325 (Cited on page [48.](#))
- Aldering, G., Antilogus, P., Bailey, S., et al. 2006, *Astrophysical Journal*, 650, 510 (Cited on page [23.](#))
- Alsabti, A. W. & Murdin, P. 2017, Supernovae and Supernova Remnants: The Big Picture in Low Resolution, ed. A. W. Alsabti & P. Murdin, 3 (Cited on page [6.](#))
- Amenouche, M., Rosnet, P., Smith, M., et al. 2025, *Astronomy & Astrophysics*, 694, A3 (Cited on page [viii.](#))
- Anderson, J. P., González-Gaitán, S., Hamuy, M., et al. 2014, *Astrophysical Journal*, 786, 67 (Cited on page [141.](#))
- Andrews, J. E., Sand, D. J., Valenti, S., et al. 2019, *Astrophysical Journal*, 885, 43 (Cited on pages [87](#), [121](#), [122](#) and [123.](#))
- Antonucci, R. 1993, *Annual Review of Astronomy & Astrophysics*, 31, 473 (Cited on page [10.](#))
- Arcavi, I., Howell, D. A., Kasen, D., et al. 2017, *Nature*, 551, 210 (Cited on pages [124](#), [125](#) and [153.](#))
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *Astrophysical Journal*, 935, 167 (Cited on page [196.](#))
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *Astronomical Journal*, 156, 123 (Cited on page [196.](#))
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *Astronomy & Astrophysics*, 558, A33 (Cited on page [196.](#))
- Aubert, M., Rosnet, P., Popovic, B., et al. 2025, *Astronomy & Astrophysics*, 694, A7 (Cited on page [viii.](#))
- Barbary, K., Bailey, S., Barentsen, G., et al. 2021, SNCosmo (Cited on pages [64](#) and [197.](#))
- Barbon, R., Ciatti, F., & Rosino, L. 1979, *Astronomy & Astrophysics*, 72, 287 (Cited on page [141.](#))
- Barmantloo, S., Jerkstrand, A., Iwamoto, K., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 533, 1251 (Cited on page [178.](#))
- Becker, A. 2015, HOTPANTS: High Order Transform of PSF AND Template Subtraction, Astrophysics Source Code Library, record ascl:1504.004 (Cited on pages [47](#) and [98.](#))

- Bedding, T. R. & Zijlstra, A. A. 1998, *Astrophysical Journal Letters*, 506, L47 (Cited on page 12.)
- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019a, *Publications of the Astronomical Society of the Pacific*, 131, 068003 (Cited on pages 33 and 68.)
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019b, *Publications of the Astronomical Society of the Pacific*, 131, 018002 (Cited on pages 33 and 118.)
- Berry, R. & Burnell, J. 2000, *The handbook of astronomical image processing* (Cited on page 39.)
- Biswas, R., Goobar, A., Dhawan, S., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 509, 5340 (Cited on page 78.)
- Bose, S., Dong, S., Pastorello, A., et al. 2018, *Astrophysical Journal*, 853, 57 (Cited on page 123.)
- Brennan, S. J. & Fraser, M. 2022, *Astronomy & Astrophysics*, 667, A62 (Cited on pages 47, 158 and 166.)
- Brown, P. J., Hosseinzadeh, G., Jha, S. W., et al. 2019, *Astrophysical Journal*, 877, 152 (Cited on pages 121, 122 and 123.)
- Bulla, M., Sim, S. A., Pakmor, R., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 455, 1060 (Cited on page 19.)
- Burgaz, U., Maeda, K., Kalomeni, B., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 502, 4112 (Cited on page 123.)
- Burgaz, U., Maguire, K., Dimitriadis, G., et al. 2025a, *Astronomy & Astrophysics*, 694, A9 (Cited on pages viii, 75 and 160.)
- Burgaz, U., Maguire, K., Dimitriadis, G., et al. 2025b, *Astronomy & Astrophysics*, 694, A13 (Cited on page viii.)
- Buzzoni, B., Delabre, B., Dekker, H., et al. 1984, *The Messenger*, 38, 9 (Cited on pages 36 and 97.)
- Cai, Y. Z., Pastorello, A., Fraser, M., et al. 2022, *Astronomy & Astrophysics*, 667, A4 (Cited on page 140.)
- Camacho-Neves, Y., Jha, S. W., Barna, B., et al. 2023, *Astrophysical Journal*, 951, 67 (Cited on page 81.)
- Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *Astrophysical Journal*, 345, 245 (Cited on pages 66 and 69.)
- Carreres, B., Rosselli, D., Bautista, J. E., et al. 2025, *Astronomy & Astrophysics*, 694, A8 (Cited on page vii.)
- Chambers, K. C., Boer, T. D., Bulger, J., et al. 2021, *Transient Name Server Discovery Report*, 2021-2390, 1 (Cited on page 126.)
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560 (Cited on pages 36 and 158.)
- Chandra, P., Chevalier, R. A., James, N. J. H., & Fox, O. D. 2022, *Monthly Notices of the Royal Astronomical Society*, 517, 4151 (Cited on page 123.)
- Chandrasekhar, S. 1931, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 11, 592 (Cited on page 2.)
- Chen, P., Dong, S., Stritzinger, M. D., et al. 2020, *Astrophysical Journal Letters*, 889, L6 (Cited on pages 175 and 183.)

- Chen, T.-W., Schady, P., Xiao, L., et al. 2017, *Astrophysical Journal Letters*, 849, L4 (Cited on page 123.)
- Chevalier, R. A. & Fransson, C. 2017, in *Handbook of Supernovae*, ed. A. W. Alsabti & P. Murdin, 875 (Cited on page 22.)
- Chugai, N. N., Blinnikov, S. I., Cumming, R. J., et al. 2004, *Monthly Notices of the Royal Astronomical Society*, 352, 1213 (Cited on page 25.)
- Chugai, N. N. & Yungelson, L. R. 2004, *Astronomy Letters*, 30, 65 (Cited on page 142.)
- Coughlin, M. W., Bloom, J. S., Nir, G., et al. 2023, *Astrophysical Journal Supplemental Series*, 267, 31 (Cited on pages 78, 126 and 156.)
- Dahiwal, A. & Fremling, C. 2021, *Transient Name Server Classification Report*, 2021-135, 1 (Cited on page 21.)
- De, K., Kasliwal, M. M., Polin, A., et al. 2019, *Astrophysical Journal Letters*, 873, L18 (Cited on page 8.)
- de Vaucouleurs, G., de Vaucouleurs, A., Corwin, Herold G., J., et al. 1991, *Third Reference Catalogue of Bright Galaxies* (Cited on page 4.)
- Deckers, M., Graur, O., Maguire, K., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 4414 (Cited on page vii.)
- Deckers, M., Groh, J. H., Boian, I., & Farrell, E. J. 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 3726 (Cited on page 123.)
- Deckers, M., Maguire, K., Shingles, L., et al. 2025, *Astronomy & Astrophysics*, 694, A12 (Cited on page viii.)
- Dekany, R., Smith, R. M., Riddle, R., et al. 2020, *Publications of the Astronomical Society of the Pacific*, 132, 038001 (Cited on pages 33 and 69.)
- Dexter, J. & Kasen, D. 2013, *Astrophysical Journal*, 772, 30 (Cited on page 10.)
- Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *Astronomical Journal*, 157, 168 (Cited on pages 36 and 98.)
- Dhawan, S., Mortsell, E., Johansson, J., et al. 2024, arXiv e-prints, arXiv:2406.01434 (Cited on pages viii, 81, 125 and 148.)
- Di Valentino, E., Mena, O., Pan, S., et al. 2021, *Classical and Quantum Gravity*, 38, 153001 (Cited on page 9.)
- Dilday, B., Howell, D. A., Cenko, S. B., et al. 2012, *Science*, 337, 942 (Cited on pages 8, 19, 23 and 141.)
- Dimitriadis, G., Burgaz, U., Deckers, M., et al. 2025, *Astronomy & Astrophysics*, 694, A10 (Cited on pages viii, 14, 15, 75 and 160.)
- Dimitriadis, G., Sullivan, M., Kerzendorf, W., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 468, 3798 (Cited on pages 62, 65 and 81.)
- Dong, Y., Valenti, S., Polin, A., et al. 2022, *Astrophysical Journal*, 934, 102 (Cited on page vii.)
- Dubay, L. O., Tucker, M. A., Do, A., Shappee, B. J., & Anand, G. S. 2022, *Astrophysical Journal*, 926, 98 (Cited on pages 26 and 149.)

- Duev, D. A., Mahabal, A., Masci, F. J., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 489, 3582 (Cited on page 78.)
- Dutta, A., Sahu, D. K., Anupama, G. C., et al. 2022, *Astrophysical Journal*, 925, 217 (Cited on pages 8, 75 and 81.)
- Dutta, A., Singh, A., Anupama, G. C., Sahu, D. K., & Kumar, B. 2021, *Monthly Notices of the Royal Astronomical Society*, 503, 896 (Cited on page 123.)
- Dyer, M. J., Ackley, K., Jiménez-Ibarra, F., et al. 2024, arXiv e-prints, arXiv:2407.17176 (Cited on page 192.)
- Elias-Rosa, N., Brennan, S. J., Benetti, S., et al. 2024, *Astronomy & Astrophysics*, 686, A13 (Cited on page 25.)
- Feindt, U., M., R., V., B., & J., N. 2019a, ufeindt/simsurvey: 0.6.0 (Cited on pages 64 and 197.)
- Feindt, U., Nordin, J., Rigault, M., et al. 2019b, , 2019, 005 (Cited on pages 64, 68 and 69.)
- Filippenko, A. V. 1982, *Publications of the Astronomical Society of the Pacific*, 94, 715 (Cited on page 32.)
- Filippenko, A. V., Richmond, M. W., Matheson, T., et al. 1992, *Astrophysical Journal Letters*, 384, L15 (Cited on page 15.)
- Fink, M., Röpke, F. K., Hillebrandt, W., et al. 2010, *Astronomy & Astrophysics*, 514, A53 (Cited on page 18.)
- Fisher, R. & Jumper, K. 2015, *Astrophysical Journal*, 805, 150 (Cited on page 15.)
- Forster, F., Bauer, F. E., Arredondo, J., et al. 2019, Transient Name Server Discovery Report, 2019-2185, 1 (Cited on page 126.)
- Fox, O. D., Filippenko, A. V., Skrutskie, M. F., et al. 2013, *Astronomical Journal*, 146, 2 (Cited on page 25.)
- Fraser, M., Inserra, C., Jerkstrand, A., et al. 2013a, *Monthly Notices of the Royal Astronomical Society*, 433, 1312 (Cited on page 135.)
- Fraser, M., Magee, M., Kotak, R., et al. 2013b, *Astrophysical Journal Letters*, 779, L8 (Cited on page 25.)
- Fremling, C., Miller, A. A., Sharma, Y., et al. 2020, *Astrophysical Journal*, 895, 32 (Cited on pages 77 and 154.)
- Friesen, B., Baron, E., Parrent, J. T., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 467, 2392 (Cited on page 66.)
- Frohmaier, C., Sullivan, M., Maguire, K., & Nugent, P. 2018, *Astrophysical Journal*, 858, 50 (Cited on pages 14 and 102.)
- Frohmaier, C., Sullivan, M., Nugent, P. E., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 2308 (Cited on pages 68 and 102.)
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *Astronomy & Astrophysics*, 595, A1 (Cited on page 36.)
- Gal-Yam, A. 2019, *Annual Review of Astronomy & Astrophysics*, 57, 305 (Cited on page 10.)
- Gal-Yam, A., Bruch, R., Schulze, S., et al. 2022, *Nature*, 601, 201 (Cited on page 24.)

- Gal-Yam, A., Yaron, O., & Schulze, S. 2024, *Transient Name Server AstroNote*, 239, 1 (Cited on page 25.)
- Gibson, B. K., Stetson, P. B., Freedman, W. L., et al. 2000, *Astrophysical Journal*, 529, 723 (Cited on page 12.)
- Ginolin, M., Rigault, M., Copin, Y., et al. 2025a, *Astronomy & Astrophysics*, 694, A4 (Cited on page viii.)
- Ginolin, M., Rigault, M., Smith, M., et al. 2025b, *Astronomy & Astrophysics*, 695, A140 (Cited on page viii.)
- Gobierno de España. 1988, Ley 31/1988, de 31 de octubre, sobre Protección de la Calidad Astronómica de los Observatorios del Instituto de Astrofísica de Canarias, Boletín Oficial del Estado, núm. 262, de 1 de noviembre de 1988, pp. 31203-31205 (Cited on page 51.)
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019a, *Publications of the Astronomical Society of the Pacific*, 131, 078001 (Cited on page 33.)
- Graham, M. L., Fremling, C., Perley, D. A., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 511, 241 (Cited on pages 78 and 125.)
- Graham, M. L., Harris, C. E., Nugent, P. E., et al. 2019b, *Astrophysical Journal*, 871, 62 (Cited on pages 26, 27, 56, 66, 67, 71, 100, 104, 107, 110, 134, 135, 136, 141, 150, 152, 159 and 179.)
- Graham, M. L., Nugent, P. E., Sullivan, M., et al. 2015a, *Monthly Notices of the Royal Astronomical Society*, 454, 1948 (Cited on page 66.)
- Graham, M. L., Valenti, S., Fulton, B. J., et al. 2015b, *Astrophysical Journal*, 801, 136 (Cited on page 178.)
- Graur, O., Zurek, D., Shara, M. M., et al. 2016, *Astrophysical Journal*, 819, 31 (Cited on pages 23 and 82.)
- Groh, J. H., Meynet, G., Georgy, C., & Ekström, S. 2013, *Astronomy & Astrophysics*, 558, A131 (Cited on page 24.)
- Guillochon, J., Parrent, J., Kelley, L. Z., & Margutti, R. 2017, *Astrophysical Journal*, 835, 64 (Cited on page 114.)
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *Astronomical Journal*, 131, 2332 (Cited on pages 36, 90 and 158.)
- Gutiérrez, C. P., Pastorello, A., Jerkstrand, A., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 499, 974 (Cited on page 123.)
- Guy, J., Astier, P., Baumont, S., et al. 2007, *Astronomy & Astrophysics*, 466, 11 (Cited on page 64.)
- Hammerstein, E., van Velzen, S., Gezari, S., et al. 2023, *Astrophysical Journal*, 942, 9 (Cited on pages 88, 90 and 92.)
- Hamuy, M., Phillips, M., Suntzeff, N., & Maza, J. 2003a, , 8151, 2 (Cited on page 23.)
- Hamuy, M., Phillips, M. M., Suntzeff, N. B., et al. 2003b, *Nature*, 424, 651 (Cited on pages 23 and 139.)
- Harris, C. E., Nugent, P. E., Horesh, A., et al. 2018, *Astrophysical Journal*, 868, 21 (Cited on pages 100 and 152.)

- Harris, C. E., Nugent, P. E., & Kasen, D. N. 2016, *Astrophysical Journal*, 823, 100 (Cited on page 134.)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357 (Cited on page 197.)
- Harvey, L., Maguire, K., Burgaz, U., et al. 2025, *Astronomy & Astrophysics*, 695, A264 (Cited on page viii.)
- Harvey, L., Maguire, K., Magee, M. R., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 4444 (Cited on pages vii and 8.)
- Hatsukade, B., Morokuma-Matsui, K., Hayashi, M., et al. 2020, *Publications of the Astronomical Society of Japan*, 72, L6 (Cited on page 123.)
- Hicken, M., Garnavich, P. M., Prieto, J. L., et al. 2007, *Astrophysical Journal Letters*, 669, L17 (Cited on page 15.)
- Hill, J. M., Green, R. F., & Slagle, J. H. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6267, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. L. M. Stepp, 62670Y (Cited on page 65.)
- Hills, J. G. 1975, *Nature*, 254, 295 (Cited on page 11.)
- Hinkle, J. T. 2024, *Monthly Notices of the Royal Astronomical Society*, 531, 2603 (Cited on page 12.)
- Hinkle, J. T., Holoiën, T. W. S., Shappee, B. J., et al. 2022, *Astrophysical Journal*, 930, 12 (Cited on pages 12, 140, 148, 168, 180 and 181.)
- Hinkle, J. T., Shappee, B. J., Auchettl, K., et al. 2024, arXiv e-prints, arXiv:2405.08855 (Cited on pages 12 and 139.)
- Holtzman, J. A., Marriner, J., Kessler, R., et al. 2008, *Astronomical Journal*, 136, 2306 (Cited on page 56.)
- Hosseinzadeh, G., Valenti, S., McCully, C., et al. 2018, *Astrophysical Journal*, 861, 63 (Cited on page 123.)
- Howell, D. A., Sullivan, M., Nugent, P. E., et al. 2006, *Nature*, 443, 308 (Cited on page 15.)
- Howell, S. B. 1989, *Publications of the Astronomical Society of the Pacific*, 101, 616 (Cited on page 43.)
- Howell, S. B. 2006, *Handbook of CCD Astronomy*, Vol. 5 (Cited on pages 29, 37 and 39.)
- Hoyle, F. & Fowler, W. A. 1960, *Astrophysical Journal*, 132, 565 (Cited on page 7.)
- Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90 (Cited on page 197.)
- Hviding, R. E., Hainline, K. N., Rieke, M., et al. 2022, *Astronomical Journal*, 163, 224 (Cited on pages 83, 85, 128, 149, 164 and 168.)
- Iben, I., J. & Tutukov, A. V. 1984, *Astrophysical Journal Supplemental Series*, 54, 335 (Cited on page 18.)
- Inserra, C., Fraser, M., Smartt, S. J., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 459, 2721 (Cited on pages 19, 141 and 142.)
- Irani, I., Prentice, S. J., Schulze, S., et al. 2022, *Astrophysical Journal*, 927, 10 (Cited on page 107.)
- Isern, J., Mochkovitch, R., García-Berro, E., & Hernanz, M. 1997, *Astrophysical Journal*, 485, 308 (Cited on page 2.)

- Itagaki, K. 2024, Transient Name Server Discovery Report, 2024-39, 1 (Cited on page 4.)
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *Astrophysical Journal*, 873, 111 (Cited on pages 109, 154, 185 and 191.)
- Izzo, L., Thöne, C. C., García-Benito, R., et al. 2018, *Astronomy & Astrophysics*, 610, A11 (Cited on page 123.)
- Jayasinghe, T., Stanek, K. Z., Kochanek, C. S., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 961 (Cited on page 36.)
- Jordan, George C., I., Perets, H. B., Fisher, R. T., & van Rossum, D. R. 2012, *Astrophysical Journal Letters*, 761, L23 (Cited on pages 14 and 81.)
- Kankare, E., Kotak, R., Mattila, S., et al. 2017, *Nature Astronomy*, 1, 865 (Cited on page 12.)
- Kasen, D. 2006, *Astrophysical Journal*, 649, 939 (Cited on page 8.)
- Kasen, D., Woosley, S. E., & Heger, A. 2011, *Astrophysical Journal*, 734, 102 (Cited on page 10.)
- Kashi, A. & Soker, N. 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 1466 (Cited on page 17.)
- Kasliwal, M. M., Cannella, C., Bagdasaryan, A., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 038003 (Cited on pages 78 and 194.)
- Kasliwal, M. M., Kulkarni, S. R., Gal-Yam, A., et al. 2012, *Astrophysical Journal*, 755, 161 (Cited on page 14.)
- Katsuda, S., Mori, K., Maeda, K., et al. 2015, *Astrophysical Journal*, 808, 49 (Cited on page 24.)
- Kawabata, M., Kawabata, K. S., Maeda, K., et al. 2018, *Publications of the Astronomical Society of Japan*, 70, 111 (Cited on page 81.)
- Kelly, P. L., Hicken, M., Burke, D. L., Mandel, K. S., & Kirshner, R. P. 2010, *Astrophysical Journal*, 715, 743 (Cited on page 9.)
- Khokhlov, A. M. 1991, *Astronomy & Astrophysics*, 245, L25 (Cited on page 17.)
- Kilpatrick, C. D., Drout, M. R., Auchettl, K., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 2073 (Cited on page 87.)
- Kippenhahn, R., Weigert, A., & Weiss, A. 2013, *Stellar Structure and Evolution* (Cited on pages 1, 3 and 4.)
- Kool, E. C., Johansson, J., Sollerman, J., et al. 2023, *Nature*, 617, 477 (Cited on pages 24, 77, 105, 106 and 107.)
- Krause, O., Tanaka, M., Usuda, T., et al. 2008, *Nature*, 456, 617 (Cited on page 23.)
- Kromer, M., Fink, M., Stanishev, V., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 429, 2287 (Cited on pages 14 and 81.)
- Kulkarni, S. R. 2013, *The Astronomer's Telegram*, 4807, 1 (Cited on page 36.)
- Kumar, B., Eswaraiah, C., Singh, A., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 3089 (Cited on page 123.)
- Kuncarayakti, H., Maeda, K., Ashall, C. J., et al. 2018, *Astrophysical Journal Letters*, 854, L14 (Cited on pages 121, 122 and 123.)

- Kwok, L. A., Siebert, M. R., Johansson, J., et al. 2024, *Astrophysical Journal*, 966, 135 (Cited on page [vii](#).)
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 1395 (Cited on pages [36](#), [64](#) and [66](#).)
- Leloudas, G., Hsiao, E. Y., Johansson, J., et al. 2015, *Astronomy & Astrophysics*, 574, A61 (Cited on pages [24](#) and [26](#).)
- Leonard, D. C. 2007, *Astrophysical Journal*, 670, 1275 (Cited on page [17](#).)
- Leonard, D. C., Filippenko, A. V., Barth, A. J., & Matheson, T. 2000, *Astrophysical Journal*, 536, 239 (Cited on pages [176](#) and [183](#).)
- Li, S., Liang, Y.-F., Liao, N.-H., Lei, L., & Fan, Y.-Z. 2024, arXiv e-prints, arXiv:2407.05968 (Cited on page [123](#).)
- Lin, W., Wang, X., Yan, L., et al. 2023, *Nature Astronomy*, 7, 779 (Cited on page [123](#).)
- Lintott, C., Schawinski, K., Bamford, S., et al. 2011, *Monthly Notices of the Royal Astronomical Society*, 410, 166 (Cited on page [107](#).)
- Liu, C., Miller, A. A., Boos, S. J., et al. 2023a, *Astrophysical Journal*, 958, 178 (Cited on page [vii](#).)
- Liu, Z.-W., Röpke, F. K., & Han, Z. 2023b, *Research in Astronomy and Astrophysics*, 23, 082001 (Cited on page [18](#).)
- Livne, E. & Arnett, D. 1995, *Astrophysical Journal*, 452, 62 (Cited on page [18](#).)
- Maeda, K., Tanaka, M., Nomoto, K., et al. 2007, *Astrophysical Journal*, 666, 1069 (Cited on page [10](#).)
- Maeda, K., Taubenberger, S., Sollerman, J., et al. 2010, *Astrophysical Journal*, 708, 1703 (Cited on page [183](#).)
- Maguire, K., Sim, S. A., Shingles, L., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 477, 3567 (Cited on page [183](#).)
- Marcowith, A., Bret, A., Bykov, A., et al. 2016, *Reports on Progress in Physics*, 79, 046901 (Cited on page [22](#).)
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 018003 (Cited on pages [33](#) and [120](#).)
- Mauerhan, J. C., Smith, N., Williams, G. G., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 6090 (Cited on page [123](#).)
- Maund, J. R., Steele, I., Jermak, H., Wheeler, J. C., & Wiersema, K. 2019, *Monthly Notices of the Royal Astronomical Society*, 482, 4057 (Cited on page [123](#).)
- Mazzali, P. A., Röpke, F. K., Benetti, S., & Hillebrandt, W. 2007, *Science*, 315, 825 (Cited on page [17](#).)
- Mazzali, P. A., Sullivan, M., Filippenko, A. V., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 2631 (Cited on page [66](#).)
- Mazzali, P. A., Sullivan, M., Hachinger, S., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 439, 1959 (Cited on pages [65](#), [68](#) and [87](#).)
- McCully, C., Crawford, S., Kovacs, G., et al. 2018, *astropy/astrocrappy: v1.0.5 Zenodo Release* (Cited on page [40](#).)

- McCully, C., Jha, S. W., Scalzo, R. A., et al. 2022, *Astrophysical Journal*, 925, 138 (Cited on page 81.)
- Merline, W. J. & Howell, S. B. 1995, *Experimental Astronomy*, 6, 163 (Cited on pages 43, 44 and 45.)
- Miller, A. A., Silverman, J. M., Butler, N. R., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 404, 305 (Cited on page 25.)
- Miller, A. A., Yao, Y., Bulla, M., et al. 2020, *Astrophysical Journal*, 902, 47 (Cited on pages 57, 117 and 157.)
- Mirizzi, A., Tamborra, I., Janka, H. T., et al. 2016, *Nuovo Cimento Rivista Serie*, 39, 1 (Cited on page 4.)
- Mo, G., De, K., Wiston, E., et al. 2025, *Astrophysical Journal Letters*, 980, L33 (Cited on page 149.)
- Mochkovitch, R. 1983, *Astronomy & Astrophysics*, 122, 212 (Cited on page 2.)
- Moran, S., Fraser, M., Kotak, R., et al. 2023, *Astronomy & Astrophysics*, 669, A51 (Cited on page 123.)
- Nagao, T., Cikota, A., Patat, F., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 489, L69 (Cited on page 123.)
- Nakaoka, T., Kawabata, K. S., Maeda, K., et al. 2018, *Astrophysical Journal*, 859, 78 (Cited on page 123.)
- Nicholl, M., Berger, E., Margutti, R., et al. 2017, *Astrophysical Journal Letters*, 845, L8 (Cited on page 123.)
- Nomoto, K. 1982, *Astrophysical Journal*, 253, 798 (Cited on page 16.)
- Nomoto, K., Suzuki, T., Deng, J., Uenishi, T., & Hachisu, I. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 342, 1604-2004: Supernovae as Cosmological Lighthouses, ed. M. Turatto, S. Benetti, L. Zampieri, & W. Shea, 105 (Cited on pages 19 and 23.)
- Nordin, J., Brinnel, V., Santen, J. V., et al. 2021, *Transient Name Server Discovery Report*, 2021-2254, 1 (Cited on page 126.)
- Nugent, P. E., Sullivan, M., Cenko, S. B., et al. 2011, *Nature*, 480, 344 (Cited on page 65.)
- Ofek, E. O., Laher, R., Law, N., et al. 2012, *Publications of the Astronomical Society of the Pacific*, 124, 62 (Cited on page 44.)
- Pakmor, R., Kromer, M., Röpke, F. K., et al. 2010, *Nature*, 463, 61 (Cited on page 18.)
- Pakmor, R., Kromer, M., Taubenberger, S., et al. 2012, *Astrophysical Journal Letters*, 747, L10 (Cited on page 18.)
- pandas development team, T. 2020, pandas-dev/pandas: Pandas (Cited on page 197.)
- Parrent, J. T., Howell, D. A., Friesen, B., et al. 2012, *Astrophysical Journal Letters*, 752, L26 (Cited on page 65.)
- Patat, F. 2005, *Monthly Notices of the Royal Astronomical Society*, 357, 1161 (Cited on pages 23 and 82.)
- Patat, F., Chandra, P., Chevalier, R., et al. 2007, *Science*, 317, 924 (Cited on page 23.)
- Patat, F., Cordiner, M. A., Cox, N. L. J., et al. 2013, *Astronomy & Astrophysics*, 549, A62 (Cited on page 66.)

- Patnaude, D. J., Badenes, C., Park, S., & Laming, J. M. 2012, *Astrophysical Journal*, 756, 6 (Cited on page 24.)
- Pauli, W. 1925, *Zeitschrift fur Physik*, 31, 765 (Cited on page 1.)
- Pereira, R., Thomas, R. C., Aldering, G., et al. 2013, *Astronomy & Astrophysics*, 554, A27 (Cited on pages 65 and 66.)
- Perets, H. B., Gal-Yam, A., Mazzali, P. A., et al. 2010, *Nature*, 465, 322 (Cited on page 14.)
- Perley, D. A., Fremling, C., Sollerman, J., et al. 2020, *Astrophysical Journal*, 904, 35 (Cited on pages 77 and 154.)
- Perley, D. A., Schulze, S., Fremling, C., et al. 2021, *Transient Name Server AstroNote*, 156, 1 (Cited on page 21.)
- Peterson, B. M. & Horne, K. 2004, *Astronomische Nachrichten*, 325, 248 (Cited on page 11.)
- Phillips, M. M. 1993, *Astrophysical Journal Letters*, 413, L105 (Cited on page 9.)
- Piersanti, L., Gagliardi, S., Iben, Icko, J., & Tornambé, A. 2003a, *Astrophysical Journal*, 598, 1229 (Cited on page 18.)
- Piersanti, L., Gagliardi, S., Iben, Icko, J., & Tornambé, A. 2003b, *Astrophysical Journal*, 583, 885 (Cited on page 18.)
- Piersanti, L., Tornambé, A., & Yungelson, L. R. 2014, *Monthly Notices of the Royal Astronomical Society*, 445, 3239 (Cited on page 17.)
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *Astronomy & Astrophysics*, 641, A6 (Cited on page 28.)
- Podsiadlowski, P., Joss, P. C., & Hsu, J. J. L. 1992, *Astrophysical Journal*, 391, 246 (Cited on page 6.)
- Popovic, B., Rigault, M., Smith, M., et al. 2025, *Astronomy & Astrophysics*, 694, A5 (Cited on page viii.)
- Pozzo, M., Meikle, W. P. S., Fassia, A., et al. 2004, *Monthly Notices of the Royal Astronomical Society*, 352, 457 (Cited on pages 176 and 183.)
- Prajs, S., Sullivan, M., Smith, M., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 464, 3568 (Cited on page 102.)
- Prentice, S. J., Inserra, C., Schulze, S., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 4342 (Cited on page vii.)
- Prentice, S. J., Maguire, K., Flörs, A., et al. 2020, *Astronomy & Astrophysics*, 635, A186 (Cited on pages 175 and 183.)
- Prieto, J. L., Chen, P., Dong, S., et al. 2017, *Research Notes of the American Astronomical Society*, 1, 28 (Cited on page 123.)
- Prochaska, J. X., Hennawi, J., Cooke, R., et al. 2020, *pypeit/PypeIt: Release 1.0.0* (Cited on pages 158 and 196.)
- Prochaska, J. X., Hennawi, J. F., Westfall, K. B., et al. 2020, *Journal of Open Source Software*, 5, 2308 (Cited on pages 158 and 196.)
- Raskin, C. & Kasen, D. 2013, *Astrophysical Journal*, 772, 1 (Cited on page 19.)

- Rau, A., Kulkarni, S. R., Law, N. M., et al. 2009, *Publications of the Astronomical Society of the Pacific*, 121, 1334 (Cited on pages 36, 64 and 66.)
- Ravi, A. P., Rho, J., Park, S., et al. 2023, *Astrophysical Journal*, 950, 14 (Cited on page 87.)
- Rees, M. J. 1988, *Nature*, 333, 523 (Cited on page 11.)
- Reusch, S. 2020, ztfpps (Cited on page 56.)
- Reusch, S. 2023, simeonreusch/fpbot: Release 1.1.3 (Cited on pages 49, 115 and 156.)
- Ricci, C. & Trakhtenbrot, B. 2023, *Nature Astronomy*, 7, 1282 (Cited on pages 11 and 149.)
- Richardson, D., Jenkins, Robert L., I., Wright, J., & Maddox, L. 2014, *Astronomical Journal*, 147, 118 (Cited on page 6.)
- Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, *Astrophysical Journal Letters*, 934, L7 (Cited on page 9.)
- Rigault, M. 2018, ztfquery, a python tool to access ZTF data (Cited on page 62.)
- Rigault, M., Smith, M., Goobar, A., et al. 2025a, *Astronomy & Astrophysics*, 694, A1 (Cited on pages vii, 14, 55, 56, 75, 81, 105, 106, 154 and 156.)
- Rigault, M., Smith, M., Regnault, N., et al. 2025b, *Astronomy & Astrophysics*, 694, A2 (Cited on page viii.)
- Rosswog, S., Kasen, D., Guillochon, J., & Ramirez-Ruiz, E. 2009, *Astrophysical Journal Letters*, 705, L128 (Cited on page 18.)
- Rudolf, N., Günther, H. M., Schneider, P. C., & Schmitt, J. H. M. M. 2016, *Astronomy & Astrophysics*, 585, A113 (Cited on page 47.)
- Ruiz-Lapuente, P., Cappellaro, E., Turatto, M., et al. 1992, *Astrophysical Journal Letters*, 387, L33 (Cited on page 15.)
- Ruiz-Lapuente, P., Damiani, F., Bedin, L., et al. 2018, *Astrophysical Journal*, 862, 124 (Cited on page 17.)
- Ruppin, F., Rigault, M., Ginolin, M., et al. 2025, *Astronomy & Astrophysics*, 694, A6 (Cited on page viii.)
- Saha, A., Thim, F., Tammann, G. A., Reindl, B., & Sandage, A. 2006, *Astrophysical Journal Supplemental Series*, 165, 108 (Cited on page 12.)
- Saito, S., Tanaka, M., Moriya, T. J., et al. 2020, *Astrophysical Journal*, 894, 154 (Cited on page 123.)
- Salmaso, I., Cappellaro, E., Tartaglia, L., et al. 2023, *Astronomy & Astrophysics*, 673, A127 (Cited on pages 124 and 153.)
- Sasdelli, M., Mazzali, P. A., Pian, E., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 445, 711 (Cited on page 15.)
- Schaefer, B. E. & Pagnotta, A. 2012, *Nature*, 481, 164 (Cited on page 17.)
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *Astrophysical Journal*, 500, 525 (Cited on pages 68 and 90.)
- Schulze, S., Gal-Yam, A., Dessart, L., et al. 2024, arXiv e-prints, arXiv:2409.02054 (Cited on page 25.)

- Scolnic, D., Brout, D., Carr, A., et al. 2022, *Astrophysical Journal*, 938, 113 (Cited on page 9.)
- Senzel, R., Maguire, K., Burgaz, U., et al. 2025, *Astronomy & Astrophysics*, 694, A14 (Cited on page viii.)
- Shappee, B. J., Prieto, J. L., Grupe, D., et al. 2014, *Astrophysical Journal*, 788, 48 (Cited on page 36.)
- Shappee, B. J. & Stanek, K. Z. 2011, *Astrophysical Journal*, 733, 124 (Cited on page 64.)
- Shappee, B. J., Stanek, K. Z., Kochanek, C. S., & Garnavich, P. M. 2017, *The Astrophysical Journal*, 841, 48 (Cited on page 65.)
- Sharma, Y., Sollerman, J., Fremling, C., et al. 2023, *Astrophysical Journal*, 948, 52 (Cited on pages vii, 23, 25, 77, 96, 104, 105, 106 and 107.)
- Shen, K. J. & Bildsten, L. 2009, *Astrophysical Journal*, 699, 1365 (Cited on page 18.)
- Shi, D.-W., Wang, S.-Q., Gan, W.-P., & Liang, E.-W. 2024, *Astrophysical Journal*, 969, 32 (Cited on page 123.)
- Siebert, M. R., Kwok, L. A., Johansson, J., et al. 2024, *Astrophysical Journal*, 960, 88 (Cited on page vii.)
- Silverman, J.M., Nugent, P. E., Gal-Yam, A., et al. 2013, *Astrophysical Journal Supplemental Series*, 207, 3 (Cited on page 23.)
- Smartt, S. J., Valenti, S., Fraser, M., et al. 2015, *Astronomy & Astrophysics*, 579, A40 (Cited on pages 36 and 97.)
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *Astronomical Journal*, 146, 32 (Cited on pages 36 and 90.)
- Smith, N. 2017, in *Handbook of Supernovae*, ed. A. W. Alsabti & P. Murdin, 403 (Cited on page 22.)
- Smith, N. & Andrews, J. E. 2020, *Monthly Notices of the Royal Astronomical Society*, 499, 3544 (Cited on page 123.)
- Smith, N., Andrews, J. E., Milne, P., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 405 (Cited on pages 121, 122 and 123.)
- Smith, N., Chornock, R., Li, W., et al. 2008, *Astrophysical Journal*, 686, 467 (Cited on page 22.)
- Smith, N., Mauerhan, J. C., Cenko, S. B., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 449, 1876 (Cited on pages 25, 105, 176 and 183.)
- Soker, N. 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 1333 (Cited on page 19.)
- Solar, M., Michałowski, M. J., Nadolny, J., et al. 2024, *Nature Communications*, 15, 7667 (Cited on page 24.)
- Sollerman, J., Taddia, F., Arcavi, I., et al. 2019, *Astronomy & Astrophysics*, 621, A30 (Cited on pages 124, 125 and 153.)
- Sollerman, J., Yang, S., Schulze, S., et al. 2021, *Astronomy & Astrophysics*, 655, A105 (Cited on page 137.)
- Soraisam, M., Matheson, T., Lee, C.-H., et al. 2022, *Astrophysical Journal Letters*, 926, L11 (Cited on page 125.)
- Stahl, B. E., Zheng, W., de Jaeger, T., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 492, 4325 (Cited on pages 65 and 66.)

- Stanishev, V., Goobar, A., Amanullah, R., et al. 2018, *Astronomy & Astrophysics*, 615, A45 (Cited on page 69.)
- Steehhs, D., Galloway, D. K., Ackley, K., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 511, 2405 (Cited on page 192.)
- Sternberg, A., Gal-Yam, A., Simon, J. D., et al. 2011, *Science*, 333, 856 (Cited on page 23.)
- Stetson, P. B. 1992, , 86, 71 (Cited on page 43.)
- Stritzinger, M., Mazzali, P. A., Sollerman, J., & Benetti, S. 2006, *Astronomy & Astrophysics*, 460, 793 (Cited on page 7.)
- Strotjohann, N. L., Ofek, E. O., Gal-Yam, A., et al. 2021, *Astrophysical Journal*, 907, 99 (Cited on pages 60, 73 and 157.)
- Strubbe, L. E. & Quataert, E. 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 2070 (Cited on page 11.)
- Subrayan, B. M., Milisavljevic, D., Chornock, R., et al. 2023, *Astrophysical Journal Letters*, 948, L19 (Cited on page 11.)
- Sullivan, M., Conley, A., Howell, D. A., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 782 (Cited on page 9.)
- Taam, R. E. 1980, *Astrophysical Journal*, 242, 749 (Cited on page 18.)
- Tartaglia, L., Pastorello, A., Sollerman, J., et al. 2020, *Astronomy & Astrophysics*, 635, A39 (Cited on pages 121, 122 and 123.)
- Taubenberger, S. 2017, *The Extremes of Thermonuclear Supernovae*, ed. A. W. Alsabti & P. Murdin, 317 (Cited on pages 12, 13 and 15.)
- Taubenberger, S., Elias-Rosa, N., Kerzendorf, W. E., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 448, L48 (Cited on page 66.)
- Terwel, J. H., Maguire, K., Dimitriadis, G., et al. 2025, *Astronomy & Astrophysics*, 694, A11 (Cited on pages vii, 27 and 55.)
- Tonry, J., Denneau, L., Heinze, A., et al. 2019, *Transient Name Server Discovery Report*, 2019-2239, 1 (Cited on page 126.)
- Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 130, 064505 (Cited on page 36.)
- Tripp, R. 1998, *Astronomy & Astrophysics*, 331, 815 (Cited on page 9.)
- Tsebrenko, D. & Soker, N. 2015, *Monthly Notices of the Royal Astronomical Society*, 447, 2568 (Cited on pages 21 and 26.)
- Tsvetkov, D. Y., Volkov, I. M., Shugarov, S. Y., et al. 2022, *Contributions of the Astronomical Observatory Skalnat Pleso*, 52, 46 (Cited on page 123.)
- Tucker, M. A., Shappee, B. J., Vallely, P. J., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 1044 (Cited on page 17.)
- Urry, C. M. & Padovani, P. 1995, *Publications of the Astronomical Society of the Pacific*, 107, 803 (Cited on page 10.)
- Utrobin, V. P., Chugai, N. N., Andrews, J. E., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 116 (Cited on page 123.)

- van Dalen, J. N. D., Levan, A. J., Jonker, P. G., et al. 2024, arXiv e-prints, arXiv:2409.19056 (Cited on page [viii](#).)
- van der Walt, S. J., Crellin-Quick, A., & Bloom, J. S. 2019, *Journal of Open Source Software*, 4 (Cited on pages [78](#), [126](#) and [156](#).)
- van Dokkum, P. G. 2001, *Publications of the Astronomical Society of the Pacific*, 113, 1420 (Cited on pages [40](#) and [41](#).)
- Van Dyk, S. D., de Graw, A., Baer-Way, R., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 471 (Cited on page [24](#).)
- van Velzen, S., Gezari, S., Hammerstein, E., et al. 2021, *Astrophysical Journal*, 908, 4 (Cited on page [87](#).)
- Wandel, A., Peterson, B. M., & Malkan, M. A. 1999, *Astrophysical Journal*, 526, 579 (Cited on page [11](#).)
- Wang, B. 2018, *Research in Astronomy and Astrophysics*, 18, 049 (Cited on pages [17](#) and [19](#).)
- Wang, L., Baade, D., & Patat, F. 2007, *Science*, 315, 212 (Cited on page [19](#).)
- Webbink, R. F. 1984, *Astrophysical Journal*, 277, 355 (Cited on page [18](#).)
- Wes McKinney. 2010, in *Proceedings of the 9th Python in Science Conference*, ed. Stéfan van der Walt & Jarrod Millman, 56 – 61 (Cited on page [197](#).)
- Wevers, T., Nicholl, M., Guolo, M., et al. 2022, *Astronomy & Astrophysics*, 666, A6 (Cited on page [vii](#).)
- Wheeler, J. C., Chatzopoulos, E., Vinkó, J., & Tuminello, R. 2017, *Astrophysical Journal Letters*, 851, L14 (Cited on page [123](#).)
- Whelan, J. & Iben, Icko, J. 1973, *Astrophysical Journal*, 186, 1007 (Cited on page [16](#).)
- Wiseman, P., Wang, Y., Hönig, S., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 3992 (Cited on page [11](#).)
- Wiseman, P., Williams, R. D., Arcavi, I., et al. 2025, *Monthly Notices of the Royal Astronomical Society*[arXiv:2406.11552] (Cited on pages [12](#), [139](#), [149](#) and [181](#).)
- Wood-Vasey, W. M., Wang, L., & Aldering, G. 2004, *Astrophysical Journal*, 616, 339 (Cited on page [23](#).)
- Woosley, S. E., Langer, N., & Weaver, T. A. 1995, *Astrophysical Journal*, 448, 315 (Cited on page [6](#).)
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *Astronomical Journal*, 140, 1868 (Cited on pages [36](#), [85](#), [128](#) and [164](#).)
- Yamanaka, M., Kawabata, K. S., Kinugasa, K., et al. 2009, *Astrophysical Journal Letters*, 707, L118 (Cited on page [15](#).)
- Yan, L., Lunnan, R., Perley, D. A., et al. 2017, *Astrophysical Journal*, 848, 6 (Cited on page [10](#).)
- Yang, S., Sollerman, J., Chen, T. W., et al. 2021, *Astronomy & Astrophysics*, 646, A22 (Cited on pages [124](#), [125](#) and [153](#).)
- Yao, Y., Miller, A. A., Kulkarni, S. R., et al. 2019, *Astrophysical Journal*, 886, 152 (Cited on pages [57](#), [117](#) and [157](#).)

- Yaron, O. & Gal-Yam, A. 2012, *Publications of the Astronomical Society of the Pacific*, 124, 668 (Cited on pages 64, 65 and 114.)
- York, D. G., Adelman, J., Anderson, John E., J., et al. 2000, *Astronomical Journal*, 120, 1579 (Cited on pages 36, 90 and 158.)
- Zackay, B. & Ofek, E. O. 2017a, *Astrophysical Journal*, 836, 187 (Cited on pages 31 and 48.)
- Zackay, B. & Ofek, E. O. 2017b, *Astrophysical Journal*, 836, 188 (Cited on pages 31 and 48.)
- Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, *Astrophysical Journal*, 830, 27 (Cited on pages 47, 48 and 56.)
- Zeng, X., Wang, X., Esamdin, A., et al. 2021a, *Astrophysical Journal*, 919, 49 (Cited on page 123.)
- Zeng, X., Wang, X., Esamdin, A., et al. 2021b, *Astrophysical Journal*, 909, 176 (Cited on page 123.)
- Zhang, K., Wang, X., Zhang, J., et al. 2016, *Astrophysical Journal*, 820, 67 (Cited on pages 65, 66, 81 and 82.)
- Zhu, J., Jiang, N., Dong, S., et al. 2023, *Astrophysical Journal*, 949, 23 (Cited on pages 121, 122 and 123.)
- Zimmerman, E. A., Irani, I., Chen, P., et al. 2024, *Nature*, 627, 759 (Cited on page vii.)