

HISTORIC SUPERNOVA LIGHT ECHO
IDENTIFICATION WITH MACHINE
LEARNING

CREATING EFFECTIVE TRAINING SETS FOR MACHINE
LEARNING PACKAGE ALED WITH DRAGONFLY TELEPHOTO
ARRAY IMAGES TO IDENTIFY HISTORIC SUPERNOVA LIGHT
ECHOES AROUND SUPERNOVA 1054 (CRAB)

By NICOLE MULYK, BSc

A Thesis Submitted to the School of Graduate Studies in Partial
Fulfillment of the Requirements for
the Degree Master of Science

McMaster University © Copyright by Nicole Mulyk, August 2024

Abstract

Advances in machine learning for visual recognition and ultra-low surface brightness imaging have made it possible to detect older and fainter historic supernova light echoes (SN LEs). We are particularly interested in the historic core-collapse SN (CCSN) Crab (SN 1054), as it is the only CCSN with records of direct-light observations in the last 1000 years. We have improved the SN LE machine-learning Python package ALED (Automated Light Echo Detection), created by Bhullar et al. (2021), by adding false positive masks as an additional input. ALED is visual recognition software that identifies and locates LEs in difference images. Before the invention of ALED, LE images had to be categorized by visual inspection, which was a very time-consuming task. Additionally, we have developed a method for manufacturing and augmenting LE training sets, which has previously not been applied to LEs. We manufactured Dragonfly Telephoto Array (DTA) LEs by extracting LEs from Canada-France-Hawaii Telescope difference images and overlaying them on DTA difference images. The DTA is a promising tool for LE detection because of its ability to observe ultra-low surface brightness structures. Additionally, we augmented the

only existing DTA LE image by overlaying it on other DTA images. Both of these procedures provided options for further augmentation, such as changing the LE's brightness and width. We also created a process to mask the bright star difference artifacts in DTA images. These stars are typically mislabeled as LEs, and hence masking them makes LE identification simpler. We have created an effective DTA training set for ALED, which is prepared to search for LEs around the historic CCSN Crab (SN 1054), once more DTA images in that region are procured.

Acknowledgements

I would like to start by thanking my supervisor, Dr. Doug Welch, for his constant support throughout this project. Doug - you always knew exactly what to say to give me confidence. Your positivity and passion made me excited to work on my research. I am so grateful that we have had the chance to work together over the past two years. I would also like to thank my committee members: Dr. Alison Sills and Dr. Laura Parker for their encouragement and advice throughout my time at McMaster University. Finally, I would like to thank the SHARCNET, Digital Research Alliance of Canada, and Dragonfly Telephoto Array Teams for supporting my research throughout this process.

I want to give a big thank you to all of my friends and family who supported me throughout my master's degree. Especially my husband, Garon Mulyk, who has pushed me to be the best version of myself. Garon - I am so grateful for all the love and compassion you have given and will continue to give me. I could not have done this without you.

Table of Contents

Abstract	iii
Acknowledgements	v
Abbreviations	xiii
1 Introduction to Supernovae	1
1.1 Types of Supernovae	2
1.2 Historic Supernovae in the Milky Way	17
2 Introduction to Supernova Light Echoes	23
2.1 History	24
2.2 Light Echo Geometry	29
2.3 Light Echoes as Possible Tools for Studying Supernovae	31
2.4 Core-Collapse Supernova Light Echoes	35
3 Light Echo Imaging and Detection	36

3.1	Light Echo Imaging	36
3.2	Light Echo Detection and Machine Learning	41
4	False Positives	50
4.1	False Positive Masks	53
4.2	DECam False Positives	56
5	Manufacturing and Augmenting Light Echo Images	59
5.1	Overlaying CFHT Light Echoes onto Dragonfly Images	60
5.2	Augmenting Existing Light Echoes	62
5.3	Masking Out Bright Stars Difference Artifacts	64
6	Dragonfly Light Echoes	66
6.1	Dragonfly False Positives	69
6.2	Results with Manufactured and Augmented Dragonfly LEs	70
6.3	Results with Masked Bright Stars Diffraction Artifacts	71
6.4	ALED Adaptations for Dragonfly Light Echoes	71
6.5	The Size of Light Echoes Successfully Identified	74
6.6	The Brightness of Light Echoes Successfully Identified	75
7	Discussion and Conclusion	79
A	Versions of ALED	83
B	Fractions of Correctly Identified False Positive, Light Echo, and Background Images in Versions of ALED with Dragonfly Training Sets	90

**C Fractions of Correctly Identified Light Echo Images using DF5 and
DF10 Weights with Test Sets that have Various Brightness Levels 102**

List of Figures

1.1	SNe Classification	3
1.2	SNe Spectra	8
1.3	SNe Type Ia Spectra	9
1.4	Hubble Diagram	11
1.5	SNe Light Curves	17
2.1	Tycho and Cas A LE Apparent Motion Vectors	26
2.2	SN 1987A LEs	28
2.3	LE Geometric Diagram	30
2.4	Multiple LEs from SN 1987A	31
2.5	Image Differencing and LE Apparent Proper Motion Vectors	33
3.1	The Dragonfly Telephoto Array	39
3.2	ALED Input Image Cropping	45
3.3	The Architecture of ALED	46
3.4	CFHT LE Difference Image and RPV	48
4.1	CFHT False Positive and Masks with Varying Levels of Detail	51

4.2	Workflow Diagram	54
4.3	DECam FP and RPV	56
5.1	Real DTA LE	60
5.2	CFHT LE Overlaid on a DTA Image	63
6.1	DTA Bright Star Diffraction Artifacts FPs and RPV	70

List of Tables

2.1	Alternative LE Sources	29
4.1	Original ALED Training Set	50
4.2	Variations on Versions 7 and 12	55
4.3	Average Failure Rates, Fractions of Correctly Identified FPs and LEs for CFHT Versions	58
5.1	Differences between DTA and CFHT Images	61
5.2	Average Fraction of LE Pixels in Cropped Images	61
6.1	Cropped Image Size for DTA Versions	72
6.2	Pixel Threshold <i>num_ones</i> for DTA Versions	73
6.3	Fraction of Correctly Identified LEs for DTA LEs of Different Sizes	75
6.4	Fractions of Correctly Identified LEs with Test Sets that have the LE Brightness Varied	76
6.5	Average Failure Rates, Fractions of Correctly Identified FP, LE, and BG images with a Probability Threshold of > 0.00042	77

6.6	Average Failure Rates, Fractions of Correctly Identified FP, LE, and BG images with a Probability Threshold of > 0.001	78
-----	---	----

Abbreviations

ALED	Automated Light Echo Detection
BG	Background
BSDA	Bright Star Difference Artifacts
CapsNet	Capsule Neural Network
Cas A	Cassiopeia A
CNN	Convolution Neural Network
CCSN	Core-Collapse Supernova
DECam	Dark Energy Camera
DTA	Dragonfly Telephoto Array
CFHT	Canada France Hawaii Telescope
FP	False Positive

LE	Light echo
LMC	Large Magellanic Cloud
LOS	Line-of-sight
RPV	Routing Path Visualization
SN	Supernova (singular)
SNe	Supernovae (plural)
SNR	Supernova Remnant
WD	White Dwarf

Chapter 1

Introduction to Supernovae

A supernova (SN) is a fantastically energetic event that expels the majority of the star's material (Leibundgut, 2008). They can briefly reach 100 million times the luminosity of the Sun, outshining the combined light of all other stars in their galaxy (Stephenson & Green (2002); Leibundgut (2008)). These brief, catastrophic events are one of the final stages of a star's life, either as it collapses into a compact object or as its dense core remnant is disrupted. The resulting SN remnant (SNR) can be observed through electromagnetic radiation for several thousand years (Stephenson & Green, 2002).

The interest in supernovae (SNe) goes beyond their internal processes - they can trigger star formation and influence galactic chemical evolution, as they are a dominant source of heavy elements recycled into the interstellar medium (Filippenko, 1997). They are involved in compact object production, emit cosmic rays, and are

used in the calibration of the cosmic distance scale (Filippenko (1997); Branch & Tammann (1992); Riess et al. (1998)).

1.1 Types of Supernovae

Minkowski (1941) identified two classifications of SNe by observing the optical spectra of fourteen supernovae in IC 4182, NGC 4636, and NGC 4725. SNe are sorted into these two groups by the nonexistence (Type I) or existence (Type II) of H absorption lines in their optical spectra (Filippenko, 1997), see Figure 1.1. The Type I SNe are further divided into those with singly ionized silicon lines (Type Ia), those without silicon lines but with He lines (Type Ib), and those without silicon or He lines (Type Ic) (Leibundgut, 2008). Further investigation showed that the Type I and II classification system had little physical relevance. A clear distinction between Type Ia and Type Ib/c categories, both in luminosity and emission, indicated that they were produced by different processes (Weiler & Sramek, 1988). Type Ia SNe have a more luminous peak by $\sim 1.5 - 2$ magnitudes and they do not emit in radio frequencies, unlike Type Ib/c SNe (Weiler & Sramek (1988); Weiler et al. (2002)). Furthermore, observations of SNe, such as SN 1993J, that have H lines in early stages and lose them in later stages, show a possible transition from Type II to Type Ib/c SNe (Leibundgut, 2008). This metamorphosing class of SNe has been labelled Type Iib and suggests a connection between Type Ib/c and Type II SNe. Therefore, it was finally concluded that two physical classifications can be made: Thermonuclear SNe (Type Ia) and Core-collapse SNe (Type Ib, Ic, and II). Zwicky (1965) proposed the addition of Type III, IV, and V classifications after observing a few unusual SNe. However, Doggett & Branch (1985) argued that since all of their spectra had H lines,

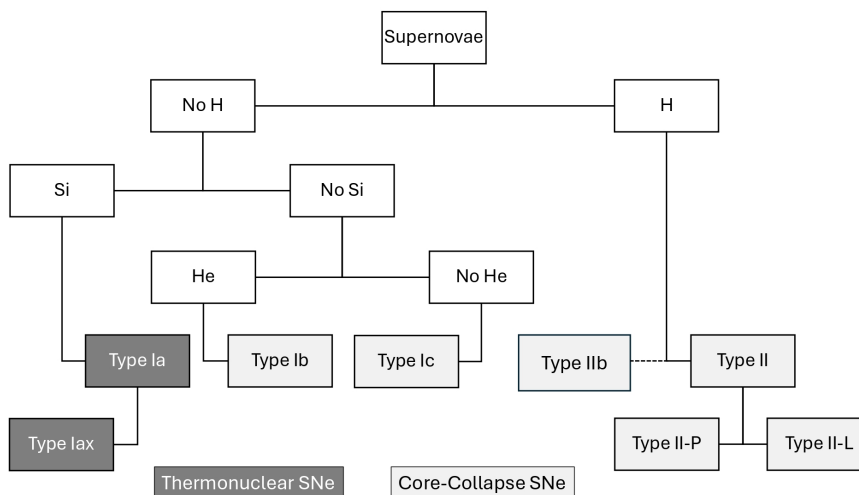


Figure 1.1: Supernova classification based on the existence of certain absorption lines. Adapted from Leibundgut (2008).

these SNe should be classified as Type II SNe.

Clearly, the spectrum of the SN is required for classification. However, for the historic SNe in Section 1.2, the direct light was observed several hundred years ago - before the advent of astronomical spectroscopy (Rest et al., 2008a). Therefore, modern, alternative methods are required to reconstruct the spectra of historic SNe, see Section 2.3.2.

1.1.1 Thermonuclear Supernovae

Thermonuclear SNe, or more commonly Type Ia SNe, do not have H or He lines in their spectra because their progenitor stars have shed those outer layers, leaving only their compact core, a white dwarf (WD) (Leibundgut, 2008), refer to Figure 1.1. If the WD becomes more massive than its internal pressure can support, the WD's collapse launches a thermonuclear explosion, which disrupts the object. These SNe are extremely luminous, with an absolute magnitude of -19.4 ± 0.5 and can

produce $\sim 10^{51}$ ergs of energy (Branch & Tammann (1992); Maoz et al. (2014)). They are known for having very homogeneous light curves which are easily calibrated for measuring cosmological distances (Leibundgut, 2008).

Thermonuclear Supernovae Progenitors

As thermonuclear SNe do not have any H lines, the progenitor must have $< 0.1M_{\odot}$ of H in its atmosphere (Hillebrandt & Niemeyer, 2000). However, a deeper investigation of the spectra leads to a compact progenitor with a mass of $\sim 1M_{\odot}$ and a radius $< 10^4$ km, composed mainly of C and O; likely a C-O WD (Hillebrandt & Niemeyer, 2000).

Additionally, Type Ia SNe have been observed in many types of galaxies and different locations within those galaxies (Filippenko, 1997). Elliptical galaxies typically do not contain massive stars (Hillebrandt & Niemeyer, 2000). Therefore, observations of Type Ia SNe in elliptical galaxies support WDs as progenitors, as they form from low-mass stars. However, it cannot be ruled out that Type Ia SNe that are observed in spiral galaxies have a different class of progenitors. For example, Type Ia SNe with bright peaks and slow declines are not found in elliptical galaxies and hence could be a separate class. Hillebrandt & Niemeyer (2000) states that different explosion mechanisms and white dwarf progenitors could account for the minor variations in the light curves.

Aside from minor variations, Type Ia SNe have uniform light curves, which is further evidence for only one class of progenitors (Hillebrandt & Niemeyer, 2000). Chandrasekhar-mass WDs are a likely candidate. WDs are supported by electron-degeneracy pressure, therefore, there is a maximum mass (Chandrasekhar mass limit)

that can be endured before they succumb to gravitational pressure and collapse (Leibundgut, 2008). Therefore, all progenitors would have approximately the same mass before erupting, resulting in a similar light curve (Hillebrandt & Niemeyer, 2000).

Currently, a C-O WD has been accepted as the probable progenitor of thermonuclear SNe (Hillebrandt & Niemeyer, 2000). As C-O WDs typically form with a mass of $\sim 0.6 M_{\odot}$, a companion is required to increase the WDs mass to the Chandrasekhar limit ($\sim 1.4 M_{\odot}$), which is required to launch the explosion. Two scenarios have been proposed to achieve this: a single-degenerate binary or a double-degenerate binary.

1. **Single-Degenerate Scenario:**

In this scenario, a WD is in a binary with a companion (Hillebrandt & Niemeyer, 2000). The companion has overflowed its Roche lobe and the WD is accreting material onto its surface until it reaches critical mass. Although this model has been generally accepted, there are a few problems that are discussed in the literature, mainly regarding the accretion rate. Various ranges of accretion rates could result in smaller (nova) explosions. If more mass is ejected than is being accreted, then the critical mass will not be reached. The companion could also envelope the WD if the accretion rate is high enough. However, a better understanding of WD winds and required accretion rates that result in H burning without novae could potentially dismiss these issues (Hillebrandt & Niemeyer, 2000).

2. **Double-Degenerate Scenario:**

Double-degenerate models consist of two C-O WDs in a binary, which merge to launch the supernova (Hillebrandt & Niemeyer (2000); Pakmor et al. (2012)).

This model would explain the lack of H but also has a few complications (Hillebrandt & Niemeyer, 2000). First, there is little observational evidence for WD binaries that are massive enough to reach the Chandrasekhar limit. Second, it is difficult to explain how merging WDs would result in a consistent light curve, as the progenitor WDs could have varying mass and composition. Third, merger models suggest that gravitational collapse is more likely than thermonuclear disruption due to off-center ignition.

Several studies of Type Ia SN remnants (Badenes et al., 2007), rates (Ruiter et al., 2009), delay times (Maoz et al., 2010), and emission (Hancock et al. (2011); Leonard (2007); Gilfanov & Ákos Bogdán (2010)) lead to the double degenerate scenario being the favoured model (Pakmor et al., 2012). However, no definitive evidence has yet been found.

The Thermonuclear Supernovae Explosion Mechanism

The explosion mechanism for the single-degenerate model has historically been more thoroughly researched (Pakmor et al., 2012). In this scenario, the companion star is accreting $4 \times 10^{-6} - 10^{-8} M_{\odot}$ of material per year onto the C-O WD (Branch & Nomoto, 1986). Through H and He flashes, the accreted material is converted to C and O. As the mass of the WD approaches the Chandrasekhar limit, the internal temperature and pressure increase in the core, eventually igniting the C core. This event launches a subsonic (slower than the speed of sound) front of nuclear burning outward from the core, which converts the C and O into Ni-56 and other isotopes (Branch & Nomoto (1986); Leibundgut (2008)). This slow-moving front, or deflagration, transitions to detonation, a supersonic burning front (Leibundgut, 2008). These

runaway thermonuclear reactions result in the complete detonation of the star, releasing approximately $\sim 10^{51}$ ergs of kinetic energy, with no dense remnant left behind (Branch & Nomoto, 1986).

Recently, high-resolution simulations have made it possible to explore double-degenerate merger models (Pakmor et al., 2012). It was initially thought that the larger of the two WDs could tidally disrupt the smaller WD companion and accrete its remains (Maoz et al., 2014). However, more recent models predict violent, head-on collisions, and He WD accretion. Violent mergers are more likely to occur in dense environments, as the probability of collision increases (Maoz et al., 2014). These mergers could result in a Type Ia SN, even if the total mass of the binary is below the Chandrasekhar limit. A head-on collision between two WDs is most likely to occur in a dense environment or a tertiary system. A system with three WDs, or a hierarchical triple, typically has two WDs in a stable binary and a low-mass WD with an eccentric outer orbit. However, as only 10-20% of all stars can be found in tertiary systems, it is unlikely that there are enough tertiary WD systems to produce the observed number of Type Ia SNe. Furthermore, such a system would need to form WDs before merging. Lastly, similar to the single-degenerate scenario, a He WD can accrete He onto its degenerate companion. The accreted He forms an outer layer on the WD, which can ignite He burning and launch the SN.

Thermonuclear Supernovae Light Curves and Spectra

Thermonuclear SNe emissions have been primarily observed in optical and near-infrared wavelengths (Leibundgut, 2008). During the outburst peak, singly ionized Si, Ca, Mg, O, and S lines can be found in the optical spectra, along with some Fe and

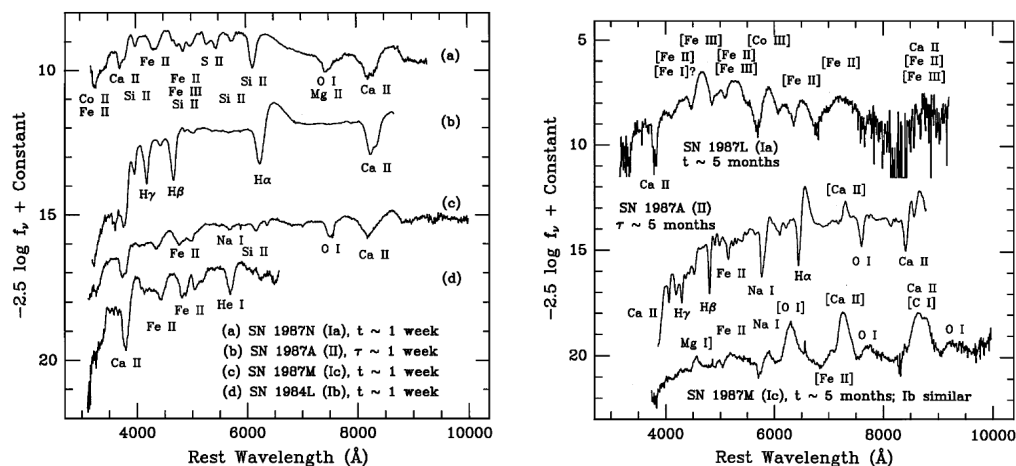


Figure 1.2: The spectra of a Type Ia (SN 1987N & SN 1987L), Ib (SN 1984L), Ic (SN 1987M), and II (SN 1987A) SNe. The left plot shows the early-time spectra (one week after the optical peak) and the right plot shows the late-time spectra (five months after the optical peak). Filippenko (1997), Figures 1 & 2.

Co lines at shorter wavelengths (Hillebrandt & Niemeyer (2000); Filippenko (1997)), as can be seen in Figure 1.2 (left). These early emission lines are quite broad, due to fast ejecta velocities (Filippenko, 1997). The emission is powered by the decay of Ni-56 into Co-56 and Fe-56 (Leibundgut, 2008). A $1.4 M_{\odot}$ (Chandrasekhar mass) WD is expected to produce $\sim 0.6 M_{\odot}$ of Ni-56 during C and O burning (Maoz et al., 2014). The decay of this uniform amount of Ni-56 results in the standardized post-outburst light curve that is observed for Type Ia SNe (Weiler & Sramek, 1988). In the weeks following the peak, permitted Fe II lines begin to dominate, followed by forbidden Fe II, Fe III, and Co III lines during the tail of the light curve (Leibundgut, 2008), see Figure 1.2 (right). After the ejecta has cooled, thermal infrared emission is more predominant than optical and near-infrared.

As mentioned previously, Type Ia SNe have very homogeneous light curves and spectra, as shown in Figure 1.3. Hillebrandt & Niemeyer (2000) notes that 85% of




Figure 1.3: The spectra of three Type Ia SNe (SN 1990N, SN 1987N, and SN 1987D) one week after the optical peak. All three SNe show similar peaks and troughs. Filippenko (1997), Figure 6.

thermonuclear SNe have a standard light curve and spectrum, and show 0.2-0.3 magnitudes of dispersion in their peak luminosity. Although Type Ia SNe are considered to be homogeneous, some variations and correlations do exist. In general, explosions with lower luminosity peaks also decline from the peak more rapidly, are redder, and have lower ejection velocities (Branch, 1998). There is a correlation in the light curves between the width and height of the peak (Hillebrandt & Niemeyer, 2000). Furthermore, there is a relation between the SNe and their host galaxies. Elliptical galaxies host Type Ia SNe that are dimmer (by 0.2-0.3 magnitudes), with more rapidly evolving light curves and slower ejecta. However, similar dimmer Type Ia SNe have been found in the outer parts of spiral galaxies.

Thermonuclear Supernovae as Standard Candles

Type Ia SNe have a characteristic light-curve shape and are extremely luminous, with an absolute magnitude of -19.4 ± 0.5 (Branch & Tammann (1992); Stephenson & Green (2002)). Both of these factors make Type Ia SNe helpful standard candles for measuring cosmological distances, up to $z \approx 1$, making them useful tools for studying the expansion of the universe (Stephenson & Green (2002); Krause et al. (2008b); Hillebrandt & Niemeyer (2000); Leibundgut (2008)). Standard candles are useful for determining distances because their absolute magnitude M is known (Leibundgut, 2001). The distance modulus ($m - M$), see Equation 1.1.1, assumes the universe expands linearly, hence only holds true in the local universe.

$$m - M = 5 \log z + 5 \log \frac{c}{H_0} + 25 \quad (1.1.1)$$

Due to the homogeneity of their light curves, Type Ia SNe are useful for finding the Hubble constant H_0 (Hillebrandt & Niemeyer, 2000). The Hubble diagram, see Figure 1.4, is a plot of the distance modulus (Equation 1.1.1) or normalized distance modulus versus redshift (Leibundgut, 2008). By plotting the Type Ia SNe data, the Hubble constant can be estimated. It is assumed that the local expansion is linear and all Type Ia SNe have the same absolute magnitude (i.e. standard candles). Type Ia SNe are good (especially in infrared), but not perfect, standard candles. In practice, each distance modulus must be determined individually, by normalizing the peak luminosity with the shape of the light curve (Leibundgut (2001); Leibundgut (2008)).

Additionally, Type Ia standard candles have been used for estimating the density parameters of matter Ω_M and dark energy Ω_Λ (Hillebrandt & Niemeyer, 2000), and




Figure 1.4: Hubble diagram constructed with Type Ia SNe. The distance modulus (top panel) or the normalized distance modulus (middle and lower panels) are plotted against redshift. In addition to the SNe, four cosmological models are also plotted: the empty universe (solid), the matter-dominated model (long-dashed), the vacuum-dominated model (dashed), and the flat universe (dotted). The middle and lower panels show the High- z SN Search Team data and the Supernova Cosmology Project data, respectively. Leibundgut (2001), Figure 1.

therefore are useful for determining the equation of state of the universe (Hillebrandt & Niemeyer, 2000). Type Ia SNe were fainter than was expected as a function of z , implying the existence of “dark energy”, which has negative pressure (Leibundgut, 2001). As shown in Figure 1.4, Type Ia SNe measurements are consistent with a flat universe that contains matter and dark energy ($\Omega_M = 0.3$ and $\Omega_\Lambda = 0.7$).

Riess et al. (1998) performed a cosmological study with 16 high-redshift Type Ia SNe and 34 local SNe. They measured the distance modulus of these SNe to constrain cosmological parameters, such as H_0 , Ω_M and Ω_Λ . They determined that

the universe had a dynamical age of 14.2 ± 1.7 Gyr and the expansion of the universe is accelerating.

1.1.2 Core-Collapse Supernovae

Core-collapse SNe (CCSNe) comprise of all other classifications (Type Ib, Ic, and II) (Leibundgut, 2008). All CCSNe have one common progenitor, a massive star ($> 7 M_{\odot}$) (Smartt, 2009). However, the various possible combinations of absorption lines observed are dependent on the layers of the star present before collapse (Leibundgut, 2008). For example, Type II SNe still have a H layer prior to collapse, whereas Type Ib SNe have already shed this layer through stellar winds or binary interactions. A Type Ic SN occurs if the He layer has also been shed. Type IIb SNe have been proposed as a transitional class, initially showing a Type II typical spectra (with H lines) and transforming into a Type Ib/c spectra later. In some of the literature, Type II SNe have been divided into Type II-P and Type II-L depending on whether there is a plateau in the light curve directly after the peak, see Figure 1.5 (Smartt, 2009).

CCSNe occur when nuclear burning in the core of the massive star ends (Smartt, 2009). Without outward thermal pressure to support the star, the star collapses. The outer layers rebound off the dense core, launching massive shock waves that blow away the outer layers of the star (Leibundgut, 2008). In general, CCSNe are far less homogeneous than Type Ia SNe, showing variation in their spectra, light curves, peak luminosity, and evolution (Weiler & Sramek, 1988). CCSNe have an average magnitude of -18.0 ± 0.8 , about 2 magnitudes fainter than Type Ia. Additionally, CCSNe, unlike Type Ia SNe, also emit in radio.

Core-Collapse Supernovae Progenitors

Massive stars ($> 7 M_{\odot}$) have been accepted as the progenitors of CCSNe (Smartt, 2009). However, given the variations in CCSNe spectra, further investigation into each sub-type's progenitor is required. The CCSNe progenitor stars must be massive so that the core pressure and temperature are high enough to produce Fe or O-Mg-Ne cores. The location where CCSNe are detected may also reveal information about the progenitor. Unlike Type Ia SNe, CCSNe are preferentially found in spiral arms and HII regions, where young stars have recently formed (Filippenko, 1997). This is especially true for Type Ic SNe, which are frequently found in HII regions. HII regions are associated with young stars that are ionizing the regions around them. Such an environment suggests a massive progenitor, as massive stars evolve quickly and die young. Type Ib SNe have been found in HII regions, although not as exclusively as Type Ic SNe. Type II SNe do not have a preference toward HII regions. Hence, Type Ic SNe have the youngest progenitors, followed by Type Ib and Type II, in order of increasing age of progenitor.

Type Ibc SNe differ from Type II SNe by the absence of H lines in the ejecta. Therefore, the progenitor of Type Ibc SNe likely lost its H envelope prior to collapse (Smartt, 2009). One possible progenitor class is Wolf-Rayet stars, which shed their H envelope during earlier stages of their evolution. Another potential progenitor class is a massive star in a binary with a low-mass star. The H envelope of the massive star is accreted onto the low-mass companion, through Roche lobe overflow, before it erupts. The literature favours interacting binaries as Type Ibc progenitors because such binary systems are common in the Milky Way. Too few Wolf-Rayet stars have been observed to agree with the rate of Type Ibc SNe. Furthermore, Wolf-Rayet stars

are massive enough to collapse into black holes, which are unlikely to eject enough material to produce bright electromagnetic emission.

Type II-P SNe are classified by the presence of a plateau directly after the peak of the light curve, refer to Figure 1.5 (Smartt, 2009). These SNe are the most frequently observed core-collapse SNe and therefore the best understood. Red supergiants have been accepted as the progenitors of Type II-P SNe. Such red supergiants form as massive, $8 - 30M_{\odot}$, stars expand and cool when they have begun He burning in their core. Smartt (2009) used the Salpeter Initial Mass Function, with a slope of $\alpha = -2.35$, to predict that Type II-P have progenitors with masses between $8.5 - 16.5 \pm 1.5 M_{\odot}$. No Type II-P red supergiant progenitors have been detected with masses between $17 - 30 M_{\odot}$. This is the “Red Supergiant Problem”. These massive stars may produce other types of CCSNe or extremely faint explosions. Additionally, the assumed α value, red supergiant metallicities, or other assumptions could be incorrect. Despite such ambiguities, these progenitors have been observed before the SNe launches, making them the most thoroughly researched SNe progenitors.

Type II-L SNe do not have a plateau following the peak of the light curve. Compared to other types of SNe, Type II-L SNe are extremely rare. Therefore, little is known about their progenitors and whether they differ from Type II-P.

Core-Collapse Supernovae Explosion Mechanism

While fusion is ongoing in the core of a massive star, hydrostatic equilibrium is maintained (Leibundgut, 2008). But after iron peak elements are produced, exothermic reactions cease and there is no longer thermal pressure to support the star against collapse (Leibundgut (2008); Smartt (2009)). In the core, electrons and protons merge

into neutrons, as densities increase (Leibundgut, 2001). The inward falling material bounces off the proto-neutron star, causing outward shock waves (Leibundgut (2008); Stephenson & Green (2002)). However, these shock waves do not have sufficient energy to breach the surface of the star (Smartt, 2009). Escaping neutrinos carry away a large fraction of the available energy, 3×10^{53} ergs, and deposit some energy in the outer layers, which assists with launching the SNe (Stephenson & Green (2002); Smartt (2009)). Ni-56 is produced as the shock wave heats the outer layers, initiating Si and O explosive burning. These explosions have immense kinetic energy, $\sim 10^{51}$ ergs (Smartt, 2009). Once the outer layers of the star are blown away, only the neutron-dense core remains (Stephenson & Green, 2002).

Core-Collapse Supernovae Light Curves and Spectra

Historically, there has been significantly more interest in Type Ia SNe light curves, due to their use as standard candles in the study of cosmology. As a consequence, CCSNe light curves have been followed less completely after discovery. However, more recent work has found that CCSNe have great diversity in their light curves.

Type Ib SNe are set apart from Type Ic SNe by the presence of He in their spectra. These He I lines become easy to identify in the weeks following peak brightness (Filippenko, 1997), see Figure 1.2. Unlike Type Ib SNe, Type Ic SNe do not show H or strong He lines.

At the early stages of Type II-P SNe, they have high temperatures ($\geq 10^4$ K) and blue spectra (Filippenko, 1997). At later times, H α lines are strong during the tail of the light curve. Type II-L SNe also have a blue near-featureless peak, with faint H α lines. The H α lines become more dominant as time passes. However, Filippenko

(1997) notes that $H\alpha$ absorption features have not been observed for Type II-L SNe. This could be an identifying feature for Type II-L SNe, however too few have been observed to be conclusive.

Type II SNe are divided into two classes by features in their light curves, as shown in Figure 1.5. As previously mentioned, Type II-P SNe have plateaus directly following their peak (Woosley & Weaver, 1986). Type II-L SNe have a near-linear decline in their light curve after the peak, followed by a change to a more gradual linear slope after 100 days. The plateau in Type II-P SNe light curves is thought to be caused by a rapid injection of energy into the H envelope (Smartt, 2009). Whereas, Type II-L SNe have very low-mass H envelopes, hence do not have a plateau. During the plateau phase of a Type II-P SNe, H recombination occurs as the photosphere cools (Filippenko (1997); Smartt (2009)). Following this, the tail of the light curve is powered by the decay of Co-56 to Fe-56 (Smartt, 2009). Depending on the quantity of Ni-56 ejected, and correspondingly the amount of Co-56, there can be variation in the luminosity of the tail of the light curve.

Observations of the direct light from SNe contain emission along one line-of-sight (Rest et al., 2011b). Therefore, it is not possible to study the asymmetry of SN explosions from the direct light. However, there are other techniques to study asymmetry, explained in Section 2.3.3. Theoretical evidence suggests that CCSNe emit asymmetrically (Wang & Wheeler, 2008). As discussed in more detail in later sections, this was verified using light echo spectra by Rest et al. (2011b) and Sinnott et al. (2013). They discovered that two CCSNe, Cas A (Rest et al., 2011b) and SN 1987A (Sinnott et al., 2013), had asymmetric explosions.




Figure 1.5: The light curves of Type Ia, Ib (and Ic), II-P, and II-L SNe. SN 1987A has also been included. Filippenko (1997), Figure 3.

1.2 Historic Supernovae in the Milky Way

The appearance of a “new star”, or “stellae novae” in Latin, would have been an intriguing discovery for ancient astronomers (Leibundgut, 2008). These SNe were comparable to, or brighter than the most luminous stars in the sky (Stephenson & Green, 2002), and some were even visible during the day. The early pre-telescopic observations are important to modern astronomers because they cover a longer timescale than modern observations and are therefore more likely to capture rare local events, such as supernovae (Stephenson & Green, 2002). Although these observations have lower precision than what is attainable with modern instrumentation, they often include useful information about the brightness, location, and duration of the supernovae. Information about the change in brightness, whether it was smooth or rapid, can help modern astronomers approximate the light curve and determine whether a

historic event was a SN or a nova. Additionally, proximity to a SNR can also indicate whether it was a SNe. Ancient astronomers often compared the guest star with nearby stars and planets, making it possible for modern astronomers to estimate their apparent magnitude.

Although thousands of extragalactic SNe have been detected, few Milky Way SNe have been observed (Stephenson & Green, 2002). The predicted SN rate in the Milky Way is 1-2 per century. However, throughout history, only five galactic supernovae have reliable observational records: SN 1006, 1054, 1181, 1572, and 1604, where the SN identification corresponds to the year in which it was observed. Additionally, there are a few potential SNe sightings recorded in ancient texts prior to 1000 AD. Although the SN was not directly observed, evidence of a SNR called Cassiopeia A (Cas A), would suggest that an additional galactic SN occurred around 1680 AD. Each of these events is discussed in Sections 1.2.1 - 1.2.6. There are many other (200+) SNRs in the Milky Way observable in radio wavelengths, similar to Cas A, that have no recorded direct detections.

Ancient astronomers across Asia, Europe, and northern Africa recorded observations of new stars in historic texts. Stephenson & Green (2002) comments that the historic SNe records were mainly found in the northern hemisphere. While there is a bias for recorded SN observations in the northern hemisphere, SNe were likely observed in the southern hemisphere as well. However, these observations were probably not recorded or the records were destroyed. Chinese astronomers has been recording celestial events since 700 BC (Stephenson & Green, 2002). Detailed records dating back to 200 BC include the appearance of several new stars. Following 1000 AD, several SNe recordings were made in Japan and Korea, including SNe 1006, 1054,

1181, and 1604. SN 1006 was also observed by several Arab and European historians. There was a notable shift to more detailed astronomical records during the Renaissance period in Europe. Astronomers, such as Tycho Brahe and Johannes Kepler, recorded SN 1572 and 1604, along with many other astronomical events. Their research showed that these “new stars” appeared at distances far beyond the planets. The idea that distant objects could evolve was revolutionary, as it was thought that distant stars were “fixed” in the sky.

1.2.1 SN 1006

SN 1006 was observed for several years following its appearance in 1006 AD (Stephenson & Green, 2002). Multiple records from Europe, the Middle East, northern Africa, and Asia noted this exceptionally bright “new star”. Using these historic records, it was possible to find the approximate coordinates of the SN at RA 15.23^h and Dec $-43.1^\circ \pm 1^\circ$ in the constellation Lupus, identify the SNR as G327-6+14-6, and classify it as a Type Ia SNe (Stephenson & Green (2002); Branch & Tammann (1992)).

1.2.2 SN 1054 (Crab)

SN 1054 is well-known for leaving behind the Crab Nebula (Messier 1) (Stephenson & Green, 2002). The Crab Nebula is optically very bright and is therefore the earliest SNR to be discovered, in 1731 AD. Hubble (1928) and Lundmark (1921) were the first to suggest that the Crab Nebula was a remnant of the “guest star” recorded in China and Japan in 1054 AD. These ancient records included the duration and location of SN 1054, making it possible to verify the connection between SN 1054 and the Crab Nebula (Stephenson & Green, 2002). The Japanese and Chinese records suggest that

it was visible for nearly two years. Additionally, SN 1054 would have been bright enough to see during daylight for ~ 23 days, with a magnitude of ~ 3 . Chevalier (1977) classified SN 1054 as a CCSNe. It is important to note that SN 1054 is the only historic CCSNe with reliable observations in the last 1000 years.

1.2.3 SN 1181

SN 1181 was observed for six months in China and Japan in 1181-1182 AD (Stephenson & Green, 2002). Although there is no evidence that it was visible during the day, reports of a large guest star would indicate that it was very bright, with a luminosity of at least magnitude 0 during its peak. Its estimated location is $2^h 5^m$ and $+64.5^\circ$, based on the relative position of nearby stars. Ritter et al. (2021) provides evidence that it is a Type Iax SN, formed in a double-degenerate merger. Type Iax is a recently added SN sub-classification that includes Type Ia SNe with slower expansion velocities ($2 - 7 \times 10^4 \text{ km s}^{-1}$ rather than the typical 10^5 km s^{-1}), see Figure 1.1. Ko et al. (2024) has confirmed that IRAS 00500+6713 is the SNR of SN 1181.

1.2.4 SN 1572 (Tycho)

Tycho Brahe's SN 1572 is a well-known historic example of a thermonuclear SN (Krause et al., 2008b). Records from Europe, China, and Korea indicate that SN 1572 was located in the constellation Cassiopeia and was visible during the day (Stephenson & Green, 2002). It was observed in Korea and China for 18 months, but no records were found in Japan during this time. SN 1572 also had a luminosity of magnitude -4, which is comparable to Venus (-4.5 during its peak). Tycho Brahe's measurements made it possible to firmly determine the SNR as G120-1+1-4 and approximate the

light curve.

1.2.5 SN 1604 (Kepler)

Johannes Kepler observed SN 1604 extensively, and so his name is associated with it (Stephenson & Green, 2002). Observers of this SN in Korea and Europe commented on how it appeared to “twinkle”, as a result of its low altitude in the sky. The Korean and Chinese records span six months. Both Kepler and David Fabricius’ positional observations were incredibly precise, indicating that SN 1604 was positioned at RA $17^h7.1^m$ and Dec $-21^\circ5'$. The European measurements included several dated brightness comparisons with nearby stars and planets. Although it was not visible during the day, at its peak SN 1604 would have had a luminosity of about magnitude -3.0, which is brighter than Jupiter (-2.0 at the time of observations). A light curve has been estimated using the European and Korean measurements. These detailed observations have made it possible to confidently classify it as a Type Ia SN (Branch & Tammann, 1992).

1.2.6 Cas A

Cassiopeia A (Cas A) is a young SNR located in the constellation of Cassiopeia (Stephenson & Green, 2002). Unlike the other SNe previously discussed, the SN that created Cas A was not directly observed. Cas A was first discovered as a bright radio source in the late 1940s. There are visible circular filaments surrounding the central radio source. The uniform expansion of these filaments confirms that it is the remnant of a SN that erupted in 1671.3 ± 0.9 . The SN should have been observable in Europe and Asia, and circumpolar in the northern latitudes. Therefore, the lack of

reliable observational evidence is unexpected. However, if the progenitor had shed its H envelope, then the resulting SN may have been optically subluminal, accounting for the lack of its detection (Woosley & Weaver, 1986). The Cas A SNe was classified as a Type IIb SN by Krause et al. (2008a) using light echoes.

1.2.7 Potential SN Records prior to 1000 AD

Early records in 70, 185, 369, 386, 393, and 837 AD include potential SNe sightings (Stephenson & Green, 2002). Little is known about the guest star observed in 70 AD. It is thought to have been a nova, but, as the position is poorly known, that conclusion is not definitive. The guest star from 185 AD has been convincingly identified as a SN, due to its long duration (between 8-20 months) and many potential nearby SNRs. SN 185 is often categorized as a potential Type Ia SN in the literature (Branch & Tammann, 1992). It was not possible to determine whether the guest star in 369 AD was a SN because the position of the star is not accurately known (Stephenson & Green, 2002). Despite having a very short duration (~ 3 months), the guest star from 386 AD is potentially a SN, with several potentially associated SNRs. Of the three guest stars from this century, the guest star of 393 AD is the most promising SN candidate, due to its long duration (~ 8 months) and proximity to several potential SNRs. Two individual guest stars were observed in 837 AD. Both stars were determined to have been novae as they both had a short duration (22 days and 75 days).

Chapter 2

Introduction to Supernova Light Echoes

Light echoes (LEs) occur when light from a transient source scatters off the surrounding interstellar or circumstellar dust (Sugerman, 2003). As the light must travel first to the dust and then to the observer, the scattered light takes a longer path and reaches the observer after the direct light (Rest et al., 2015). Therefore, LEs are potentially observable centuries after the direct light (Rest et al., 2011a). One common source of LEs is a SN, hence LEs are useful for studying historic SNe (Rest et al., 2005). Without LEs, there is only a brief period when the direct light is visible from Earth, during which spectra, light curves, and other information about the SN can be recorded (Rest et al., 2012b). Therefore, light echoes have the potential to greatly expand the data collection period for temporary astronomical events. LEs have been found near a few of the historic Milky Way SNe introduced in Section 1.2, however, there are several remaining historic Milky Way SNe without LE detections.

LEs can provide information about the distance, spectra, classification, and asymmetry of SNe, as well as the structure of the surrounding dust. When combined with observations of the remnant, LEs may provide an opportunity to observe both light from the event and the aftermath simultaneously. LEs should not be confused with so-called infrared echoes. Such features, like ones discovered by Krause et al. (2005), are caused by dust absorbing light in the optical part of the spectrum and re-radiating the energy as thermal infrared radiation (Rest et al., 2012b).

2.1 History

The first LEs were observed by Ritchey (1901) around Nova Persei. Kapteyn (1902) and Perrine (1903) recognized that these expanding annuli were created by light scattering off the dust, hence a so-called “echo”. However, a deeper understanding of LEs was established by Couderc (1939), who described much of the LE physics outlined in Section 2.2. Over the past 100 years, LEs have been discovered around several sources including variable stars, novae, Cepheids, young stars, and SNe, see Table 2.1. Zwicky (1940) was the first to propose that LEs could be observed around ancient SNe. There were several attempts to do this, including van den Bergh (1965) and van den Bergh (1966), in the twentieth century that were unsuccessful. Although SNe are the brightest objects known to produce LEs, they are also very rare in a given galaxy so few local candidates are available (Sugerman, 2003). Additionally, LEs are very faint, extended features making them very difficult to image (Rest et al., 2012b). With the invention of CCDs and larger telescopes, wide-field time-domain surveys made the detection of LEs from historic SNe possible (Rest et al., 2013). Before this point, LEs had only been discovered while the SN’s direct light was visible (Crotts

(1988); Suntzeff et al. (1988)). Rest et al. (2005) made the first observations of historic SN LEs around three SNe in the Large Magellanic Cloud (LMC) with the SuperMACHO survey. Since then, they have found LEs around SN 1572 (Tycho) and Cas A (Rest et al., 2008b), and other non-SN sources. Further work by Rest et al. (2008a) refined the process for classifying the SN based on the LE spectrum. Additionally, as multiple angles of the SN can be viewed with its LEs, Rest et al. (2011b) and Sinnott et al. (2013) were able to show the asymmetry of Cas A and SN 1987A, respectively.

2.1.1 Historic Supernova Light Echoes in the Milky Way

Of the six SNe introduced in Section 1.2, only two have successfully identified LEs (Rest et al., 2008b). Rest et al. (2008b) discovered LEs around Cas A, a CCSN, and SN 1572 (Tycho), a thermonuclear SN, with the Mayall 4 m Telescope at the Kitt Peak National Observatory. They observed several light echo arclets and used their apparent motion vectors to determine the origin of the emitted light, as described more thoroughly in Section 2.3.1. They found twelve LE clusters, half converged on the Cas A remnant and the remaining six on the Tycho remnant, see Figure 2.1. Although Cas A, Tycho, and SN 1181 are all in a similar region of the sky, no LEs were associated with SN 1181. Due to their relative line-of-sight (LOS) proximity, it is suspected that some of the Cas A and Tycho LEs were produced in the same dust structure. Krause et al. (2008b) measured Tycho’s LE spectrum with the Subaru telescope. They identified broad emission and absorption Si lines, classifying Tycho as a Type Ia SN. As described in Section 2.3.2, the LE spectrum is the time-integrated SN spectrum, and hence can be used for SN classification. Rest et al. (2011b) explored




Figure 2.1: The apparent motion vectors of several LE arclets surrounding SN 1572 (Tycho), Cas A, and SN 1181. Six clusters of arclets were found to originate at each of the Tycho and Cas A remnants. None of the LEs were associated with SN 1181.

The yellow arrow indicates the apparent motion vector of the LEs and blue line indicates the reverse of this vector, showing the origin of the light. Panel A includes all motion vectors and Panel B has the average of each cluster. The green boxes represent the fields that were imaged. Rest et al. (2008b), Figure 2.

the asymmetry of the Cas A SN explosion by examining the spectra of LEs scattered at three different locations, as described in Section 2.3.3. Each LE has a perspective of the SN from a different angle. They were able to confirm that Cas A had an asymmetric explosion.

LEs have yet to be found for SNe 1006, 1054, 1181, or 1604. Rest et al. (2008b) and McDonald (2012) both attempted to look for SN 1181 LEs. However, SN 1181 is considerably older than Tycho and Cas A, and would likely have LEs with a fainter surface brightness (Rest et al., 2008b). This work focuses on a survey of SN 1054 (Crab) fields, which is the only CCSN with recorded historic observations of the

direct light over the past 1000 years. Similarly to SN 1181, SN 1054 is quite old and likely had a low intrinsic luminosity. McDonald (2012) was unable to detect SN 1054 or SN 1181 LEs with the Canada-France-Hawaii Telescope’s MegaCam and visual inspection.

2.1.2 Extra-galactic Supernova Light Echoes

Thus far, only SNe within the Milky Way have been discussed. However, LEs have been detected for extra-galactic SNe, including several in the LMC (Rest et al., 2011a). As previously noted, Rest et al. (2005) used the SuperMACHO survey to capture the first historic SN LE images from SNe 0519-69.0, 0509-67.5, and 0509-68.7 (N103B), in addition to the more recent SN 1987A (see Figure 2.2). They began by extrapolating the apparent proper motion vectors, similar to Figure 2.1, to find the origin of the LEs. In addition to the location of SN 1987A, there were three other points of convergences, that were each in proximity to a known SNR. All three of these SNRs, 0519-69.0, 0509-67.5, and 0509-68.7 (N103B), are suspected to have been produced by Type Ia SNe. LE spectra obtained for 0509-67.5 by Rest et al. (2008a) confirmed the association. Rest et al. (2005) also estimated the age of two of the SNRs as $\sim 600 \pm 200$ years and $\sim 410 \pm 120$ years, respectively.

LEs have been observed from non-Milky Way SNe 1980K, 1991T, 1993J, 1995E, 1998bu, 2002hh, 2003gd, 2004et, 2006X, 2006bc, 2006gy, 2007it, 2008bk (Rest et al., 2013). SN 1980K was located in NGC 6946, known as the “Fireworks Galaxy” because of the nine SNe observed in the galaxy (Sugerman et al., 2012). This Type II-L SN had evidence of LEs from light interacting with a thin circumstellar shell, created by the progenitor. Sugerman & Crots (2002) discovered two LEs in the vicinity of SN




Figure 2.2: LEs from SN 1987A. Rest et al. (2005), Figure 1.

1993J with archived Hubble Space Telescope data. SN 1993J was previously classified as a Type IIb SN. As it was located in the spiral arm of M81, it was in an ideal position to produce LEs in the interstellar and circumstellar medium. Another LE was also found near SN 2003gd with archival Hubble Space Telescope data (Sugerman, 2005). SN 2002hh and SN 2003gd were both Type II-P SNe located in NGC 6946 and NGC 628 (M74), respectively.

2.1.3 Light Echoes from Other Sources

In addition to SNe, LEs have been detected from novae, Cepheids, variable stars, and young stellar objects. Several examples of non-SN LEs can be found in Table 2.1.

η Carinae was a significant LE discovery. It is a massive, binary variable star in the Milky Way. (Rest et al., 2012a). In the mid-nineteenth century, η Carinae ejected $> 10M_{\odot}$ and greatly increased its luminosity for two decades (Rest et al. (2012a); Rest et al. (2015)). This event was named the “Great Eruption” and has been called a “SN imposter” (Rest et al., 2012a). This LE discovery, with the Blanco CTIO 4-m

Name	Type of object	Citation
Nova Persei 1901	Nova	Ritchey (1901)
Sagittarii 1936	Nova	Swope (1940)
Nova 1975 Cygni	Nova	Bode & Evans (1985)
V838 Monocerotis	Variable star	Bond et al. (2003)
η Carinae	Variable star	Rest et al. (2012a)
RS Puppis	Cepheid	Westerlund (1961)
S CrA	Young star (T Tauri)	Ortiz et al. (2010)
R CrA	Young star (Herbig Ae/Be)	Ortiz et al. (2010)

Table 2.1: Non-SN LE Sources.

telescope, was significant because of the long duration of the transient (Rest et al., 2015). As described in Section 2.3.2, a LE spectrum can contain several epochs of the transient’s spectrum. This is significantly easier for long-term transients, such as η Carinae, making it possible to study how the spectrum evolved.

2.2 Light Echo Geometry

Figure 2.3 shows the imaginary ellipsoid which represents all possible scattering locations that will produce LEs visible for the observer (Rest et al., 2011a). The ellipsoid has the source and observer at the focal points. As light is emitted from a source, it expands out in all directions. If it encounters dust along the imaginary ellipsoid, then it has the potential to be scattered toward the observer. The ellipsoid expands as the light continues to travel outward. As the scattered light must first travel a distance r to the dust located on the ellipsoid and then to the observer, it takes a longer path than the direct light, which travelled distance D along the LOS. Therefore, the scattered light arrives at the observer with a delay time t after the direct light. The

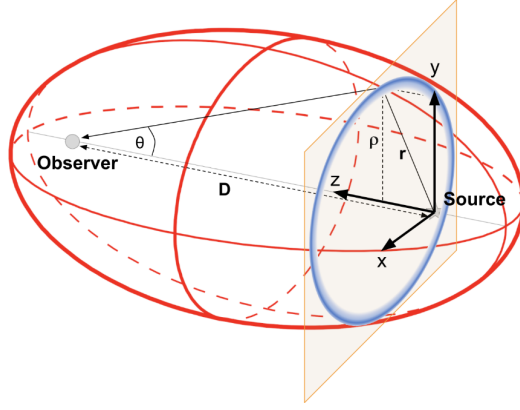


Figure 2.3: LE Geometric Diagram. Modified from Partoush et al. (2024), Figure 1.

horizontal distance of the plane of the dust from the source, z , can be represented by the LE equation when $z \ll D$ (Couderc, 1939):

$$z = \frac{\rho^2}{2ct} - \frac{ct}{2} \quad (2.2.1)$$

where ρ is the vertical distance of the dust from the LOS, such that $r^2 = \rho^2 + z^2$ and $\rho = (D - z) \tan(\theta)$ (Rest et al. (2008b); Partoush et al. (2024)). Therefore, if D and t are known, then the structure of the dust can be mapped with Equation 2.2.1 (Rest et al., 2005).

If the light scatters at multiple places along the ellipsoid, each LE contains a perspective from a different LOS (Rest et al., 2011a), as shown in Figure 2.4. Hence, it is possible to compare the spectra of the SN along different LOS. As SN 1987A had dust sheets nearly perpendicular to the LOS, as shown in Figure 2.3, nearly complete circular LEs were observed in Figure 2.2. Furthermore, if light scatters off dust at multiple planes $D - z$ from the observer, as shown in Figure 2.4, then multiple concentric circles of LEs will be visible, as is seen with SN 1987A in Figure 2.2.




Figure 2.4: Path of the light from seven different LEs observed around SN 1987A. Each LE was emitted in a different direction and scattered at a different location. Each LE contains spectral information along a different LOS. Rest et al. (2013), Figure 1.

2.3 Light Echoes as Possible Tools for Studying Supernovae

LEs have proved to be useful tools for studying historic supernovae. Examination of SN LEs and SNRs provides an unusual opportunity to examine the explosion and resulting remnant simultaneously (Rest et al., 2015). Furthermore, historic SNe can sometimes be studied even if their direct emission was not observed with modern instrumentation (Rest et al., 2011a). The observation and detection of LEs from historic SNe can expand our current understanding of SNe and their remnants.

2.3.1 LE Apparent Motion

The apparent proper motion of the LE can reveal the age and the origin of the LE. As shown in red in Figure 2.5, the LE arclet can be fit with a line (Rest et al., 2008b). The apparent proper motion vector (yellow) is assumed to be perpendicular to this arclet. If multiple LE features associated with the same SN are found, it is possible to determine the origin of the light, and hence the location of the SN, by extrapolating these motion vectors backward. If the SNR has not migrated since the explosion, it may be possible to find the remnant in the vicinity of the SN origin (Rest et al., 2005). This method was used by Rest et al. (2005) to determine that SNRs 0519-69.0, 0509-67.5, and 0509-68.7 (N103B) produced a set of LE arclets.

2.3.2 Spectroscopy

Historic SNe can be classified by analyzing the LE spectra (Rest et al., 2008b). This is possible because the LE spectrum represents a weighted, time-integrated spectrum of the original SN outburst (Rest et al., 2008a). More plainly, the annulus of scattered light is made up of many rings, each with the spectrum of a different epoch of the SN. The outermost ring was emitted earliest and the innermost ring was emitted most recently (Partoush et al., 2024). Hence, the SN's light curve is projected onto the sky. However, interpreting the LE spectrum is further complicated by observational factors, such as the point-spread function, finite slit size and orientation, and physical factors, such as the orientation of the dust sheet and LE arclet. Additionally, if the dust sheet was infinitely thin, then light would be scattering at one point on the ellipsoid. However, it is more realistic to have a dust sheet with a finite width,




Figure 2.5: The images in the top panel show a region near Tycho taken at two different epochs: Oct. 20, 2006 (left) and Dec. 13, 2007 (right). The lower left is the subtraction of the above images with the bright star difference artifacts masked with grey. The difference image is repeated in the lower right with a linear fit to the arclet (red), the apparent motion vector (yellow), and the reverse direction (blue).

Rest et al. (2008b), Figure 1.

where the light is scattered at multiple locations and angles. Spectral models, such as those by Rest et al. (2011a) and Partoush et al. (2024), use an effective light curve to counteract these effects, making it possible to approximate the original SN light curve.

Even with early models, it was possible to classify a SN’s type and subtype (over-luminous, under-luminous, etc.) with its LE spectrum hundreds of years after it was directly observed (Rest et al. (2011a); Rest et al. (2008b)). Rest et al. (2008a) was the first to achieve this, by identifying that the spectrum of a LE near the SNR 0509-67.5 matches the profile of an over-luminous thermonuclear SN, similar to SN 1991T. Rest et al. (2008a) compared the LE spectrum to 28 time-integrated Type Ia and 6 Type

Ib/c spectra by minimizing the χ^2 value. Krause et al. (2008b) was able to verify that Tycho (SN 1572) was a Type Ia SN with the LE spectra. The LE spectra showed clear silicon lines, which is a Type Ia identifier.

2.3.3 Supernova Asymmetry

Other than ancient SN classification, LE 3D spectroscopy can be used to study the asymmetry of the SN photospheric explosion (Rest et al., 2013). SNe can be observed from multiple perspectives with LEs if there are several concentrations of dust surrounding the SN (Rest et al., 2011b). Each LE contains hemispheric light from the SN with a different photospheric centre location.

Rest et al. (2011b) were the first to use observations of LEs to provide evidence that the Cas A SN was an asymmetric explosion. They were working with three different LEs evenly spaced around the SNR. They found that the spectra of one of the LEs differed from the other two, showing strong He I and H α absorption lines. Furthermore, that entire spectrum was blue-shifted, compared to the other two LEs, suggesting an asymmetric explosion. As the remnant has been thoroughly studied, the asymmetry of the SN could be compared to its (projected) structure.

Sinnott et al. (2013) also found evidence for an asymmetric structure in SN 1987A LEs. There are three dust sheets in the vicinity of SN 1987A, all roughly parallel to the plane of the sky. Given the large amount of dust and the recent occurrence of the SN, many LEs were detected (Rest et al., 2015), see Figure 2.4. SN 1987A also showed varying strengths of H α lines in different directions, which suggests a one-sided asymmetric outburst (Sinnott et al., 2013).

2.4 Core-Collapse Supernova Light Echoes

Of the five historic SNe with reliable observations from the last 1000 years, only SN 1054 (Crab) is a CCSN (Stephenson & Green, 2002). Cas A is also a CCSN, although it does not have historic observations of the direct light, it remains the only historic CCSN in the Milky Way with LE detections (Rest et al., 2008b). Currently, many more thermonuclear SN LEs have been observed than CCSN LEs (McDonald, 2012). The current non-detection of CCSN LEs is likely due to their lower average intrinsic luminosity. Furthermore, SN 1054 is nearly a thousand years old. Older SN LEs have had more time for the light to expand outward. Hence the intensity of the light decreases by $\frac{1}{r^2}$, where r is the distance between the SN origin and the LE. Therefore, older CCSNe LEs are considerably fainter and a wider angle of the sky may need to be searched to locate them. Prior attempts by McDonald (2012) to find LEs near the regions of the Crab SN (SN 1054) in the Milky Way with the Canada-France-Hawaii Telescope were unsuccessful. McDonald (2012) relied on visual inspection to detect LEs. This project is the preparatory work to begin the search for Crab LEs with the Dragonfly Telephoto Array, presented in Section 3.1.2, and with the machine-learning code ALED, discussed in Section 3.2.3. The discovery of LEs from the Crab SN would be a significant opportunity to research a CCSN that occurred several centuries before modern instrumentation had been invented.

Chapter 3

Light Echo Imaging and Detection

3.1 Light Echo Imaging

LEs are difficult to image because they are faint, extended sources (Rest et al., 2012b). Therefore, the telescope required for LE searching must have a large field of view and sufficient sensitivity to detect objects with low surface brightness. Although there were several attempts to image historic SN LEs in the twentieth century, including van den Bergh (1965) and van den Bergh (1966), the invention of mosaic CCDs made historic SN LE imaging and detection easier. Rest et al. (2005) used the wide-field time-domain survey, SuperMACHO, to image the first historic SN LEs in the LMC. They also utilized difference imaging, a technique involving subtracting two images of the same field taken at different epochs. Difference imaging removes the background and constant sources but highlights objects that have changed brightness and/or

position between epochs (Rest et al., 2005). As LEs constantly expand outward from the location of the SN and will show changes on a scale of months to years (for observable regions in the plane of the Milky Way disk), difference imaging can draw attention to LEs while removing much of the distracting background. Unfortunately, other transient objects, such as variable stars, reveal themselves in difference images. Figure 2.5 illustrates this process with a field containing Tycho SN LEs. Both images in the top panel are of the same field, but the images were taken more than a year apart. The resulting subtracted image, in the bottom left, has the bright stellar artifacts removed or masked and the LE prominently revealed diagonally from top-left to bottom-right.

3.1.1 Imaging Historic Supernova Light Echoes in the Milky Way

This project focuses on SN LEs in the Milky Way. Imaging LEs within our galaxy has advantages and disadvantages. A high dust density environment is required to scatter the light and produce LEs (Rest et al., 2012b). Such dust can be found in locations in the plane of the Milky Way (Stephenson & Green, 2002). However, dust can also make it difficult to detect more distant SNe. The distance to the SN and size of the LE search radius are inversely proportional for the same forward-scattering angle (Rest et al., 2015). Therefore, SN LE surveys in the Milky Way must cover a larger fraction of the sky than extra-galactic SN LE surveys. Hence, wide-field surveys are necessary for the detection of historic Milky Way SN LEs.

3.1.2 Light Echo Telescopes

Canada-France-Hawaii Telescope (CFHT)

Recent work by McDonald (2012) and its reanalysis by Bhullar et al. (2021) searched for LEs with MegaCam images from the Canada-France-Hawaii Telescope (CFHT). The CFHT is stationed on Maunakea, Hawaii (CFHT Observing Assistants, 2003). One of the instruments at this facility is MegaCam, a wide-field optical imager (CFHT Team, 2024). MegaCam’s features include:

- 40 2048 x 4612 pixel CCDs
- 378 megapixels
- Field-of-view of 1 square degree
- 0.187 arcseconds per pixel resolution

MegaCam was ideal for such a search because of its wide field-of-view and high angular resolution due to excellent conditions at Maunakea (McDonald, 2012).

Dark Energy Camera (DECam)

The Dark Energy Camera (DECam) was originally built for the Dark Energy Survey, to study the observational imprint of cosmological models on the distribution of distant galaxies (Abbott et al., 2016). DECam is located in Cerro Tololo, Chile and is part of the 4 m Blanco Telescope facility at the Cerro Tololo Inter-American Observatory (NOIRLab, 2023a). DECam is also a wide-field optical imager with the following properties (NOIRLab (2023b); NOIRLab (2024)):



Figure 3.1: The Dragonfly Telephoto Array. This image shows 24 of the 48 telephoto lenses and CCD cameras.

Source: <https://www.dragonflytelescope.org/>

- 62 CCDs
- 520 megapixels
- Field-of-view of 3 square degrees
- 0.263 arcseconds per pixel resolution

Due to the nature of the Dark Energy Survey, DECam was required to perform large-scale surveys with high resolution, making it similarly ideal for LE searches (Abbott et al., 2016). The LE DECam surveys covered a few hundred square degrees near known remnants. Our work with DECam LEs is discussed in Section 4.2.

The Dragonfly Telephoto Array (DTA)

The Dragonfly Telephoto Array (DTA) is a unique visible wavelength facility that is designed to capture extended, faint structures (Abraham & van Dokkum, 2014). The DTA is designed to be sensitive to ultra-low surface brightnesses and consequently has a high dynamic range, meaning it can differentiate between minor variations in

surface brightness. The DTA has the distinctive ability to capture objects with a surface brightness brighter than $\mu_B = 32$ mag arcsec². At very low surface brightnesses, diffraction and scattering light “ghosts” can become problematic, as discussed thoroughly in Chapter 4. To mitigate such issues, the enhanced anti-reflection coatings of the DTA optics and the use of all-refractive surfaces minimize light scattering and the appearance of ghosts (Abraham & van Dokkum (2014); Danieli et al. (2020)). For wide-angle observation, refractive lenses reduce scattering by a factor of 10 more than reflective mirrors (Danieli et al., 2020). Furthermore, the DTA does not have a central obstruction, as these do cause significant diffraction patterns.

A single-fast-large refracting telescope would be ideal for low surface brightness imaging (Abraham & van Dokkum, 2014). However, such optics are very costly and difficult to create. On a smaller scale, high-end commercially made telephoto lenses have many of the same features as refracting telescopes, as they can image distant, quickly moving objects. The configuration of the DTA used for these observations employed 48 Canon telephoto lenses, each mounted on a commercially-available, thermoelectrically-cooled CCD camera (Abraham & van Dokkum (2014); Danieli et al. (2020)), see Figure 3.1. These lenses are all pointed at the same field and the final image is a combination of the images from all of the cameras. Each lens has a 143 mm aperture ($d_{individual}$) (Abraham & van Dokkum, 2014). However, combining multiple cameras and lenses creates a very low f-ratio system with an effective aperture $D_{effective} = \sqrt{n} d_{individual}$ (~ 1 m) (Danieli et al., 2020). The combined image contains the entire dynamic range of all images but with the noise reduced by \sqrt{n} . The combined field-of-view of the entire array is $2.6^\circ \times 1.9^\circ$.

DTA is suited for observing LEs. Hence, utilizing DTA for LE detection and

imaging is the main focus of this work. The DTA is a promising tool for the detection of older and less luminous LEs. At the time of writing, we have 31 DTA difference images. These images were differenced with Properimage (Zacharias et al., 2013) and include one LE from the Cas A SN (Figure 5.1) and 30 non-LE images near the Crab SNR.

3.2 Light Echo Detection and Machine Learning

Identifying faint LEs in an image containing other bright objects (stars, etc.) is not easily accomplished. Difference imaging can assist with this process, by highlighting regions that have changed in brightness and/or position between two epochs. To create a difference image, two images of the same field that are taken months or years apart are required. By subtracting these images, any differences will be emphasized, as shown in Figure 2.5. However, LEs can be difficult to detect visually even in such difference images.

3.2.1 Visual Inspection

In work prior to 2021, detecting LEs in difference images was achieved solely by visual inspection. This method requires manually inspecting each image on a high-resolution display monitor to look for LE features (McDonald, 2012). Visual inspection requires a trained eye, is a very time-consuming task, and is susceptible to the usual challenges of the human attention span and fatigue. To provide a reference to the scale of this process, McDonald (2012) visually inspected 24,000 difference images in regions around the Crab and 1181 SNe, which did not result in any detections attributable

to the Crab SN or SN 1181. However, they did discover LEs from the Tycho and Cas A SNe, in the same portion of the sky where SN 1181 was seen.

3.2.2 Introduction to Machine Learning and Image Recognition

When discussing machine learning, we say that the computer is "learning" because it can improve its ability to perform an automated task without being explicitly told how to improve (Brown, 2021). A machine learning package utilizes a so-called neural network. A neural network, modelled after the human brain, is made up of many nodes (or neurons) which are organized into layers. The specific architecture of the network can vary greatly. For example, Convolutional Neural Networks (CNN) are commonly used for image recognition, however, Capsule Neural Networks (CapsNet) can also be used for this purpose in cases with small training sets (Hinton et al., 2018). Both of these networks contain several layers which each learns to identify different features in the images, but their specific organization is different.

A CapsNet is made of several layers of capsules (Bhullar et al., 2021). Each capsule contains the parameters required to identify/recreate an object in the form of a vector. In each subsequent capsule layer, there are more degrees of freedom, hence higher-dimension filters. The first layer can only identify simple objects, and the subsequent layers can recognize more complex objects while being influenced by previous capsules. For binary classification, the final layer is only required to identify whether or not a specific feature is included, as is the case with ALED. However, intermediate layers may identify other features while learning.

Why use a Capsule Neural Network?

Both CapsNets and CNNs require training sets to be input, that include several different sample images of the objects they are identifying. In cases with few training images, a CapsNet is preferential because CNNs require more weights and hence large training sets in practice. “Over-fitting” occurs when irrelevant features of the individual training examples are too heavily weighted. Furthermore, CapsNets can identify features in relation to each other and can locate which regions of an image are being triggered, using the following methods:

1. Pose Invariance

A CapsNet can recognize the relative spacing, orientation, and ordering of several features (Hui, 2017). Because CNNs are not designed to do this, they struggle to identify a recognized object at different orientations and scales. CNN would require multiple layers for each transformation of the object, making it more computationally expensive. More practically, this would require a large training set or risk over-fitting. CapsNets perform better with recognizing objects that have undergone transformations because the capsules can extrapolate from the training set examples to potential variations through what is known as a pose matrix. For example, both a CapsNet and a CNN could learn that a face is made up of two eyes, a nose, and a mouth, but only a CapsNet would learn that eyes are typically found above and on either side of the nose, and the mouth is directly below the nose. This property is called “pose invariance”.

2. Routing Path Visualization

The purpose of routing path visualization (RPV) is to identify the regions of the

image that are being triggered. For binary classification, the triggered regions are those that have a high probability of containing the desired object. These images are useful for locating the objects and identifying potential mistakes that the network is making.

3.2.3 Automated Light Echo Detection (ALED)

Recent advancements in machine learning have made it possible to semi-automate LE detection. Bhullar et al. (2021) created ALED (Automated Light Echo Detection), a Python package that applies supervised machine learning to identify LEs in difference images. ALED uses TensorFlow 2 packages (Abadi et al., 2016) in a Capsule Neural Network (CapsNet) to identify and locate LEs within an image. ALED learns to identify LEs through a training set of human-tagged images and performs binary classification on candidate LE images. Such an approach greatly reduces the number of images that must be visually inspected.

Bhullar et al. (2021) created a training set with 22 2400 x 4900 pixel CFHT images containing LEs discovered by McDonald (2012). Each 2400 x 4900 image was cropped into 350 200 x 200 pixel images, 175 of which contained at least a fraction of a LE. ALED had great success identifying LEs from CFHT difference images, with an accuracy of 90%, and was able to detect a set of Cas A LEs that had been missed during visual inspection by McDonald (2012).

ALED's Training Set

In its original version, ALED required a training set of LE example images, masks that identified the locations of those LEs, and a few false positive (FP) example images.




Figure 3.2: Each image that is input into ALED is cropped into 200 x 200 pixel images as shown.

Masks are a transparent overlay that indicates the location of an object of interest. Further discussion of the purpose of FPs in the training set and our improvements regarding FPs can be found in Section 4. The input images were cropped into multiple 200 x 200 pixel images, as shown in Figure 3.2 (Bhullar et al., 2021). If the number of pixels was not divisible by 200, then the edges of the image were padded with zeroes. Then, ALED sorted the cropped images into those with and without probable LEs. LEs account for only a small fraction of the entire image, 0.1 - 8 % of the total pixels, and the remaining pixels include the surrounding background sky. Therefore, many of the cropped images will not contain LEs. A cropped image was considered to contain a LE if it had at least 2500 masked pixels. Such background (BG) cropped images became part of the non-LE training set, along with the cropped false positive images. The resulting LE and non-LE sets were further divided into training, validation, and testing sets.

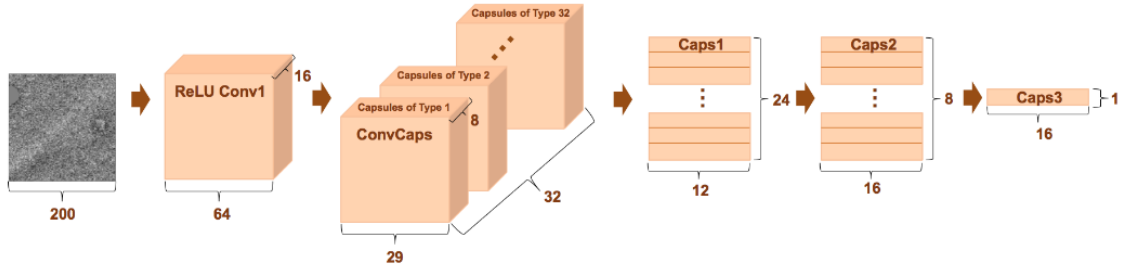


Figure 3.3: The layers of ALED: input image, ReLU activation layer, ConvCaps layer, and three capsule layers. Bhullar et al. (2021) Figure 2.

The Architecture of ALED

The layers of ALED are shown in Figure 3.3 (Bhullar et al., 2021). The input image is fed into the activation ReLU (Rectified Linear Unit) function. Then, a tensor, or “feature map”, is created that represents the image. This tensor is then convolved with 256 filters in the ConvCaps layer. The final three layers are capsule layers, containing 24, 8, and 1 capsules. The first capsule layer detects many features, including LEs and other artifacts and the second layer has a preference for saturated (and therefore zeroed-out) locations of bright stars. Only the capsule in the final layer is intended to specifically look for LEs. The identification of bright stars by the second capsule layer is informative but may influence the LE detection done by the final layer.

ALED’s Outputs

After the CapsNet was trained, a failure rate was output using the test set. The failure rate is the inverse of the accuracy:

$$accuracy = \frac{True\ Positives + True\ Negatives}{Number\ of\ Images} \times 100\% \quad (3.2.1)$$

The weights created during training are used to test a new set of images and sort them into those with at least one LE and those without any (Bhullar et al., 2021). During this testing stage, a RPV image is produced, that shows regions of the image with a high probability of containing a LE. Additionally, several values that represent the likelihood of an image containing a LE, called counts, are printed for each image. In the original ALED, images with $Count1 > 0$ are considered LE candidates. Then, only the LE candidates need to be visually inspected, greatly reducing the time required.

Why does ALED use a Capsule Neural Network?

Bhullar et al. (2021) chose to use a CapsNet for ALED because the number of LE images available for the training set was limited, so it was necessary to be able to train ALED with a small training set. CapsNets perform better than CNNs with small training sets and hence are the ideal choice here.

The Importance of Pose Invariance

Pose invariance is particularly important for LE identification. Dust is rarely located around the entire circumference of the SN, as with Figure 2.2. In most cases, only a very small portion of the annulus of a LE is captured. These LE segments may have varying lengths and orientations, but can all nominally be identified by ALED.

Routing Path Visualizations of LEs

Figure 3.4, includes a sample RPV image for CFHT (right). The yellow/green pixels have a high probability of being part of a LE, while the blue/purple pixels have a low




Figure 3.4: The image on the left is a CFHT difference image containing a large LE that arches from the upper left to the lower right. The image on the right is the corresponding RPV. The yellow/green regions, which have a high probability of being a LE, line up very well with the LE in the left image.

probability. The corresponding difference image (left) shows a LE diagonally across the upper half of the image. This example was chosen due to the clarity of the LE. Therefore, the average LE will not be so easy to identify visually. Figure 5.1 is a better example of a hidden LE, which may require a RPV to detect. During the process explained in Chapter 4, we often used original ALED RPV images to identify false positives for the training sets.

The Importance of Adapting Training Sets

identify Although ALED has previously performed very well with data from the Canada-France-Hawaii Telescope, we should not expect ALED to perform well with a new set of data without some adaptations. Most importantly, the training set will

need to be adjusted to include images from the new dataset. It is ideal to include LE and potential false positive images taken with the new telescope. If it is not possible to capture images in regions where LEs have already been detected with the new telescope, then steps can be taken to manufacture LE images - an innovation not used in the original ALED package. This augmentation approach is discussed in depth in Chapter 5. Different telescopes can produce a variety of LE-type artifacts, depending on their structure. False positives can be produced by several processes, such as light internally scattering within the optical path or diffraction from optical obstructions. These diffraction patterns appear in the difference images, as the alignment of the pattern in each image is rarely the exactly the same. Therefore, several distinct false positives need to be selected for the training set, as discussed in Chapter 4. Additionally, masks will need to be created for these new images and added to the training set.

Chapter 4

False Positives

False positives (FPs) are regions in a difference image that are known not to be LEs but share one or more properties of actual LEs. They are often mislabelled as LEs by the original version of ALED. Some FPs are known to be generated by the light from bright stars reflecting off the surface of the CCD and other surfaces, internal scattering in the telescope, diffraction patterns from central obstructions, and satellite trails. Difference images from a variety of telescopes will contain different mixes of FPs, so training with similar data is essential.

Type	No. of cropped images	Fraction of training set
LEs	155	0.50
BG	152	0.49
FPs	3	0.01

Table 4.1: Original ALED Training Set




Figure 4.1: An example of a CFHT diffraction pattern FP (left), a detailed FP mask (middle), and a less detailed FP mask (right). These masks were hand drawn with GIMP. We found that a more detailed mask produced better results and had a lower average failure rate.

Bhullar et al. (2021) found that 40% of the LE candidates identified by ALED were actually slightly offset diffraction patterns due to bright stars. To improve ALED performance regarding FPs, we focused on the choice and treatment of FPs in the training set. In its original version, ALED allowed for the inclusion of FP images. However, the FP images included in the training set were randomly chosen BG images, without any significant LE-type FP features. Additionally, the location of each image had to be individually added, and the cropped FP images were chosen manually. As a result, the influence of FPs on the overall training was small (only 1% of the training set), see Table 4.1.

A description of our versions of ALED can be found in Appendix A. We began by reviewing the LE examples in the original ALED training set, removing redundant and incorrectly labelled images (Version 1). Our first consideration to improve the

treatment of FPs in the training set was the selection of FP images. We examined the objects that ALED had mislabelled in the past, particularly diffraction patterns from bright stars. We searched through our collection of CFHT images and looked for potential FP examples with great diversity. An example of the images we chose can be seen in the left portion of Figure 4.1.

Table 4.3 includes the average failure rates of all CFHT versions of ALED for 28-30 retraining runs. We chose to retrain ~ 30 times because there was significant variation in the results between individual trainings. After ~ 30 retraining runs, the average failure rate had an uncertainty of approximately ± 0.02 .

Table 4.3 also contains the fraction of correctly identified FPs and LEs, using $Count1 > 0$ as an indicator of LE identification, as suggested by Bhullar et al. (2021). These *Counts* represent the likelihood that there is at least one LE in the image and are determined using the RPV. In an RPV image, such as the one shown on the right of Figure 3.4, the value of each pixel indicates the probability of the corresponding image pixel being part of a LE. The *Counts* are determined by counting the number of pixels in the RPV with a probability above a threshold. The thresholds in the original ALED are as follows:

$$\text{Threshold for } Count1 = 0.00042$$

$$\text{Threshold for } Count2 = 0.00037$$

$$\text{Threshold for } Count3 = 0.00030$$

The FP and BG images should be labelled with $Count1 = 0$ for non-LE and the LE images should be labelled with $Count1 > 0$.

In Version 3, we added a set of nine false positive images, and carefully chose 20 cropped images. As was expected, the average failure rate of Version 3 was higher

than the original version due to the addition of more false positives into the training set. This increase was not surprising as Bhullar et al. (2021) carefully adjusted the parameters of ALED to maximize performance with CFHT images. We did not optimize the parameters with CFHT FP data, although we did this for the DTA training sets, see Chapters 6. However, the 20 cropped FP images remain a significantly smaller fraction of the training set because manually selecting the cropped images was time-consuming. Our solution was to create false positive masks for the original full-size difference images.

4.1 False Positive Masks

We have updated ALED so that it accommodates masks for FPs. The mask contains the same number of pixels as its corresponding image but with a pixel value of one in all regions of importance and zero elsewhere. Therefore, these masks tell ALED which regions of the images contain FPs. By using the masks, ALED can automatically include the false positive containing cropped images into the training set. By incorporating FP masks, we can increase the fraction of FP cropped images in the training set without the increased time required to manually select cropped images. The final version (12) has 16 FP images in the training set, with a choice of up to 819 FP containing cropped images to include. A consequence of this is that we were able to choose what fraction of FP images we wanted to include in the non-LE portion of the training set. A more detailed discussion of this choice can be found below. With the addition of FP masks, ALED can learn to ignore commonly misidentified objects, such as satellite trails, if they are relatively common in the FPs.

Before searching for FPs in a new dataset, ALED is run with just LEs in the

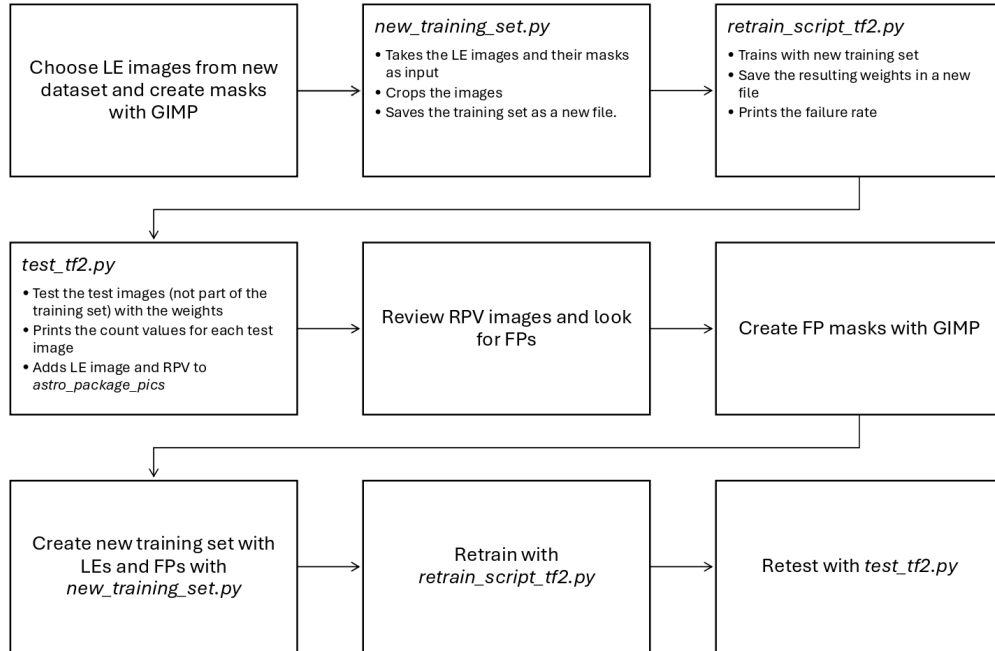


Figure 4.2: Workflow Diagram for retraining ALED with FP images.

training set. Figure 4.2 shows the details of this process. First ALED must be fully trained and tested with a LE-only test set. The RPVs from this trial will indicate the sorts of objects that are mislabelled as LE by ALED. After these FPs are revealed, they can be masked and used in the training set for the next trial. The FP masks were made with the same method as the LE masks, with GIMP (GIMP Team, 2019). GIMP is an image editing software, that we used to create a transparent layer over the image and paint black over the regions of interest. Then, the original image layer is removed, leaving only the transparent overlay with important pixels indicated by black.

We gave some consideration to the level of detail that these masks should contain. Figure 4.1 shows two masks for the same FP image (left): the detailed mask (middle) has the features of the diffraction pattern masked and the imprecise mask (right) has the entire diffraction pattern covered. When comparing the results of version 5 (which included only the imprecise version) and version 7A (which had a similar fraction of detailed masks), there is not a significant difference in the average failure rate for detailed and non-detailed masks. However, the detailed mask training sets were moderately better at identifying FPs as non-LEs. Therefore, we chose to use detailed masks for the remaining trials.

We also needed to consider whether the fraction of FPs and BG cropped images in the non-LE training set affect ALED’s ability to correctly categorize LEs and non-LEs. The A-E variations of versions 7 and 12 have different fractions of FP and BG cropped images in the non-LE training set, as shown in Table 4.2. Versions 7 and 12 are identical, except for 10 additional FPs in the non-LE training set added to Version 12. See Appendix A for details about these versions.

Variation Label	Fraction of FPs	Fraction of BG
A	80%	20%
B	75%	25%
C	85%	15%
D	90%	10%
E	70%	30%

Table 4.2: Fractions of FP and BG cropped images in the non-LE training set in Versions 7 and 12.

As shown in Table 4.3, variations A-E result in no major differences in the average failure rate for either Versions 7 or 12. Furthermore, no variation consistently




Figure 4.3: A poorly differenced DECam FP (left) and its corresponding RPV (right). ALED is triggered by the regions with high contrast, which are unfortunately a result of poor differencing.

performs better at categorizing LEs and non-LEs. Therefore, we can conclude that the exact fraction of FPs and BG cropped images included in the non-LE training set does not have a significant effect on ALED’s performance. Therefore, our default fractions for the remaining training sets had 80% FP and 20% BG images.

It was not surprising that all of our CFHT versions of ALED did not perform as well as the original ALED. Bhullar et al. (2021) adjusted the parameters of their version to minimize the failure rate. We did not do any specific fine-tuning for our CFHT versions, as we were focused on adapting ALED to work with DTA data.

4.2 DECam False Positives

Although ALED has performed very well with data from the Canada-France-Hawaii Telescope, some adaptations needed to be made for ALED to perform equally well on a new set of data. While waiting for the DTA data, we decided to attempt using ALED with DECam images. The DECam images of SN 1987A contain a large selection of LEs from full, obvious LE arcs to LEs near the detection limit. We created masks for these LEs with GIMP and added them to the existing CFHT training set (Versions

9-11 in Appendix A). However, the overall LE detection results with DECam and ALED did not meet expectations.

Although the DECam images contained full arcs of LEs generated by dust sheets between SN1987A and the observer, there were many differencing errors that ALED misidentified as LEs. Figure 4.3 shows an example of the gradient that can occur in images due to poor differencing (top). The gradients can be created by scattered light which was not accounted for during differencing, such as the difference images being taken with varying levels of background light. Once the new DECam training set was tested by ALED, the resulting RPV (Figure 4.3, right) showed that ALED is being triggered by these high contrasting regions. Differenced LEs often contain adjacent dark and light pixels, so it makes sense that ALED is triggered by this phenomenon. Even when we added FP images such as these to the training set, we still saw ALED being triggered by differencing errors. Unfortunately, without redoing the DECam differencing, it is unlikely that ALED would be able to be used to identify LE in DECam images.

Version No.	Average Failure Rate (± 0.02)	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs
Original	0.06	0.12	0.99
3	0.07	0.02	0.98
5	0.09	0.09	0.99
7A	0.08	0.12	1.00
7B	0.09	0.03	1.00
7C	0.07	0.08	0.99
7D	0.08	0.08	1.00
7E	0.08	0.11	0.99
12A	0.10	0.04	0.98
12B	0.08	0.08	0.97
12C	0.09	0.08	1.00
12D	0.11	0.08	0.96
12E	0.10	0.06	1.00
12A retest		0.06	0.98
12B retest		0.07	0.97
12C retest		0.05	1.00
12D retest		0.09	0.96
12E retest		0.05	1.00

Table 4.3: The average failure rates, fractions of correctly identified FPs and LEs based on ~ 30 retraining runs with CFHT training sets. Each version was tested with ~ 180 cropped FP images and ~ 450 cropped LE images. Version 12 was retested with 300 additional cropped FP images, which are labelled as “retest”. Unsurprisingly, our CFHT versions do not perform as well as the original ALED by Bhullar et al. (2021), as they carefully adjusted the parameters to maximize performance with those images.

Chapter 5

Manufacturing and Augmenting Light Echo Images

To test whether the DTA would be able to image LEs, a single region where LEs had already been detected by CFHT, near the Cas A SN, was imaged. Figure 5.1 is to date the only LE captured by the DTA and is notably almost a degree long. Clearly, one DTA LE does not comprise a sufficient training set. Even if the CFHT LE images from the previous training sets are included, a DTA training set requires the majority of examples to be captured with the DTA. Fortunately, augmentation can be used to expand such modest training sets. Sections 5.1 - 5.3 outline our methods to create a usable DTA training set. All of the code we used for these methods can be found in the GitHub repository *LE-Augmentation* (Mulyk & Welch, 2024).




Figure 5.1: The only LE captured with the DTA to date. The LE was produced by the Cas A SN and is circled in red. As Cas A LEs had already been detected by CFHT in this region, this image was a test that the DTA could capture LEs.

5.1 Overlaying CFHT Light Echoes onto Dragonfly Images

The first method we used to extend our DTA training set was manufacturing DTA LEs from the CFHT detections. These manufactured DTA LEs were made by isolating and overlaying the CFHT LEs onto DTA images, see *manufacture_LE.py* (Mulyk & Welch, 2024).

There are two major differences between the CFHT and DTA LE images, shown in Table 5.1, which need to be considered during this process. First, the time interval between the difference images is significantly longer for the DTA. While the CFHT difference images were taken only months apart, the DTA difference images had intervals of a few years. In the latter case, the LE has more time to change its apparent position between the captured images, so that it won't be partially or completely self-subtracted during differencing.

Image Differences	CFHT	DTA
Scale (arcsec/pixel)	0.185	2.5
Time interval between difference images	months	years

Table 5.1: Differences between DTA and CFHT Images

Second, the resolution of the DTA images is significantly larger than CFHT images. Therefore, the LEs in the DTA images cover a smaller fraction of the overall image than LEs in the CFHT images, as shown in Table 5.2. Additionally, the DTA cropped image with the largest fraction of LE pixels was only 27% of the cropped image. Where as CFHT had LEs which entirely consumed their cropped images.

Telescope	Average Fraction of LE Pixels	Largest Fraction of LE Pixels
CFHT	0.27	1.00
DTA	0.04	0.27

Table 5.2: Average Fraction of LE Pixels in 200 x 200 Cropped Images

The following steps were taken to manufacture a DTA LE from an existing CFHT LE:

1. Isolate the CFHT LE pixels by using the LE mask to set non-LE pixels to zero.
2. Pad the edges of the CFHT image so it is the same shape as the DTA image.
3. Reproject the isolated CFHT LE to rescale it to match the DTA's scale.
4. Optional: Adjust the brightness of the LE. This option is used to test whether a dimmer LE can be recognized by ALED.
5. Optional: Randomly change the LE's orientation and location. This step helps

increase the variation in the manufactured images but, due to the pose invariance property of capsule networks, is not expected to linearly improve training.

6. Change the resolution of the LE with a Gaussian filter of $\sigma = 1$ pixel to match the DTA image.
7. Overlay the LE onto the DTA image. As the DTA image only contains the sky and artifacts from differencing bright stars, it serves as the background for the new image.
8. Create a mask for the new LE by setting all LE-containing pixels to one and all non-LE pixels to zero.

Figure 5.2 contains the original CFHT LE (left) and the LE reprojected onto the DTA image (right). The LE (circled in red) is considerably smaller in the DTA image. The change in scale required significant changes in ALED as these small LEs are often visually indistinguishable from some of the artifacts left by bright star differencing. These adaptations are discussed more completely in Section 6.4.

5.2 Augmenting Existing Light Echoes

We also wanted to utilize the only real DTA LE image to expand our training set. We augmented Figure 5.1 in a similar method as described above, by overlaying the isolated real LE onto different DTA images while varying the position, orientation, brightness, and width of the gap between linear segments of the LE, see *augment_real_LE.py* (Mulyk & Welch, 2024). The steps undertaken were as follows:

1. Create a detailed mask for the real LE using GIMP.




Figure 5.2: The original CFHT LE (left) was overlaid on a DTA image. The reprojected LE is circled in red in the DTA image (right). Note the noticeable change in the size of the LE.

2. Isolate the real LE pixels by using the LE mask to set non-LE pixels to zero.
3. Pad the edges of the real LE image so it is the same shape as the DTA image.
4. Reproject the isolated real LE onto the new DTA image.
5. Optional: Adjust the brightness of the LE.
6. Optional: Adjust the width of the gap between line segments of the LE. This is done by rotating the LE so that the two LE line segments are vertical, isolating the pixels associated with each line segment, and shifting the pixels horizontally. By doing this, it mimics a SN with long emission (wide) or short emission (narrow).
7. Optional: Randomly change the LE's orientation and location.
8. Overlay the LE onto the new DTA image.

9. Create a mask for the new LE by setting all LE-containing pixels to one and all non-LE pixels to zero.

This process allows us to view the same LE with several different backgrounds, orientations, brightnesses, and widths.

5.3 Masking Out Bright Stars Difference Artifacts

As discussed more thoroughly in Section 6.1, many of the objects being misidentified as LEs by ALED in the DTA data were bright star difference artifacts (BSDA). The DTA LEs are difficult to distinguish from BSDA as they have comparable sizes and contrast. Therefore, by masking out these artifacts before training, it is less important for ALED to identify these objects correctly, see *mask_bright_stars.py* (Mulyk & Welch, 2024). We located the stars using the Fourth U.S. Naval Observatory CCD Astrograph Catalog (UCAC4) (Zacharias et al., 2013). The following steps were used to mask out the stars in a DTA image:

1. Locate the region encompassed by the image using the image header.
2. Find the same region in the UCAC4 data and list the locations of all stars in this region.
3. Create a circular aperture around each star with the radius r found with a

piece-wise function of the stars' magnitude m :

$$r(m) = \begin{cases} 5.0 & m > 15.0 \\ 7.0 & 13.0 < m \leq 15.0 \\ 10.0 & 11.0 < m \leq 13.0 \\ 13.0 & 10.0 < m \leq 11.0 \\ 18.0 & m \leq 10.0 \end{cases} \quad (5.3.1)$$

where the units of r are pixels and the units of m are magnitudes.

4. Mask out the stars by setting all pixels within the circular apertures to zero.

After the stars are masked out, the previously described LE manufacturing and LE augmentation processes can be used to create a large training set without stellar FP difference artifacts.

Chapter 6

Dragonfly Light Echoes

Our original DTA training set (DF1 in Appendix A) had 18 CFHT LEs that had been overlaid and re-scaled onto DTA images, as explained in Section 5.1, and the single real Cas A DTA LE. These 19 images were cropped into 49 100 x 100 pixel images. We realized that more DTA LE images were required to create a training set with a similar number of LE containing cropped images as we had in the CFHT versions, because of the scale of the DTA LEs. Therefore, the next training set (DF2) contained an additional 35 (54 total) manufactured DTA LEs and the real LE, for a total of 115 cropped images.

We chose to crop the images in the training set to 100 x 100 pixels, rather than 200 x 200 pixels as in the original ALED. We chose a smaller cropped image size because the DTA LEs contributed to a considerably smaller fraction of the image, as shown in Table 5.2. Section 6.4.1 contains more discussion about the choice of

cropped image size.

An additional parameter that needed to be adjusted to adapt ALED for smaller LEs was the threshold number of pixels for which a cropped image is considered to contain a LE. In *new_training_set.py*, this parameter is labelled *num_ones*, as in the number of ones in a cropped mask. *num_ones* indicates the number of pixels which are associated with a LE. For the original ALED, if a cropped image contained more than 2499 LE pixels, then it was placed in the LE-containing training set. When we used this threshold for the DTA images, we found that many LE-cropped images were being categorized with the non-LE training set because they contained too few LE pixels. Therefore, we initially chose a threshold of *num_ones* = 100 pixels for the DTA training sets. However, Section 6.4 discusses other trials with different thresholds.

Versions DF3 and DF4, contained versions DF1 and DF2, respectively, with the addition of the CFHT LE images from Version 1. Table 6.5 shows the average failure rate and the fraction of correctly identified FP, LE, and BG images for 20 retraining runs for all DTA versions. These first four DTA training sets (DF1, DF2, DF3, and DF4) did not perform very well, with average failure rates > 0.1 . DF1 and DF2 only identified 16.3% and 9.5% of the tested LE images. DF3 and DF4 performed considerably better, and recognized 88.4% and 91.9% of the tested LE images. The addition of CFHT LEs, creating a larger training set, greatly improved the results. DF4 shows the most improvement with the addition of the CFHT images, as it already had a larger training set. We chose to continue improving DF4 with the addition of FPs, see Section 6.1.

Most DTA versions, except for DF21, were tested with a standard set of test

images that included 70 cropped FP images, 290 cropped LE images (including manufactured and augmented LEs, and the real LE), and 220 cropped background images. The results of these tests can be found in Tables 6.5 and 6.6, at the end of this chapter.

In the original ALED, an image was considered to contain at least one LE if it had a $Count1 > 0$. This $Count$ suggested the likelihood of the presence of at least one LE in an image. It was calculated by tallying the number of pixels in the corresponding RPV with a probability > 0.00042 . Table 6.5 uses this threshold to determine whether the LE or non-LE (BG or FP) images had been correctly categorized. However, this threshold had been chosen to determine the presence of LEs in CFHT training sets, without a prominent focus on FPs. Therefore, we expected that this threshold was not ideal for our DTA training sets. The trials in Table 6.5 with 100% of the LEs correctly identified, and none of the non-LE images correctly identified, are not very informative. Therefore, Table 6.6 uses $Count > 0$ determined by counting the RPV pixels with a probability > 0.001 . This threshold is more informative for comparing two DTA versions of ALED, although it may underestimate the performance. Appendix B contains all of the DTA versions with the fraction of correctly identified LE, FP, and BG images for probability thresholds ranging from 0.0003 - 0.0010. It is difficult to determine which probability threshold provides the clearest picture of ALED's performance with DTA data. Larger thresholds show moderately better performance correctly labelling non-LE features (FPs and BG), but poor performance with LEs. However, the failure rate, calculated during retraining, can be used as a consistent measure of success across all versions. For this work, we use the threshold > 0.00042 to compare the DTA versions with CFHT versions, and the threshold > 0.001 to compare DTA versions to each other.

As has been said multiple times throughout this work, the training set must contain images from the dataset that will be tested. To prove this statement, we tested the newly manufactured DTA LEs in DF1 with the CFHT weights from the original ALED. ALED was able to identify 78.2% of the DTA LE images with the CFHT weights, as shown in Table 6.5. This moderate result shows that ALED can satisfactorily identify LEs with a generalized training set, but an individualized training set is required for peak performance.

6.1 Dragonfly False Positives

We found that the objects that ALED was typically misidentifying as LEs in DTA images were very different from those in CFHT images. Our CFHT FPs were typically bright star diffraction pattern differences, such as the one on the left of Figure 4.1. We used the RPV images from the first four DTA trials to search for DTA FPs. As the DTA LEs were very similar in size to BSDA, these stars were the most common FPs in DTA images, therefore ALED had a difficult time distinguishing BSDA from LEs. An example of these BSDA is shown in Figure 6.1 (left). The RPV on the right side of Figure 6.1 shows how BSDA can trigger ALED.

DF5 was the first DTA model that contained FPs. Table 6.5 shows that the average failure rate decreased significantly to 0.043 after the addition of FPs in the training set. Since the inclusion of CFHT LEs in the DTA training set greatly improved the fraction of correctly identified LEs, we also included CFHT FPs in the training set for DF6. However, the addition of CFHT FPs greatly increased the average failure rate, to similar levels as versions without FPs. Therefore, we performed further tests with the DF5 training set, with some alterations. The increased average




Figure 6.1: An example of BSDA (left) which were misidentified by ALED and were used as FPs in the DTA training sets. The corresponding RPV (right) shows how the BSDA triggered ALED, see regions highlighted with yellow/green.

failure rate is unsurprising, as the CFHT FPs (Figure 4.1, left) are very different from the DTA FPs (Figure 6.1, left). This result shows the importance of individualizing a training set for a new data set.

6.2 Results with Manufactured and Augmented Dragonfly LEs

All training sets up to and including DF9 include only the manufactured LEs, which were CFHT LEs overlaid on DTA images, as explained in Section 5.1, and the unaltered real DTA LE. DF7, DF8, and DF9 were used as test sets and are discussed in more detail in Section 6.6. DF10 has the same base training set as DF5 but with the addition of the augmented real LEs, see Section 5.2 for a description of this process. When we compare DF5 and DF10 with a probability threshold > 0.001 , the addition of augmented real LEs leads to a larger fraction of non-LE (BG or FP) images being correctly identified, see Table 6.6. However, the fraction of correctly identified LEs

also decreases, leading to a larger failure rate. DF11 through DF20 are all variations on DF 10 and are discussed in Sections 6.4 and 6.6.

6.3 Results with Masked Bright Stars Diffraction Artifacts

The images with the BSDA masked out, as described in Section 5.3, were included in DF21. These images had all BSDA in the DTA images masked out first. Then, the CFHT LEs and augmented real DTA LEs were overlaid onto the masked DTA images. This process created a similar training set to DF10, including the DTA FPs, but with the BSDA masked out of the background. As shown in Table 6.6, DF21 performed similarly to its unmasked counterpart (DF10) in regard to the average failure rate. When testing with unmasked backgrounds with a probability threshold > 0.001 , this version performs very poorly at correctly categorizing non-LEs (FPs and BG), but very well at identifying LEs. This is unsurprising, as we are testing whether ALED can differentiate between LEs and BSDA without inputting a sufficient number of differenced bright star examples. There was no significant improvement when the test images also had the BSDA masked out.

6.4 ALED Adaptations for Dragonfly Light Echoes

In this section, the parameters that can be adjusted in ALED are explored. As previously discussed, the original version of ALED had the parameter set to maximize performance with CFHT images. However, the DTA images are significantly different

from the CFHT images, as discussed in Section 5.1. Most notably, the LE size in comparison to the entire image is considerably smaller, see Table 5.2. Therefore, adjustments may need to be made to improve ALED’s success with DTA images.

6.4.1 Adjusting Cropped Image Size

As explained in the introduction of this chapter, we chose to crop the images in the training set to 100 x 100 pixels, instead of 200 x 200 pixels, because the average scale of the LEs is considerably smaller in the DTA images. Another benefit of decreasing the cropped image size is that it decreases the training time and the runtime of the testing code, and it allows the LEs to cover a larger portion of the cropped images. However, we were unsure how this would affect the overall performance of ALED. Therefore, we tested two variations of DF10 with larger cropped sizes, see Table 6.1, with our standard test set.

Version No.	Cropped Image Size
DF10	100
DF13	150
DF14	200

Table 6.1: Cropped Image Size for DTA Versions

Table 6.6 shows that, although it decreased runtime, reducing the cropped image size did minimally affect performance. The average failure rate does not vary greatly between DF10, DF13, and DF14. However, DF14 correctly categorized nearly all of the LEs but misidentified nearly all of the non-LE (BG or FP) images. Given the mixed results, we continued to crop images to 100 x 100 pixels, to maximize efficiency.

6.4.2 Adjusting Pixel Threshold for Sorting Training Images

The second parameter, *num_ones*, is the number of LE pixels required for a cropped image to be placed in the LE training set. This value is determined by counting the number of ones in the corresponding cropped mask. For CFHT images, the original ALED used a threshold of 2499 pixels. However, as the fraction of LE pixels is considerably smaller in the DTA-cropped images, see Table 5.2, we found that this threshold excluded several LE examples from the LE training set. Therefore, we chose a threshold of 100 pixels for the DTA training sets. A threshold of 100 pixels was chosen because we found through trial and error that it encompassed the majority of the LE images. But, we also wanted to explore how varying *num_ones* affected the overall performance of ALED with DTA images. We varied the pixel threshold, *num_ones*, for versions DF15 - DF20 to compare performance. Table 6.2 contains the exact *num_ones* value for each version.

Version No.	<i>num_ones</i>
DF10	100
DF15	0
DF16	50
DF17	250
DF18	500
DF19	1000
DF20	2500

Table 6.2: Pixel Threshold *num_ones* for DTA Versions

Table 6.6, shows that the average failure rate decreases considerably with larger pixel thresholds. Although the fraction of correctly identified LEs initially improves with an increased threshold, the performance for identifying LEs with *num_ones* =

2500 is very poor. However, with $num_ones = 2500$, we see a larger fraction of non-LE (BG or FP) images correctly categorized than any other version.

6.5 The Size of Light Echoes Successfully Identified

Given that many DTA FPs were near-point source BSDA, we suspected that ALED would have more difficulty identifying smaller LEs. To test this assertion, we divided the LE images in the DF10 training set into three groups by the fraction of LE pixels to total pixels in the image, using the corresponding masks. The DTA image with the largest LE had only 0.185% of the total pixels contributing to the LE. Since there were more images with a small fraction of LE pixels, we grouped the images as shown in columns 1 and 2 of Table 6.3. The choice of these ranges of LE sizes was somewhat arbitrary, but chosen to try to make the size of each group roughly equal. Then, we tested these groups with the same pixel thresholds, num_ones , as used in Tables 6.5 and 6.6. Overall, there was not a significant difference in the fraction of correctly identified LEs across the three sizes. One could argue that ALED's ability to identify small LEs was marginally weaker than with larger LEs, but this distinction is insignificant.

Category	Fraction of LE Pixels	<i>Probability Threshold</i>	Fraction of Correct LE
Small	$frac < 0.0001$	0.00042	1.00
		0.001	0.68
Medium	$0.0001 < frac < 0.0003$	0.00042	1.00
		0.001	0.82
Large	$frac > 0.0003$	0.00042	0.99
		0.001	0.77

Table 6.3: The LEs were split into small, medium, and large based on the fraction of LE pixels in the image. The fraction of correctly identified LE pixels is given for probability thresholds 0.00042 and 0.001.

6.6 The Brightness of Light Echoes Successfully Identified

With the processes described in Sections 5.1 and 5.2, we could change the brightness of LEs. We chose to adjust the brightness of our two base training sets DF5 and DF10 by multiplying the LE brightness by the factors given in column 3 of Table 6.4. These new brightened/dimmed images were tested with weights from the original unaltered training sets (DF5 and DF10). Table 6.4 shows that ALED’s ability to identify LEs does not change drastically when the tested LEs are brightened or dimmed.

Weights Used	Test Set	Brightness Factor	Probability Threshold	Fraction of Correctly Identified LEs
DF5	DF5	N/A	0.00042	1.00
			0.00100	0.93
DF5	DF7	0.1	0.00042	1.00
			0.00100	0.93
DF5	DF8	10	0.00042	1.00
			0.00100	0.93
DF5	DF9	100	0.00042	1.00
			0.00100	0.95
DF10	DF10	N/A	0.00042	0.99
			0.00100	0.76
DF10	DF11	10	0.00042	0.99
			0.00100	0.77
DF10	DF12	0.1	0.00042	0.99
			0.00100	0.76

Table 6.4: The fraction of correctly identifies LEs for thresholds 0.00042 and 0.00100, with test images with the brightness of the LEs varied by the factors in column 3. These tests were performed with DF5 and DF10 weights, the results of training with images that had the original LE luminosity. A complete list of all probability thresholds can be found in Appendix C.

Version No.	Average Failure Rate (± 0.02)	Fraction of Correct FPs	Fraction of Correct LEs	Fraction of Correct BG
DF1	0.181		0.16	
DF1_orig			0.78	
DF2	0.159		0.10	
DF3	0.113		0.88	
DF4	0.101		0.92	
DF5	0.043	0.00	1.00	0.01
DF6	0.150		1.00	
DF10	0.054	0.00	0.99	0.01
DF13	0.062	0.00	1.00	0.00
DF14	0.057	0.00	1.00	0.00
DF15	0.053	0.00	1.00	0.00
DF16	0.051	0.00	0.99	0.00
DF17	0.049	0.01	0.97	0.03
DF18	0.042	0.00	1.00	0.01
DF19	0.032	0.00	1.00	0.00
DF20	0.013	0.56	0.40	0.60
DF21	0.057	0.00	1.00	0.00

Table 6.5: The average failure rates, fractions of correctly identified FP, LE, and BG images based on 20 retraining runs with DTA training sets. Most versions were tested with 70 cropped FP images, 290 cropped LE images (including manufactured and augmented LEs, and the real LE), and 220 cropped background images, except for DF21 which was tested with 265 cropped FP images, 1096 cropped LE images, and 830 cropped background images. The determination of whether or not an image contains a LE was done using a probability threshold > 0.00042 , as in the original ALED. A complete list of all count values can be found in Appendix B.

Version No.	Average Failure Rate (± 0.02)	Fraction of Correct FPs	Fraction of Correct LEs	Fraction of Correct BG
DF5	0.043	0.04	0.90	0.08
DF10	0.054	0.16	0.76	0.26
DF13	0.062	0.20	0.78	0.26
DF14	0.057	0.01	0.97	0.05
DF15	0.053	0.10	0.84	0.16
DF16	0.051	0.00	0.96	0.02
DF17	0.049	0.03	0.93	0.06
DF18	0.042	0.04	0.91	0.13
DF19	0.032	0.00	1.00	0.01
DF20	0.013	0.80	0.15	0.83
DF21	0.057	0.00	1.00	0.01

Table 6.6: The same data as Table 6.5, but using a probability threshold > 0.001 . A complete list of all count values can be found in Appendix B.

Chapter 7

Discussion and Conclusion

During this thesis, we focused on creating effective training sets with FP images, for the machine-learning package ALED. We adapted ALED to identify DTA LEs at its significantly different image scale from ALED's CFHT difference images. During this process, we made the following improvements to ALED:

1. We showed that the selection of training set images, especially FPs, is vital to successful LE identification. A general rule for comparing DTA versions of ALED is that larger training sets typically perform better. We found this to be true when comparing the first four models with DTA training sets. However, the results from these versions also revealed that the selection of images in the training set is far more important. For example, DF2 included 35 more LE images than DF1. There was moderate improvement between these versions. However, the addition of 19 substantially different CFHT LE images,

rather than 35 similar DTA LE examples, resulted in a significant performance improvement for DF3. More remarkably, adding 7 FP images in DF5 had a similar improvement as 19 CFHT LEs. These results reinforce the need for a diverse set of FP images in the training set. Furthermore, the addition of images to the training set can significantly increase the computational expense and running time. Therefore, it is most economical to carefully choose the images in the training set to maximize diversity.

2. We focused on the inclusion of FP features in the training sets by adapting ALED to accept FP masks. By incorporating FP masks, we greatly increased the fraction of cropped FP images in the non-LE training set from only 1%, in the original ALED, to 70-90%. Although the exact fraction of FP images in the training set does not matter, the more detailed masks performed better overall.
3. Given the limited number of DTA LEs, we created a process to overlay known true LEs from a different telescope (CFHT) onto a difference image to manufacture a LE image. This procedure requires that the LE be isolated and rescaled, and have the resolution changed to match the new images. Options were also included to augment the new LE images, by changing the LE's position, orientation, and brightness in the new image.
4. We also created a procedure to augment the existing DTA LE and overlay it over other DTA images. This augmentation process was similar to the previous process but also allowed the width between the LE line segments to be adjusted. Hence, replicating a LE with longer (several months or years) or shorter (weeks or a few months) emission. The addition of the augmented images into

the training set increased the fraction of non-LE images that were correctly categorized.

5. We identified BSDA as the most common FPs for DTA images and added DTA FPs to the training sets with FP masks. This significantly improved the average failure rate.
6. We created a procedure to mask out BSDA before training and testing using the UCAC4 data from Zacharias et al. (2013). We performed this masking by creating a circular aperture around the star and setting all pixels within that aperture to zero. The size of the aperture was determined as a function of the stars' magnitude by a piece-wise function. This process is useful for removing potential FPs from the training and test sets.
7. We made several adaptations to ALED to adjust for the smaller scale of DTA LEs, including the pixel threshold for LE identification and cropped image size.
8. We found that there was no substantial evidence that ALED preferentially identified small or large LEs, as well as bright or dim LEs.

With the exception of DF20, with the largest pixel threshold $num_ones = 2500$, ALED is still showing poor performance with identifying FPs as non-LEs. However, DF20 also performed very poorly at identifying LEs. Good performance with LEs corresponding to poor performance with non-LEs (FP and BG images) was common among all variations. Therefore, time would still be required to remove FPs from the LE candidates. Future work with ALED that allowed it to identify FPs as a separate category may resolve this issue. Currently, ALED performs binary classification, so it can only categorize images as containing at least one LE and not containing any

LEs. Non-binary classification could allow for three categories: LE-containing, FP-containing, and not containing either LEs or FPs. This adaptation should not require significant changes to the architecture, as ALED is already triggered by FPs in the capsule layers. Additionally, further improvements on the DTA differencing pipeline could remove many of the FP artifacts before training.

Overall, our DTA versions of ALED had a comparable or lower average failure rate than the original version. The low average failure rate shows that ALED has been successfully adapted to identify DTA LEs. This achievement is significant because the DTA's ability to observe extended ultra-low surface brightness structures makes it ideal for imaging LEs. When combined with ALED's capacity to efficiently and accurately identify LEs, it should be possible to detect older, less luminous LEs. Generally, CCSNe have a lower average intrinsic luminosity and are older than Type Ia SNe, so the light is spread over a larger area and hence has lower intensity. Therefore, fewer CCSN LEs have been detected. To date, no LEs have been found around the historic CCSN Crab (SN 1054). We have completed the preparatory work for a search for SN LEs around the Crab SN once more DTA images in the region around the Crab SN are acquired.

Appendix A

Versions of ALED

Telescope being Tested	Version Name	Description
CFHT	Original	Original version of ALED created by Bhullar et al. (2021).
	1	Removed duplicate and incorrectly labelled LEs from training set.
	2	Replaced the background FP training set with 3 images that contained diffraction patterns, and carefully selected 8 cropped images (instead of 3).

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
	3	Added 6 more images with greater diversity to FP training set, and carefully selected 20 cropped images.
	4	Added masks for FPs. See Section 4.1.
	5	Adjusted none-LE training set to contain 80% randomly chosen FPs cropped images and 20% randomly chosen background cropped images.
	6	Created more detailed masks.
	7A	Version 6 with 80% randomly chosen FPs cropped images and 20% randomly chosen background cropped images.
	7B	Version 6 with 75% randomly chosen FPs cropped images and 25% randomly chosen background cropped images.
	7C	Version 6 with 85% randomly chosen FPs cropped images and 15% randomly chosen background cropped images.

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
	7D	Version 6 with 90% randomly chosen FPs cropped images and 10% randomly chosen background cropped images.
	7E	Version 6 with 70% randomly chosen FPs cropped images and 30% randomly chosen background cropped images.
	8	Adjusted size of batches for training. Tried 5 (default), 2, and 8 elements in each batch.
DECam	9	Added DECam LE images to existing CFHT LE training set
	10	Removed CFHT LE images from training set, but retained DECam LE images. Kept CFHT FPs images.
	11	DECam LEs and FPs images only in the training set

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
CFHT	12A	Returned to version 6 training set. Added 10 more CFHT FP images. 80% randomly chosen FPs cropped images with detailed masks and 20% randomly chosen background cropped images.
	12B	Version 12 with 75% randomly chosen FPs cropped images and 25% randomly chosen background cropped images.
	12C	Version 12 with 85% randomly chosen FPs cropped images and 15% randomly chosen background cropped images.
	12D	Version 12 with 90% randomly chosen FPs cropped images and 10% randomly chosen background cropped images.
	12E	Version 12 with 70% randomly chosen FPs cropped images and 30% randomly chosen background cropped images.

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
DTA	DF1	New DTA training set with manufactured LE images, see Section 5.1 split into 49 cropped images. No FPs images included. Images cropped to 100x100 pixels. Threshold number of pixels for LE classification is 100 pixels.
	DF1_orig	Manufactured DTA LEs images tested with original CFHT weights.
	DF2	Added new manufactured LE images to DTA training set for a total of 115 cropped images. No FPs images included.
	DF3	Added CFHT LE images to version DF1.
	DF4	Added CFHT LE images to version DF2.
	DF5	Version DF4 with DTA FPs.
	DF6	Version DF4 with DTA and CFHT FPs.
	DF7	Version DF5 with the brightness of the manufactured DTA LEs dimmed by a factor of 10.

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
	DF8	Version DF5 with the brightness of the manufactured DTA LEs brightened by a factor of 10.
	DF9	Version DF5 with the brightness of the manufactured DTA LEs brightened by a factor of 100.
	DF10	Added 10 augmented real DTA LEs to DF5 training set, see Section 5.2.
	DF11	Version DF10 with the brightness of the augmented real DTA LEs brightened by a factor of 10.
	DF12	Version DF10 with the brightness of the augmented real DTA LEs dimmed by a factor of 10.
	DF13	Version DF10 images cropped to 150x150 pixels.
	DF14	Version DF10 images cropped to 200x200 pixels.

Continued on the next page

Continued from previous page

Telescope being Tested	Version Name	Description
	DF15	Version DF10 with threshold number of pixels for LE and FP classification changed to 0 pixels.
	DF16	Version DF10 with threshold number of pixels for LE and FP classification changed to 50 pixels.
	DF17	Version DF10 with threshold number of pixels for LE and FP classification changed to 250 pixels.
	DF18	Version DF10 with threshold number of pixels for LE and FP classification changed to 500 pixels.
	DF19	Version DF10 with threshold number of pixels for LE and FP classification changed to 1000 pixels.
	DF20	Version DF10 with threshold number of pixels for LE and FP classification changed to 2500 pixels.
	DF21	Recreated DTA augmented and manufactured LEs with images that have bright stars masked out, see Section 5.3. Images cropped to 100x100 pixels. Threshold number of pixels for LE classification is 100 pixels.

Appendix B

Fractions of Correctly Identified False Positive, Light Echo, and Background Images in Versions of ALED with Dragonfly Training Sets

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
DF1	0.00042		0.163	
DF1_orig			0.782	
DF2			0.095	

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
DF3			0.884	
DF4			0.919	
DF5	0.0003	0.000	0.997	0.005
	0.00037	0.000	0.997	0.005
	0.00042	0.000	0.997	0.005
	0.0005	0.000	0.997	0.005
	0.0006	0.000	0.997	0.005
	0.0007	0.000	0.997	0.005
	0.0008	0.000	0.997	0.005
	0.0009	0.000	0.968	0.023
	0.0010	0.043	0.895	0.077
DF6	0.00042	0.000	1.000	0.000
DF7	0.0003	0.000	0.976	0.014
	0.00037	0.000	0.972	0.014

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.00042	0.000	0.969	0.018
	0.0005	0.014	0.959	0.023
	0.0006	0.029	0.938	0.036
	0.0007	0.029	0.924	0.050
	0.0008	0.086	0.866	0.100
	0.0009	0.086	0.817	0.173
	0.0010	0.157	0.748	0.236
DF8	0.0003	0.000	0.993	0.005
	0.00037	0.000	0.993	0.005
	0.00042	0.000	0.993	0.005
	0.0005	0.000	0.983	0.009
	0.0006	0.000	0.979	0.009
	0.0007	0.000	0.979	0.014
	0.0008	0.000	0.979	0.014

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0009	0.000	0.979	0.014
	0.0010	0.014	0.976	0.014
DF9	0.0003	0.000	0.969	0.045
	0.00037	0.000	0.966	0.055
	0.00042	0.014	0.952	0.059
	0.0005	0.043	0.917	0.082
	0.0006	0.071	0.907	0.109
	0.0007	0.100	0.883	0.127
	0.0008	0.100	0.869	0.145
	0.0009	0.100	0.862	0.159
	0.0010	0.114	0.831	0.177
DF10	0.0003	0.000	1.000	0.005
	0.00037	0.000	0.993	0.009
	0.00042	0.000	0.993	0.009

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0005	0.000	0.976	0.018
	0.0006	0.000	0.976	0.023
	0.0007	0.029	0.945	0.055
	0.0008	0.071	0.910	0.095
	0.0009	0.114	0.855	0.173
	0.0010	0.157	0.759	0.259
DF11	0.0003	0.029	0.966	0.027
	0.00037	0.029	0.966	0.027
	0.00042	0.029	0.966	0.027
	0.0005	0.029	0.962	0.027
	0.0006	0.029	0.959	0.036
	0.0007	0.029	0.959	0.041
	0.0008	0.043	0.952	0.045
	0.0009	0.043	0.941	0.064

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0010	0.057	0.921	0.082
DF12	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.000
	0.0005	0.000	1.000	0.000
	0.0006	0.000	0.997	0.005
	0.0007	0.000	0.986	0.023
	0.0008	0.000	0.972	0.036
	0.0009	0.029	0.938	0.064
	0.0010	0.029	0.921	0.082
DF13	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.000
	0.0005	0.029	0.966	0.045

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0006	0.100	0.910	0.082
	0.0007	0.100	0.886	0.123
	0.0008	0.114	0.869	0.164
	0.0009	0.157	0.821	0.200
	0.0010	0.200	0.776	0.259
DF14	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.000
	0.0005	0.000	1.000	0.000
	0.0006	0.000	1.000	0.000
	0.0007	0.000	0.993	0.014
	0.0008	0.000	0.990	0.023
	0.0009	0.000	0.986	0.036
	0.0010	0.014	0.972	0.050

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
DF15	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.000
	0.0005	0.000	1.000	0.000
	0.0006	0.000	0.990	0.005
	0.0007	0.000	0.983	0.014
	0.0008	0.014	0.959	0.027
	0.0009	0.029	0.907	0.100
	0.0010	0.100	0.838	0.164
	DF16	0.0003	0.000	0.997
0.00037		0.000	0.993	0.000
0.00042		0.000	0.990	0.000
0.0005		0.000	0.990	0.000
0.0006		0.000	0.990	0.000

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0007	0.000	0.986	0.009
	0.0008	0.000	0.976	0.018
	0.0009	0.000	0.969	0.018
	0.0010	0.000	0.962	0.23
DF17	0.0003	0.000	0.979	0.023
	0.00037	0.000	0.972	0.027
	0.00042	0.014	0.966	0.027
	0.0005	0.014	0.962	0.027
	0.0006	0.029	0.955	0.027
	0.0007	0.100	0.886	0.123
	0.0008	0.029	0.931	0.041
	0.0009	0.029	0.931	0.055
	0.0010	0.29	0.937	0.059
DF18	0.0003	0.000	1.000	0.000

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.005
	0.0005	0.000	0.993	0.014
	0.0006	0.000	0.986	0.027
	0.0007	0.000	0.983	0.032
	0.0008	0.000	0.959	0.064
	0.0009	0.014	0.917	0.100
	0.0010	0.043	0.907	0.132
DF19	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000
	0.00042	0.000	1.000	0.000
	0.0005	0.000	1.000	0.000
	0.0006	0.000	1.000	0.000
	0.0007	0.000	1.000	0.000

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.0008	0.000	1.000	0.000
	0.0009	0.000	0.997	0.009
	0.0010	0.000	0.997	0.009
DF20	0.0003	0.500	0.448	0.550
	0.00037	0.529	0.414	0.577
	0.00042	0.557	0.397	0.595
	0.0005	0.557	0.386	0.609
	0.0006	0.600	0.362	0.627
	0.0007	0.629	0.314	0.677
	0.0008	0.743	0.241	0.745
	0.0009	0.786	0.183	0.800
	0.0010	0.800	0.148	0.832
DF21	0.0003	0.000	1.000	0.000
	0.00037	0.000	1.000	0.000

Continued on the next page

Continued from previous page

Version No.	Probability Threshold	Fraction of Correctly Identified FPs	Fraction of Correctly Identified LEs	Fraction of Correctly Identified Background
	0.00042	0.000	1.000	0.000
	0.0005	0.000	1.000	0.000
	0.0006	0.000	1.000	0.000
	0.0007	0.000	0.998	0.002
	0.0008	0.000	0.998	0.002
	0.0009	0.000	0.998	0.002
	0.0010	0.000	0.996	0.005

Appendix C

Fractions of Correctly Identified Light Echo Images using DF5 and DF10 Weights with Test Sets that have Various Brightness Levels

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
DF5	DF5	none	0.00030	0.998
			0.00037	0.998
			0.00042	0.998

Continued on the next page

Continued from previous page

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
			0.00050	0.998
			0.00060	0.998
			0.00070	0.998
			0.00080	0.998
			0.00090	0.985
			0.00100	0.931
DF5	DF7	dimmed by a factor of 10	0.00030	0.998
			0.00037	0.998
			0.00042	0.998
			0.00050	0.998
			0.00060	0.998
			0.00070	0.998
			0.00080	0.998
			0.00090	0.986

Continued on the next page

Continued from previous page

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
			0.00100	0.934
DF5	DF8	brightened by a factor of 10	0.00030	0.996
			0.00037	0.996
			0.00042	0.996
			0.00050	0.996
			0.00060	0.996
			0.00070	0.996
			0.00080	0.996
			0.00090	0.991
			0.00100	0.928
DF5	DF9	brightened by a factor of 100	0.00030	0.998
			0.00037	0.998
			0.00042	0.998

Continued on the next page

Continued from previous page

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
			0.00050	0.998
			0.00060	0.998
			0.00070	0.998
			0.00080	0.998
			0.00090	0.995
			0.00100	0.945
DF10	DF10	none	0.00030	0.997
			0.00037	0.992
			0.00042	0.992
			0.00050	0.977
			0.00060	0.975
			0.00070	0.942
			0.00080	0.903
			0.00090	0.847

Continued on the next page

Continued from previous page

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
			0.00100	0.758
DF10	DF11	brightened by a factor of 10	0.00030	0.997
			0.00037	0.994
			0.00042	0.994
			0.00050	0.977
			0.00060	0.976
			0.00070	0.948
			0.00080	0.904
			0.00090	0.846
			0.00100	0.766
DF10	DF12	dimmed by a factor of 10	0.00030	0.997
			0.00037	0.992
			0.00042	0.992

Continued on the next page

Continued from previous page

Weights Used	Test Set	Brightness Variation	Probability Threshold	Fraction of Correctly Identified LEs
			0.00050	0.978
			0.00060	0.975
			0.00070	0.946
			0.00080	0.908
			0.00090	0.846
			0.00100	0.763

Bibliography

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, <https://www.tensorflow.org/>. <http://arxiv.org/abs/1603.04467>
- Abbott, T., Abdalla, F. B., Aleksić, J., et al. 2016, MNRAS, 460, 1270, doi: 10.1093/mnras/stw641
- Abraham, R. G., & van Dokkum, P. G. 2014, PASP, 126, 55, doi: 10.1086/674875
- Badenes, C., Hughes, J. P., Bravo, E., & Langer, N. 2007, ApJ, 662, 472, doi: 10.1086/518022
- Bhullar, A., Ali, R. A., & Welch, D. L. 2021, A & A, 655, A82, doi: 10.1051/0004-6361/202039755
- Bode, M. F., & Evans, A. 1985, A & A, 151, 452
- Bond, H. E., Henden, A., Levay, Z. G., et al. 2003, Nature, 422, 405, doi: 10.1038/nature01508

- Branch, D. 1998, *Annu. Rev. Astron. Astrophys.*, 36, 17, doi: 10.1146/annurev.astro.36.1.17
- Branch, D., & Nomoto, K. 1986, *A & A*, 164, L13
- Branch, D., & Tammann, G. A. 1992, *Annu. Rev. Astron. Astrophys.*, 30, 359, doi: 10.1146/annurev.aa.30.090192.002043
- Brown, S. 2021, *Ideas Made to Matter*, MIT Management Sloan School
- CFHT Observing Assistants. 2003, *CFHT Observatory Manual (Canada-France-Hawaii Telescope)*, Section 1. https://www.cfht.hawaii.edu/Instruments/ObservatoryManual/CFHT_Observatory_Manual_TOC.html
- CFHT Team. 2024, *MegaPrime/MegaCam - The state-of-the-art in wide-field imaging at CFHT (Canada-France-Hawaii Telescope)*. <https://www.cfht.hawaii.edu/Instruments/Imaging/MegaPrime/>
- Chevalier, R. A. 1977, In *Supernovae* (Reidel)
- Couderc, P. 1939, *Annales d'Astrophysique*, 2, 271
- Crotts, A. 1988, *IAU Circ.*, 4561
- Danieli, S., Lokhorst, D., Zhang, J., et al. 2020, *ApJ*, 894, 119, doi: 10.3847/1538-4357/ab88a8
- Doggett, J. B., & Branch, D. 1985, *ApJ*, 90, 2303, doi: 10.1086/113934
- Filippenko, A. V. 1997, *Annu. Rev. Astron. Astrophys.*, 35, 309, doi: 10.1146/annurev.astro.35.1.309

- Gilfanov, M., & Ákos Bogdán. 2010, *Nature*, 463, 924, doi: 10.1038/nature08685
- GIMP Team. 2019, GIMP - GNU Image Manipulation Program (2.10.34). <https://www.gimp.org>
- Hancock, P. P., Gaensler, B. M., & Murphy, T. 2011, *ApJ*, 735, L35, doi: 10.1088/2041-8205/735/2/L35
- Hillebrandt, W., & Niemeyer, J. C. 2000, *Annu. Rev. Astron. Astrophys.*, 38, 191, doi: 10.1146/annurev.astro.38.1.191
- Hinton, G. E., Sabour, S., & Frosst, N. 2018, in *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJWLfGWRb>
- Hubble, E. P. 1928, *ASP*, 1, 55, doi: 1928ASPL....1...55H
- Hui, J. 2017, Understanding Matrix capsules with EM Routing - Based on Hinton's Capsule Networks (Jonathan Hui blog). <https://jhui.github.io/2017/11/14/Matrix-Capsules-with-EM-routing-Capsule-Network/>
- Kapteyn, J. C. 1902, *Astronomische Nachrichten*, 157, 201
- Ko, T., Suzuki, H., Kashiyama, K., et al. 2024, *ApJ*, 969, 116, doi: 10.3847/1538-4357/ad4d99
- Krause, O., Birkmann, S. M., Usuda, T., et al. 2008a, *Science*, 320, 1195, doi: 10.1126/science.1155788
- Krause, O., Tanaka, M., Usuda, T., et al. 2008b, *Nature*, 456, 617, doi: 10.1038/nature07608

- Krause, O., Rieke, G. H., Birkmann, S. M., et al. 2005, *Science*, 308, 1604, doi: 10.1126/science.1112035
- Leibundgut, B. 2001, *Annu. Rev. Astron. Astrophys.*, 39, 67, doi: 10.1146/annurev.astro.39.1.67
- . 2008, *Gen Relativ Gravit*, 40, 221, doi: 10.1007/s10714-007-0545-9
- Leonard, D. C. 2007, *ApJ*, 670, 1275, doi: 10.1086/522367
- Lundmark, K. 1921, *PASP*, 33, 225, doi: 10.1086/123101
- Maoz, D., Mannucci, F., & Nelemans, G. 2014, *Annu. Rev. Astron. Astrophys.*, 52, 107, doi: 10.1146/annurev-astro-082812-141031
- Maoz, D., Sharon, K., & Gal-Yam, A. 2010, *ApJ*, 722, 1879, doi: 10.1088/0004-637X/722/2/1879
- McDonald, B. J. 2012, The Search for Supernovae Light Echoes from the Core-Collapse Supernovae of AD 1054 (Crab) and AD 1181 (McMaster University). <https://macsphere.mcmaster.ca/handle/11375/12461>
- Minkowski, R. 1941, *PASP*, 53, 224, doi: 10.1086/125315
- Mulyk, N., & Welch, D. 2024, LE_Augmentation. https://github.com/nmulyk/LE_Augmentation
- NOIRLab. 2023a, Mid-Scale Observatories — Cerro Tololo Inter-American Observatory (National Science Foundation). <https://noirlab.edu/science/programs/ctio>

- . 2023b, Dark Energy Camera (DECam) (National Science Foundation). <https://noirlab.edu/science/programs/ctio/instruments/Dark-Energy-Camera>
- . 2024, NOIRLab Capabilities 2024 (National Science Foundation), 8. <https://noirlab.edu/public/media/archives/brochures/pdf/brochure027.pdf>
- Ortiz, J. L., Sugerman, B. E. K., de la Cueva, I., et al. 2010, *A & A*, 519, A7, doi: 10.1051/0004-6361/201014438
- Pakmor, R., Kromer, M., Taubenberger, S., et al. 2012, *ApJ*, 747, L10, doi: 10.1088/2041-8205/747/1/L10
- Partoush, R., Rest, A., Jencson, J. E., et al. 2024, *APJ*, accepted. <http://arxiv.org/abs/2310.01501>
- Perrine, C. D. 1903, *ApJ*, 17, 310, doi: 10.1086/141032
- Rest, A., Sinnott, B., & Welch, D. L. 2012a, *PASA*, 29, 466, doi: 10.1071/AS11058
- Rest, A., Sinnott, B., Welch, D. L., et al. 2011a, *ApJ*, 732, 2, doi: 10.1088/0004-637X/732/1/2
- Rest, A., Sinnott, B., Welch, D. L., Prieto, J. L., & Bianco, F. 2013, *Proceedings of the International Astronomical Union*, 9, 126, doi: 10.1017/S1743921313009356
- Rest, A., Sinnott, B., Welch, D. L., et al. 2015, *Fifty Years of Wide Field Studies in the Southern Hemisphere*, 491, doi: 10.48550/arxiv.1502.03705
- Rest, A., Suntzeff, N. B., Olsen, K., et al. 2005, *Nature*, 438, 1132, doi: 10.1038/nature04365

- Rest, A., Matheson, T., Blondin, S., et al. 2008a, *ApJ*, 680, 1137, doi: 10.1086/587158
- Rest, A., Welch, D. L., Suntzeff, N. B., et al. 2008b, *ApJ*, 681, L81, doi: 10.1086/590427
- Rest, A., Foley, R. J., Sinnott, B., et al. 2011b, *ApJ*, 732, 3, doi: 10.1088/0004-637X/732/1/3
- Rest, A., Prieto, J. L., Walborn, N. R., et al. 2012b, *Nature*, 482, 375, doi: 10.1038/nature10775
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, 116, 1009, doi: 10.1086/300499
- Ritchey, G. W. 1901, *ApJ*, 14, 293, doi: 10.1086/140868
- Ritter, A., Parker, Q. A., Lykou, F., et al. 2021, *ApJL*, 918, L33, doi: 10.3847/2041-8213/ac2253
- Ruiter, A. J., Belczynski, K., & Fryer, C. 2009, *ApJ*, 699, 2026, doi: 10.1088/0004-637X/699/2/2026
- Sinnott, B., Welch, D. L., Rest, A., Sutherland, P. G., & Bergmann, M. 2013, *ApJ*, 767, 45, doi: 10.1088/0004-637X/767/1/45
- Smartt, S. J. 2009, *Annu. Rev. Astron. Astrophys.*, 47, 63, doi: 10.1146/annurev-astro-082708-101737
- Stephenson, F. R., & Green, D. A. 2002, *Historical Supernovae and their Remnants* (Oxford University Press), doi: 10.1093/acprof:oso/9780198507666.001.0001

- Sugerman, B. E. K. 2003, *AJ*, 126, 1939, doi: 10.1086/378358
- . 2005, *ApJ*, 632, L17, doi: 10.1086/497578
- Sugerman, B. E. K., & Crofts, A. P. S. 2002, *ApJ*, 581, L97, doi: 10.1086/346016
- Sugerman, B. E. K., Andrews, J. E., Barlow, M. J., et al. 2012, *ApJ*, 749, 170, doi: 10.1088/0004-637X/749/2/170
- Suntzeff, N. B., Heathcote, S., Weller, W. G., et al. 1988, *Nature*, 334, 135, doi: 10.1038/334135a0
- Swope, H. H. 1940, *Harvard College Observatory Bulletin*, 913, 11
- van den Bergh, S. 1965, *PASP*, 77, 269, doi: 10.1086/128214
- . 1966, *PASP*, 78, 74, doi: 10.1086/128299
- Wang, L., & Wheeler, J. C. 2008, *Annu. Rev. Astron. Astrophys.*, 46, 433, doi: 10.1146/annurev.astro.46.060407.145139
- Weiler, K. W., Panagia, N., Montes, M. J., & Sramek, R. A. 2002, *Annu. Rev. Astron. Astrophys.*, 40, 387, doi: 10.1146/annurev.astro.40.060401.093744
- Weiler, K. W., & Sramek, R. A. 1988, *Annu. Rev. Astron. Astrophys.*, 26, 295, doi: 10.1146/annurev.aa.26.090188.001455
- Westerlund, B. 1961, *PASP*, 73, 72, doi: 10.1086/127622
- Woosley, S. E., & Weaver, T. A. 1986, *Annu. Rev. Astron. Astrophys.*, 24, 205, doi: 10.1146/annurev.aa.24.090186.001225

Zacharias, N., Finch, C. T., Girard, T. M., et al. 2013, *AJ*, 145, 44, doi: 10.1088/0004-6256/145/2/44

Zwicky, F. 1940, *Rev. Mod. Phys.*, 12, 66, doi: 10.1103/RevModPhys.12.66

—. 1965, In *Stars and Stellar Systems*, Vol. 8 (Univ. Chicago Press Highlight), 367–423