Discovery of jets launched by classical nova eruptions

Dominic McLoughlin
Oriel College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Hilary 2021

Abstract

Classical novae are stellar explosions that occur on the surface of white dwarfs in close binary systems. This thesis presents time-domain spectroscopy of four recent eruptions in the Milky Way, carried out by the Global Jet Watch telescopes. Careful analysis of these observations reveals the presence of a circumbinary disc around classical nova V906 Car following its cataclysmic event in March 2018. Its existence may provide an answer to the long-standing problem of explaining the white dwarf mass distribution.

This thesis also presents the discovery of line-emitting jets and an accretion disc in classical nova YZ Reticuli, using sustained follow-up spectroscopy commencing soon after the eruption in July 2020 and lasting several months. A qualitatively similar spectral signature is also observed for classical nova V6568 Sgr, hinting that jets may be a feature of classical novae more generally than previously thought.

The same jets and accretion disc model is subsequently applied to the remaining two classical novae in the Global Jet Watch archive, V5668 Sgr and V906 Car, which exhibit the same behaviour. In the case of V5668 Sgr, the jets are found to persist for at least 2000 days after the eruption in March 2015. This thesis then suggests future imaging measurements that could be obtained using current instrumentation for spatially-resolved spectroscopy, which would elucidate further properties of nova jets like those discovered in all four classical novae observed by the Global Jet Watch.

This thesis also uses Global Jet Watch time-lapse spectroscopy on classical novae V906 Car and V5668 Sgr to tightly constrain (to within hundreds of seconds) the duration of certain phases and their transitions that take place in the immediate aftermath of these recent eruptions.

Statement of Originality

I confirm that to the best of my knowledge, the contents of this thesis are my own original work, except where explicitly declared otherwise.

This thesis is predominantly based on data captured by the Aquila instruments of the Global Jet Watch network of telescopes, which were designed and commissioned by my supervisor Katherine M. Blundell and Steven Lee. I designed and wrote several software packages (in the Python programming language) to facilitate robotisation of the observations, and automatic checks, detailed in Chapter 3. I also contributed heavily to the main data processing pipeline (ENDEAVOUR); these contributions are specifically detailed in Section 3.5.1, and are the subject of Chapters 4, 5, 6 and 7. The data analysis for this thesis was carried out using a desktop spectroscopy application called POIROT which I designed and built.

This thesis also makes use of data from the Shelyak LHires instrument currently mounted at the Rainbow observatory built and operated by Steven Lee, for which I designed and wrote the entirety of the data processing software (Chapter 8).

The work contained in Chapter 9 has been peer-reveiewed, and is published as McLoughlin et al. in the Monthly Notices of the Royal Astronomical Society, 494:743–749 (2020). The work contained in Chapter 10 has also been peer-reviewed, and is published as McLoughlin et al. in the Monthly Notices of the Royal Astronomical Society, 503:704–71 (2021). The work contained in Chapter 11 has also been peer-reviewed, and is published as McLoughlin et al. in the Monthly Notices of the Royal Astronomical Society, 505:2518–2529 (2021). This thesis was entirely written by me. My doctoral supervisor, Katherine Blundell, was the second author for all three papers, and provided the raw data and contributed many valuable discussions. The third author, Steven Lee, designed and built the spectrographs, and helped capture some of the data used in this thesis. The fourth author, Chris McCowage, also helped capture data.

The light curve data presented in this thesis was recorded by the American Association of Variable Star Observers (AAVSO).

Acknowledgements

I would like to thank the Science and Technology Facilities Council for my doctoral studentship, and Oriel College for a graduate scholarship.

I also acknowledge with thanks the variable star observations from the American Association of Variable Star Observers (AAVSO) International Database contributed by observers worldwide and used in this research.

I thank my supervisor, Katherine, for your humour, patience, and curiosity that have guided me to get far more out of this process than I possibly could have without you. Thanks also to Steven Lee, for designing the spectrographs whose elegant spectra made all this possible.

I would like to thank the tutors at Pembroke, especially Alfons, Tim and Bill, who kept me engaged during my undergraduate physics degree and supported my wider development throughout.

I thank my friends, for making me laugh and for all the good times we have had.

I am grateful to Peter Allwright, for your infinite wisdom, of which only a tiny fragment can be neatly summarised by that one simple phrase: "learn some words". Thank you also for the time you gave me 29/30.

Thank you Wayne, Michelle and everyone else at Anaconda for your mentorship. Thanks for giving me the gift of freedom in water, and for teaching me what things are really important. I thank OUSC, OUWPC and OUMPA and all the wonderful people involved in them who made my time at Oxford so enjoyable. I also warmly acknowledge the staff at the Iffley Road pool for their hard and valuable work.

To Gabrielle, thank you for making every day an adventure; you make my life complete. Every moment with you is precious, and I cherish all of our past memories, as well as those waiting to be made.

To my grandparents, I am tremendously grateful to have had the opportunity to learn so much from you, and I very much appreciate the warmth, love and kindness you have all shown me over the years.

Most of all, I want to thank my mum and dad for always supporting and guiding me over the years, and teaching me the most important lessons — those not covered in academic study. Thank you for always believing in me, while never putting pressure on me, and thanks for all the food! Finally, thank you Cam; you have made me an immensely proud brother your whole life and this feeling shows no sign of fading. You always have the most interesting ideas and I love listening to what you have to say.

Contents

1 Introduction **1**
- 1.1 Telescopes and spectroscopy . . . 1
  - 1.1.1 Spectroscopy . . . 2
- 1.2 Classical novae . . . 3
  - 1.2.1 Circumbinary material . . . 4
  - 1.2.2 Jets and accretion . . . 4
- 1.3 Thesis outline . . . 5
  - 1.3.1 Technological development . . . 5
  - 1.3.2 Science results . . . 5

2 Literature review: the state of the art **7**
- 2.1 Astrophysical transients . . . 7
  - 2.1.1 Supernovae . . . 7
  - 2.1.2 Kilonovae . . . 8
- 2.2 Cataclysmic Variables . . . 8
  - 2.2.1 Classical Novae . . . 8
  - 2.2.2 Recurrent Novae . . . 9
  - 2.2.3 Dwarf Novae . . . 9
  - 2.2.4 Intermediate polars . . . 10
- 2.3 Observations of classical novae . . . 10
  - 2.3.1 Optical light curves . . . 10
  - 2.3.2 Optical spectroscopy . . . 11
  - 2.3.3 Gamma Rays . . . 12
  - 2.3.4 X-rays . . . 13
  - 2.3.5 Radio . . . 13
  - 2.3.6 Remnants . . . 13
  - 2.3.7 Stardust and enrichment . . . 14
  - 2.3.8 Distances . . . 14

2.4 Circumbinary orbits 14
2.5 Jets and accretion 15

3 Observatories and detectives: from Aluminium to Silicon 17
3.1 Observatories 19
3.2 Instruments 23
  3.2.1 Aquila 23
  3.2.2 LHires 23
3.3 Gadget 23
3.4 Ramotswe 24
3.5 Endeavour 24
  3.5.1 My contribution to Endeavour 24
  3.5.2 Types of frame 25
  3.5.3 Gradient method 25
3.6 Poirot 25
  3.6.1 Poirot: fitting data 26
  3.6.2 Poirot: displaying data 27

4 Endeavour: refactoring the pipeline 29
4.1 Scope of Endeavour 29
4.2 Scope of AnotherOne 31
  4.2.1 Readability and PEP8 31
  4.2.2 Modularisation 31
  4.2.3 Batch to continuous processing 32
  4.2.4 Parallel code 32
4.3 AnotherOne: phases of Endeavour 32
  4.3.1 Scan 33
  4.3.2 Prepare 33
  4.3.3 Calibration Arc Matching 33
  4.3.4 Identify 34
  4.3.5 Extract 35
4.4 Performance 36
  4.4.1 Python 37
  4.4.2 iMac Pro 37
  4.4.3 GJW-DM 37
  4.4.4 Performance experiments E1 37
    4.4.4.1 Results E1 37

4.4.4.2 Discussion E1 38
4.4.5 Performance experiments E2 39
4.4.5.1 Results E2 39
4.4.5.2 Discussion E2 39
4.4.6 Galactic supernova 41

5 Characterising thermal breathing on CCD detectors 42
5.1 Motivation 42
5.2 Method 43
5.2.1 Finding the peaks 44
5.2.1.1 Top peaks 44
5.2.1.2 Top Peaks Fit 2 44
5.2.1.3 Threshold Filtering 45
5.2.2 Fitting the peaks 45
5.3 Problems 45
5.3.1 Halo effect 46
5.3.2 Undispersed light 46
5.3.3 Left side illumination 47
5.3.4 Median filter offsets 47
5.3.5 Defining the dispersion axis 48
5.4 Results 49
5.4.1 Variation of curvature coefficient 49
5.4.2 Number of peaks 49
5.4.3 Daily temperature profile 50
5.4.4 Exposure time 53
5.5 Pipeline upgrade 53
5.5.1 Current technique for CAM 55
5.5.2 Novel technique for CAM 55
5.5.2.1 Sampling dispersion axis 56
5.6 Moving forwards 56

6 Refinement of the gradient method in the low signal-to-noise regime 58
6.1 Understanding the problem 60
6.1.1 Periodic blockiness in faint spectra 60
6.1.2 Interference? 60
6.2 The role of the pixel gain map 60
6.3 Old method to correct parallel dependence 63

6.4 Old method to correct perpendicular dependence 63
6.5 New pixel gain map algorithm 64
6.6 Improvement in spectral quality 67

7 Machine learning for enhanced wavelength calibration 69
7.1 Scope of algorithm: understanding the problem 70
7.2 Old algorithm 72
  7.2.1 Search-and-Destroy 72
  7.2.2 Identification 73
  7.2.3 Feedback search 73
7.3 Limitations of the old algorithm 73
  7.3.1 Time complexity 74
  7.3.2 Slow operations 74
  7.3.3 Unnecessary repetition 74
  7.3.4 Failure rate 74
  7.3.5 Accuracy 74
  7.3.6 Manual identification 74
7.4 New algorithm 75
  7.4.1 Find peaks 75
  7.4.2 Screen peaks 76
  7.4.3 Refine 78
  7.4.4 Determine polynomial mappings 78
7.5 Advantages 81
  7.5.1 Accuracy 81
  7.5.2 Time complexity 81
  7.5.3 Automatic 81
  7.5.4 Speed 83
  7.5.5 Reliability 83
  7.5.6 Robust fitting 83
7.6 Limitations of the new algorithm 83
7.7 Results of a run 84
  7.7.1 Shape of dispersion axis 84
  7.7.2 Wavelength dependence on x 84

8 **High Resolution** **88**
8.1 The science case for high resolution 88
8.1.1 Confidence 89
8.1.2 Curiosity / digging deeper 89
8.1.3 Galactic Supernova 89
8.1.4 Slit-spectroscopy advantage 89
8.2 Implementation 90
8.2.1 Mechanical details 90
8.2.2 Empty slot 90
8.3 Software design 90
8.3.1 PEP8 93
8.3.2 Dependency injection 93
8.3.3 Strategy Pattern 93
8.3.4 Interface directly to Poirot 93
8.4 High-resolution algorithm 94
8.4.1 Identify 94
8.4.2 Dispersion 97
8.4.3 Extraction strategy 99
8.4.3.1 Summation method 100
8.4.3.2 Trimmed mean method 100
8.4.3.3 Gradient method 100
8.4.4 Final output 100
8.5 Slit spectroscopy - a bonus 100

9 **Classical Nova V906 Car: Discovery of circumbinary iron and oxygen** **103**
9.1 Introduction 105
9.2 Observations 106
9.2.1 Global Jet Watch observatories 106
9.2.2 Epochs of observations 108
9.2.3 Spectral features 109
9.2.4 Iron and Oxygen disc detection 110
9.2.5 Profile variations 111
9.3 Interpretation 111
9.3.1 Spectral shape 111
9.3.2 Disc model and fitting 111

9.4 Circumbinary discs 115
9.4.1 Theoretical predictions 115
9.4.2 Significance of timing 115
9.4.3 Significance of chemistry 116
9.4.4 Circumbinary discs in classical novae 116
9.4.5 Kepler argument for disc being circumbinary 117
9.4.6 Low-density hot gas 118
9.5 Possible alternative models 118
9.5.1 Bipolar outflow model 118
9.5.2 Pinwheel model 119
9.5.3 Outflowing disc model 120
9.5.4 Spectral comparison of models 120
9.6 Conclusions 121

10 The precessing jets of classical nova YZ Reticuli 123
10.1 Introduction 125
10.1.1 Ejecta of classical novae 125
10.1.2 Jets and accretion discs in novae 126
10.1.3 Classical nova YZ Reticuli 127
10.1.4 Classical nova V6568 Sgr 128
10.2 Observations 129
10.2.1 Global Jet Watch observatories 129
10.2.2 Aquila spectroscopy 130
10.2.3 LHires spectroscopy 130
10.2.4 Epochs of observations 130
10.2.5 Data reduction 130
10.2.6 Post-processing application 132
10.3 Fitting methods 132
10.3.1 Composite fits 133
10.3.2 Importance of coverage 135
10.3.3 Robustness 135
10.4 Results 136
10.4.1 Best-fitting model 136
10.4.2 Observed trends in derived parameters 138
10.4.3 Dynamical changes 141
10.5 Physical interpretation 141

10.5.1 Jets and accretion disc 141
10.5.2 Precession of the accretion disc and its jets 143
10.5.3 The behaviour of the central component 144
10.5.4 Why are radio jets not observed in classical novae? 146
10.5.5 Why is the inner pair not just a spherical shell? 146
10.5.6 Alternative models 147
10.6 Similar characteristics in V6568 Sgr 147
10.7 Conclusions 148

11 The jets phase of classical novae 150
11.1 Introduction 152
11.1.1 Classification schemes for classical novae 152
11.1.2 Jets and shell ejecta in novae 153
11.1.3 Classical nova V5668 Sgr 154
11.1.4 Classical nova V906 Car 155
11.1.5 Classical nova YZ Reticuli 156
11.1.6 A heterogenous sample 156
11.2 Observations 157
11.2.1 Observatories 157
11.2.2 Major campaigns 157
11.2.3 Data reduction 158
11.3 Timing phase transitions for V906 Car 159
11.3.1 Timing the high-velocity absorption in V906 Car 159
11.3.2 Timing the jets phase onset in V906 Car 159
11.3.3 Evolution of nitrogen and oxygen in V906 Car 162
11.4 Timing phase transitions for V5668 Sgr 163
11.4.1 V5668 Sgr: Initial jets and accretion disc 163
11.4.2 V5668 Sgr: Transition to nebular phase 165
11.4.3 V5668 Sgr: Nebular phase dominant 165
11.4.4 V5668 Sgr: evolution of nitrogen and oxygen 165
11.5 Overview: A dynamical model for nova phases 167
11.5.1 Onset of jets and accretion disc in V906 Car 167
11.5.2 An updated prescription for classical nova optical evolution . . 169
11.5.3 Jets, not clumps, responsible for H-alpha complexes 171
11.5.4 Implications for interstellar enrichment 171
11.5.5 Prediction for spatially resolved spectroscopy 172

11.6 Conclusions 172

12 Discussion **174**

12.1 Circumbinary discs around classical novae 174

12.2 Classical nova jets 176

12.3 Closing remarks 177

Bibliography **178**

List of Figures

3.1 Inside an Aquila spectrograph 18
3.2 The GJW-WA observatory 19
3.3 View from within the GJW-CL observatory 20
3.4 Inside the GJW-IN dome 21
3.5 The solar array at GJW-IN 22
3.6 The POIROT home page 26
3.7 The POIROT select target page 28

4.1 Profiling ENDEAVOUR: E1 experiments 38
4.2 Profiling ENDEAVOUR: E2 experiments 40

5.1 Halo effect 46
5.2 Undispersed light leak 47
5.3 Median filter and centroid calculation 48
5.4 Curvature of dispersion axis over time 50
5.5 Number of peaks in ThKr over time 51
5.6 Diurnal changes in ThKr offset parameter 52
5.7 Variation of offset with physical parameters 54

6.1 Blockiness in spectrum of YZ Reticuli 59
6.2 Beats interference along dispersion axis 61
6.3 Correction to wavelength dependence of flat-field 62
6.4 Illumination and distance from dispersion axis 65
6.5 Correcting for the cross-axis illumination profile of flats 66
6.6 Corrected spectrum of YZ Reticuli 68

7.1 Demonstration of a raw ThKr frame 71
7.2 Background subtraction for ThKr frame 76
7.3 Peak detection in one dimension 77
7.4 Clustering peaks on dispersion axis 79

7.5 Centre of mass for refinement of peak centroids 80
7.6 Matching peaks against reference wavelengths 82
7.7 Approximating the dispersion axis 85
7.8 Determination of wavelength dependence 87

8.1 The Perseus instrument selector 91
8.2 How the LHires instrument will be mounted on the Perseus instrument selector. 92
8.3 Two-dimensional LHires neon arc 95
8.4 Neon with emission lines identified 96
8.5 Two-dimensional LHires spectrum of YZ Reticuli 97
8.6 Automatically locating the dispersion axis 98
8.7 Determination of pixels to use for LHires extraction 99
8.8 Extracted one-dimensional spectrum of YZ Reticuli 101
8.9 Spatially-resolved high-resolution spectroscopy of Jupiter 102

9.1 Observational campaign for V906 Car 107
9.2 Characteristic circumbinary disc signatures 109
9.3 M-shaped spectral feature in V906 Car spectra 110
9.4 Three example spectra 112
9.5 Keplerian disc fit to M-shaped profile in V906 Car 114
9.6 Model comparison for M-shaped profiles 121

10.1 Light curve for YZ Reticuli 129
10.2 Observational campaign for YZ Reticuli 131
10.3 Multiple gaussian fit to YZ Reticuli H-alpha profile 133
10.4 Sequence of H-alpha profiles and fits in YZ Reticuli 134
10.5 Using LHires to support Aquila 137
10.6 Ideogram showing anti-correlated velocity shifts in emission components 138
10.7 Time evolution of velocity separation of emission line pairs 139
10.8 Upper energy level dependence for split lines 145
10.9 Confirmation of spectroscopic jets in V6568 Sgr 148

11.1 Light curves for V5668 Sgr, V906 Car and YZ Reticuli 154
11.2 Timing the appearance of high-velocity absorption in V906 Car . . . 160
11.3 Timing the onset of emission in V906 Car 161
11.4 Disappearance of nitrogen triplet in V5668 Sgr 162
11.5 Sequence of Gaussian fits to V5668 Sgr spanning 2000 days 164

11.6 High-excitation coronal iron line in V5668 Sgr 166
11.7 Jets and accretion disc fits to V5668 Sgr, V906 Car and YZ Reticuli . 168
11.8 Time progression of jets and disc model for three classical novae . . . 170

Chapter 1

Introduction

For millennia, people have been fascinated by the cosmos. There is a certain magic that transcends human barriers in the fact that anyone can cast their gaze upwards and see the same constellations, etched into the sky, seemingly unchanging across centuries. However, this is not strictly true — everybody who looks up sees it differently, according to the weather, the season, the time of night, and their location. The constellations visible from Australia are different from those seen above Oxford, while astronauts in space have the ultimate vantage point. Beyond simple changes in how and where we conduct our observations, there is also a multiplicity of natural phenomena which can occur that subtly alter how the sky appears. This thesis explores one category of such events, called *classical novae*, and this chapter sets out the broader context of this area of astronomy for the lay audience.

1.1 Telescopes and spectroscopy

While features of the night sky are constantly changing, at first glance it does seem remarkably similar from one night to the next. To understand and describe the intricate changes that are happening, it was necessary to improve upon the human eye. Some light arriving from space is too faint to be seen, but fortunately it is possible to design machines sensitive enough to detect these signals. A key turning point in the study of the universe beyond our home planet came in 1610, when Galileo Galilei pointed a telescope at Jupiter, discovered four of its moons (Io, Europa, Ganymede and Callisto) and showed that they must be in Jupiter's orbit. Telescopes typically use an intricate system of mirrors and lenses to create magnified images of distant objects, with large collecting areas far bigger than a human pupil.

Modern telescopes do more than just enlarge images to be inspected through an eyepiece — they also capture the images, recording the data for analysis on a rainy day. This is made possible by a separate invention, the camera. Early cameras recorded telescopic observations on film, but since the advent of Charged Coupled Device (CCD), and later complementary metal–oxide–semiconductor (CMOS), detectors, it has become possible to save observations directly to a computer. These detectors are made up of a grid of pixels, each of which can count how much light lands in them. Control software opens the shutter, and light is directed onto the detector for a set duration (exposure time), and at the end, a computer records how many counts were associated with each pixel. This has considerable implications – whilst the human eye only counts light for less than a second, it is possible to leave a detector recording light for upwards of an hour in an observatory! The upshot of this is that the final image is deeper, cleaner and more detailed than the eye could see without assistance.

The frequencies (colours) of light that we are capable of seeing are restricted to a rather limited subset of those which occur in nature; imagine a musical note too high-pitched to be heard. While this range is ample for our everyday needs, it means the human eye would miss out on a lot of the details surrounding the extraordinary phenomena found in the depths of the Universe. Why is frequency useful in astrophysics? To continue the auditory analogy, imagine the sound of an approaching vehicle. Due to the so-called *Doppler shift*, while the vehicle approaches the receiver, the pitch is higher, but when it passes and recedes, the pitch is lower. Light also exhibits a Doppler shift, and by measuring the precise frequency of detected photons (light), it is possible to deduce whether a particular source is approaching the Earth, or moving away, and at what speed.

With modern telescopes, we are now able to examine and interrogate the universe in a number of ways which were previously unimaginable. Exploration of data from such instruments has uncovered a wealth of astronomical targets exhibiting a rich variety of behaviours. To truly understand how the natural world behaves in such remote systems would be impossible; the best we can do is to piece together evidence found using different examples, wavelengths, and ways of thinking.

1.1.1 Spectroscopy

A camera's job is simple — to count how many photons land in each pixel. For a standard image, the light detected in the leftmost pixels corresponds to what was in the left of the visual scene, and vice versa for the right. In spectroscopy, however, the light is split up (diffracted) according to its wavelength before it reaches the detector, and so the pixels on the left might count the low-wavelength light while the pixels on the right capture the high-wavelength light. This can be valuable in an astrophysics context, where it is useful to know how much light of each particular wavelength was emitted by the target of interest relative to that at other wavelengths. Atoms under certain conditions can emit photons of very precisely known 'rest' wavelength and hence frequency, known as *emission lines*; measuring the photon counts at wavelengths either side of the rest wavelength can therefore reveal how the object that emitted the line was moving, by considering its Doppler shift. A blue-shifted line, for example, would manifest itself as a lot of photons very slightly bluer (shorter wavelength) than the rest wavelength, and we would infer that the source was moving towards the camera. The remainder of this thesis relies heavily on this principle.

1.2 Classical novae

When looking at the night sky, very occasionally it is possible to see something that nobody has ever seen before. A *nova* manifests itself as a bright, apparently 'new', star, hence the name. By close examination, astronomers have deduced that these events are typically caused by the sudden brightening of existing systems, by several orders of magnitude. These events are subdivided into several categories depending on the nature of the brightening, and the potential causes, but the type that this thesis is primarily concerned with is known as a classical nova. Other types include supernovae, and hypernovae, but these have distinct underlying mechanisms, occur in substantially different circumstances, and release vastly disparate amounts of energy.

While our Sun is not accompanied by another star, the majority of stars in the Milky Way are thought to exist in *binary systems*, whereby two stars both orbit their common centre-of-gravity. Classical novae occur exclusively in binary systems where one star is a white dwarf (a small, compact remnant of a star that has exhausted its fuel for nuclear burning), and the other is not. When the two stars get close enough, the strong gravity of the compact white dwarf disrupts the *secondary* (companion) star, causing hydrogen gas to flow from the secondary onto the surface of the white dwarf. This process is known as *accretion*. The hydrogen at the bottom of the accreted layer heats up, until it reaches the critical temperature at which hydrogen atoms can merge together to create helium, similarly to the nuclear fusion reaction that powers the sun. While the sun is stable, on the surface of such a white dwarf the material heats up, which accelerates the reaction, leading to even more heat being produced. This process, known as *thermonuclear runaway*, causes an eruption, and the corresponding increase in brightness is sometimes sufficient for it to be seen with the naked eye.

The new data analysed in this thesis span four recent classical novae. These are: V5668 Sgr, which was discovered in March 2015; V906 Car, which was discovered in March 2018 (they nearly share a birthday!); YZ Reticuli which was discovered in July 2020; and V6568 Sgr, also discovered in July 2020 (not to be confused with V5668 Sgr, despite their similar names).

Through careful and incremental studies over the last few centuries, we have learned a tremendous amount about classical novae. This thesis uses densely time-sampled sequences of observations to shed new light on the nature of these eruptions and the events which immediately follow them.

1.2.1 Circumbinary material

Much like the planets, comets, and various asteroid belts which orbit the Sun in our solar system, stars contained in binary systems may also have particles orbiting them. Matter situated in stable orbits around such a pair of stars, beyond the radius of the inner binary, is known as *circumbinary*, as opposed to *circumstellar*, which refers to matter in orbit around a single star. One of the key findings of this thesis is that there is evidence that the V906 Car system may have circumbinary material.

1.2.2 Jets and accretion

Accretion is when matter falls in towards a central mass such as a star. In everyday life, gravity makes things fall down, but this simple picture does not suffice to explain orbital dynamics. The moon, for instance, does not fall down to the Earth, nor does the Earth fall towards the Sun. This is because of *angular momentum*, which characterises the tendency of objects to keep on spinning (the rotational analogue of linear momentum, which is the tendency of things to continue to move in a straight line). In order to fall towards the white dwarf, hydrogen must somehow lose angular momentum, and it typically does this in a flattened disc, or *accretion disc*. Here, hydrogen may fall to progressively lower orbital radii, spiralling inwards down the accretion disc, as it can lose angular momentum through friction. A remarkable feature of accretion discs is that they are often associated with the launch of so-called *jets* — the high-speed collimated outflow of gas. This thesis lays out the evidence for the existence of such jets arising following eruptions in all four of the classical novae studied (V5668 Sgr, V906 Car, YZ Reticuli, V6568 Sgr).

1.3 Thesis outline

Chapter 2 gives an overview of the literature surrounding the classical nova phenomenon, laying out the existing foundations for the new work presented in this thesis. The remainder of this thesis can then be thought of as roughly splitting into two parts — Chapters 3 to 8 discuss the technology that underpins the capturing, calibration, and correction of data from a computational perspective, while Chapters 9 to 11 present the main science results that came about through the careful analysis of those data.

1.3.1 Technological development

The technology chapters focus on a suite of new software packages which I designed and implemented. Chapter 3 provides details on the observatories used to capture data on the four targets, and introduces the various software packages that have been written to help with the collection, processing, and visualisation of data. Chapter 4 discusses my development of the main software package responsible for the data processing step. Chapter 5 describes an investigation I carried out into the variability of certain elements of the wavelength calibration step, and the resulting upgrade to the data reduction code. Chapter 6 discusses an enhancement to the extraction process for very faint observations, which was necessary for processing early spectra of YZ Reticuli. Chapter 7 discusses an enhancement to the algorithm for calibration discussed in Chapter 5, dramatically improving the robustness and speed of processing using machine learning techniques. In July 2020, an additional instrument was added to capture high resolution spectra; I designed and wrote a new package which extracts these data, and report on this in Chapter 8.

1.3.2 Science results

Chapter 9 is closely based around a paper, published as McLoughlin et al. [2020], which presents the discovery of circumbinary iron and oxygen around the classical nova V906 Car. That paper uses the name Nova Carinae 2018, another name for V906 Car. Chapter 10 is based on a second published paper, McLoughlin et al. [2021a], which reports the existence of jets emanating from classical nova YZ Reticuli, and uses data for V6568 Sgr to support the hypothesis of jets in classical novae. Chapter 11 follows as a direct sequel, and applies the jet model to V906 Car and V5668 Sgr, finding evidence that they too show signatures of jets being launched, and that therefore this may be an ingredient in the phenomenon of classical novae in general. I also propose new observations that may elucidate the issue, which could be achieved using telescopes and instruments that are already available.

Chapter 2

Literature review: the state of the art

This chapter sets out the observational foundations for the study of classical novae, describing what is already known and laying out the current paradigm. As with any active scientific field, the theories are not fixed, instead being continually tested with new observational data, often resulting in the modification and refinement of existing ideas. This chapter therefore presents the prevailing understanding, but does not claim that this is The Truth.

Section 2.1 gives a brief outline of some of the more energetic astrophysical transients. Section 2.2 describes the different types of novae in the context of cataclysmic variables. Section 2.3 covers what is known about classical novae through various avenues, including direct imaging, spectroscopy and physical evidence in the form of meteorites. Sections 2.4 and 2.5 give the theoretical bases for circumbinary discs and jets in classical novae respectively.

2.1 Astrophysical transients

The sky is not fixed; there is a veritable zoo of dynamic and evolving astrophysical phenomena, including, e.g. supernovae and kilonovae. This section briefly explores the characteristics of these other types of transient events, comparing them with classical novae.

2.1.1 Supernovae

Supernovae are violent stellar explosions that result in the collapse or total destruction of the progenitor star, unlike classical nova eruptions, which leave the host white dwarf reasonably intact. The luminosity of a supernova at its peak can temporarily outshine the entire host galaxy, despite only emitting a very small fraction of the energy output as electromagnetic radiation. There are many different names for the various types of supernova (Type Ia, Ib, Ic, II, II-L, II-P, IIn, IIn (bright),...), which depend on either the spectral features or the nature of the progenitor system.

2.1.2 Kilonovae

Kilonovae are a class of astrophysical transients with luminosities around a thousand times larger than the Eddington luminosity for a solar mass star, first described by Metzger et al. [2010], who associate them with being the electromagnetic counterpart of merging compact objects. Observationally, these have far higher luminosities of electromagnetic radiation than classical novae, which are typically approximately at the Eddington limit of luminosity but sometimes exceed this temporarily [Bode and Evans, 2008].

2.2 Cataclysmic Variables

Classical novae belong to the category of flaring binary systems known as cataclysmic variables, which contain a white dwarf as the primary, and a less evolved star as the secondary. The stars are often close enough together that the gravity of the primary can significantly distort the tidally-locked secondary star, causing matter to flow from the secondary (sometimes referred to as the 'donor') to the primary in what is known as Roche-lobe overflow (RLOF), although not all cataclysmic variables accrete like this. The closeness of the binary results in very short orbital periods, which typically range from 80 minutes to 8 hours, albeit with a notable absence of periods between two and three hours, known as the period gap and thought to be caused by enhanced accretion rates [Howell et al., 2001]. This is explored in greater depth in Chapter 9, which provides a possible explanation for this phenomenon in the case of classical nova V906 Car.

The mass and mass accretion rate of the white dwarf heavily[1] influence the evolution of the system, leading to qualitatively different behaviour. This section discusses the most significant subtypes and describes the distinctions between them.

2.2.1 Classical Novae

Classical novae, the main topic of this thesis, have relatively high white dwarf masses of typically $0.5\,M_\odot < M_{\rm WD} < 1.1\,M_\odot$ and accretion rates of approximately $\dot{M} = 10^{-9}$

---

[1] no pun intended

to 10^{-10} $M_\odot\,\mathrm{yr}^{-1}$ [Warner, 1995]. There is a hard upper limit; if the mass of the white dwarf were ever to exceed the Chandrasekhar limit of $M_{\mathrm{WD}} = 1.4\,M_\odot$, then it would collapse and explode as a Type-Ia supernova. As matter accumulates on the degenerate surface of the white dwarf, the pressure at the bottom of this layer increases and it heats up, eventually reaching the critical conditions at which the (catalytic) CNO cycle initiates. This produces more heat, leading to thermonuclear runaway. The accreted mass required for thermonuclear runaway is inversely proportional to the mass of the white dwarf, since the increased surface gravity of a higher mass white dwarf results in higher pressure at the base of the accreted layer. Since classical novae eruptions, unlike supernovae, fully disrupt neither the white dwarf nor the secondary, it is possible for these events to recur. However, a classical nova is defined to be a nova in a system which has not previously been observed to erupt, and so are presumed to have longer recurrence times than the few hundred years of scientific observations.

2.2.2 Recurrent Novae

Recurrent novae are cataclysmic variable systems which have been observed to erupt multiple times. Their rapid recurrence is due to higher white dwarf masses of typically $M_{\mathrm{WD}} > 1.3\,M_\odot$ and correspondingly high mass transfer rates of up to $\dot{M} = 10^{-7}$ $M_\odot\,\mathrm{yr}^{-1}$ [Warner, 1995]. Chapter 10 presents new evidence of classical novae exhibiting jets, drawing a link between YZ Reticuli, V6568 Sgr, and the class of recurrent novae. Jets are discussed further in Section 2.5.

2.2.3 Dwarf Novae

Dwarf novae are ‘eruptive variable stars exhibiting repetitive outbursts with a typical amplitude of 2 - 5 mag with a typical recurrence time of order of a few weeks to months’ as defined by Osaki [1997], and have an entirely different mechanism underpinning their significantly weaker flaring episodes compared to classical novae or recurrent novae. Instead of being powered by thermonuclear runaway, they are thought to be caused by accretion disc instabilities [Osaki, 1997]. Sokoloski et al. [2006] present the first evidence for a ‘combination nova’; an event which displays characteristics of both a dwarf nova and a classical nova. In that case, it was deemed that the accretion disc instability initially responsible for the dwarf nova event probably acted as the trigger for the subsequent classical nova eruption.

2.2.4 Intermediate polars

These are cataclysmic variables with medium-strength magnetic fields. These systems, which account for $< 1\%$ of cataclysmic variables [Downes et al., 2001], have an accretion disc which is disrupted at lower radii, with infalling matter redirected onto magnetic field lines. These typically have spin periods that are faster than their orbital periods, and example novae that are classified as intermediate polars include DQ Her and GK Persei.

2.3 Observations of classical novae

This section discusses the broad variety of ways in which classical novae have been studied, from the electromagnetic radiation they emit to the meteoritic evidence that ancient novae have conveniently deposited on Earth.

2.3.1 Optical light curves

Careful analysis of the light curves of optical observations can reveal important information concerning the eruption. The typical classical nova light curve very rapidly approaches a maximum, which is many orders of magnitude brighter than the quiescent object. It then declines in brightness, taking several years to fade. Strope et al. [2010] present a large collection of 93 classical nova light curves, categorised depending on the features they reveal, and with a canonical example for each. The categories consist of: smooth, plateau, dust, cusp, oscillations, flat-topped, and jitter. The categories are useful for conveying qualitative information quickly; for instance, Chapter 10 discusses the survival of the accretion disc in YZ Reticuli in the context of its plateau-type light curve.

An important observable derived from the light curve is the decline time. The most commonly referred-to characteristic numbers for light curves of classical novae are t_2 and t_3, the times taken for the magnitude to decline from the peak by two and three magnitudes respectively. A classical nova is then classified on the scale originally set out by Gaposchkin [1957], which ranges from ‘very slow’ ($t_2 > 150\,\mathrm{d}$) to ‘very fast’ ($t_2 < 10\,\mathrm{d}$). While this simple definition works well for classical novae with a smooth light curve, it proves more challenging for novae with oscillations and dust dips such as DQ Her or V5668 Sgr.

2.3.2 Optical spectroscopy

Optical spectroscopy gives key insights into the dynamics of classical novae, as it provides detailed information on the species, ionisation state, and dynamical behaviour of distinct emitting or absorbing regions near a classical nova. Many authors over the years have contributed to our understanding of the progression of the classical nova phenomenon by describing various schemes of phases derived from optical spectroscopic measurements, and a thorough exposition on these phases is given in the excellent review paper Della Valle and Izzo [2020], which summarises the established consensus. Here, I present a summary of the main phases as they pertain to this thesis, leaning extensively on the pioneering works of McLaughlin [1956], Gaposchkin [1957], Shore [2012].

Before the light curve reaches optical maximum, spectra of classical novae typically exhibit the 'pre-maximum' or 'fireball' phase, which is characterised by very blue-shifted 'P-Cygni' [Beals, 1953] absorption with Doppler speeds of around $2000\,\mathrm{km\,s^{-1}}$ to $3000\,\mathrm{km\,s^{-1}}$. This is typically associated with the optically thick expanding envelope of matter that previously made up the accreted layer on the surface of the white dwarf. While it was previously difficult to capture classical novae on the rise to maximum brightness, in recent years, the development of automated transient surveys which give rapid alerts when a new potential nova erupts has been critical in allowing observatories to get valuable data on the early evolution of classical novae. Chapter 8 discusses the addition of a new high-resolution instrument to the Global Jet Watch observatories, and how the software package which I wrote for it was designed to enable rapid turnover of spectra and to allow early fine-tuning of observational parameters. This will help guide future Global Jet Watch campaigns to capture high-quality data of the initial phases of classical novae.

The next phase is the 'early decline', during which time the light curve has peaked and begins to decay. By this point, the effective photosphere recedes as the ejecta expands. Multiple absorption systems with high blue-shifted velocities are typically present at this time. Known as 'diffuse-enhanced' or 'high-velocity component' (HVC) [Arai et al., 2016], recent high-resolution spectroscopy of the HVC of ^{7}Be in V5668 Sgr revealed that classical novae may be significant contributors to the Galactic stellar ^{7}Li [Molaro et al., 2016], discussed further in Section 2.3.7. Williams et al. [2008] speculated that transient heavy-element absorption (THEA) lines might have been formed in circumbinary material, which is concordant with the discovery of circumbinary iron and oxygen presented in Chapter 9. Aydi et al. [2019] proposed that the arrival times of new absorption systems in classical nova ASASSN-17pf (LMCN 2017-11a) were correlated with flares in the light curve. A convincing unified picture of the HVC absorption has not yet appeared, and this remains an area of active research. Strong emission lines can appear during this phase, and Chapter 11 presents the most precise timing data to date for the onset of such lines, for the case of V906 Car.

The next phase is the 'transition' phase, during which time the classical nova ejecta transition from being optically thick to optically thin, and the absorption systems disappear. During this time, the 'Orion' spectrum arises, as the now rarefied ejecta are irradiated by the hot white dwarf. These lines typically have different radial velocities from those observed during previous phases.

With continued expansion and the resultant decrease in density, the emitting regions of late-time classical novae approach the so-called 'nebular' phase, during which the spectra begin to resemble those expected of a planetary nebula. The spectrum is dominated by forbidden transitions, although Balmer lines of hydrogen can still be present in emission. Chapter 11 examines the spectra of V5668 Sgr, tracing various components of the emission for over two thousand days, clearly showing the onset of the nebular phase.

Williams [1992] prescribes a classification system for novae, depending on which of two groups of emission lines is more prominent in optical spectra. Under this scheme, a classical nova may be classified as being of either He /N or Fe II spectral type. They further proposed that the Fe II lines are formed in the wind, while the He /N lines are formed in the ejected shell, and showed that in rare cases, an object may be further designated as a 'hybrid', exhibiting both systems at different points in the evolution. Chapter 11 makes the point that the discovery of jets in all four of the targets (covering all spectral types) examined in this thesis implies that jets are a ubiquitous feature of novae.

2.3.3 Gamma Rays

The original presentation of four classical novae detected in gamma-rays (hereafter γ-rays) by Ackermann et al. [2014] established that classical novae were capable of generating high-energy particles, most likely through some kind of shocks. Li et al. [2017] found correlations between the optical light curve and γ-ray detections, concluding that reprocessed shocks are the source of most of the optical emission from classical novae, and not the white dwarf. Martin et al. [2018] showed that these shocks most likely arise as internal interactions between various parts of the ejecta. This is important to Chapter 10, based off McLoughlin et al. [2021a], which proposes that there are jets emanating from all four of the classical novae analysed for this thesis; a high-speed jet that starts days after the fireball would pierce it, providing the location of these internal shocks and potentially explaining the γ-ray evidence.

2.3.4 X-rays

The X-ray emission of classical novae has two distinct modes, hard and soft (defined by Schwarz et al. [2011] as $1 - 10\,\mathrm{keV}$ and $0.3 - 1\,\mathrm{keV}$ respectively). While the hard X-rays are thought to arise from shocks similarly to γ-rays, the soft component arises from continued nuclear processing of a thin layer of hydrogen on the white dwarf surface, which did not acquire escape velocity, instead re-establishing hydrodynamic equilibrium. The turn-off of this 'super-soft source' phase is associated with the exhausting of surface fuel, which happens quicker for classical novae with more massive white dwarfs [Schwarz et al., 2011]. Classical novae are also expected to emit a flash of X-rays lasting only hours during the fireball phase, before becoming engulfed by the advent of an optically thick envelope. However, this has not been observed (to date) due to the observational difficulty of this event preceding the light curve rise to maximum.

2.3.5 Radio

Radio observations such as the ALMA measurements of V5668 Sgr presented by Diaz et al. [2018] have confirmed the existence of roughly spherical shell outflows surrounding classical novae. Sokoloski et al. [2008] presented evidence of highly collimated outflows and lobes in radio images of the recurrent nova RS Oph using the VLBA. However, radio jets have not been seen in classical novae to date. Chapter 11 discusses the practicalities of observing jets in the radio, and why they would not necessarily be strong emitters at long wavelengths despite spectroscopic evidence for line-emitting jets.

2.3.6 Remnants

The first examples of classical novae that were originally spectroscopically studied over a century ago are now observable as mature remnants. While most classical nova remnants show the relatively slow-moving approximately spherical ejecta [Bode, 2004], bipolar outflows have also been observed [Chesneau et al., 2012]. It is challenging to observe detailed structure, because most classical novae have faded significantly by the time they are spatially resolved.

2.3.7 Stardust and enrichment

Through detailed isotopic analysis of the Murchison meteorite, Amari et al. [2001] discovered five SiC grains and one graphite grain which were formed in the ejecta of an ancient classical nova containing a massive ONe white dwarf. This has implications for the possibility of enrichment of the Milky Way; while supernovae may be more commonly thought of as the producers of stardust, classical novae do still have a role to play in the processing and distribution of certain isotopes. V5668 Sgr, one of the classical novae discussed in this thesis, was found to be a strong producer of dust [Banerjee et al., 2015]. Significant levels of ^{7}Be were detected in its spectrum, which decays to ^{7}Li, meaning classical novae might account for all of the Galactic stellar ^{7}Li. The discovery of jets for such a high-yielding system has implications for its ability to launch grains into interstellar space, discussed in Chapter 11.

2.3.8 Distances

Searches for a relation between maximum magnitude and rate of decline (MMRD) for classical novae are neither predicted nor observed to give small scatter, meaning that it is unlikely classical novae can be used as standard candles with as much success as their more powerful relatives, supernovae [Shara et al., 2017]. Distances to classical novae were often calculated using colour and Galactic extinction maps, which historically led to some targets having multiple and mutually inconsistent values for the distance. Recently this has been somewhat alleviated, as most new classical novae can be identified with a progenitor star from the GAIA catalogue [Prusti et al., 2016, Brown et al., 2018] which contains very accurate distances. Another relatively recent advance in the estimation of classical nova distances has been the discovery of multiple novae in galaxies other than the Milky Way.

2.4 Circumbinary orbits

The existence of three families of stable circumbinary orbits was shown by Doolin and Blundell [2011], using n-body simulations. As explained in McLoughlin et al. [2020], the paper which Chapter 9 is based on, stable means not that the material is confined to closed orbits whose orbital plane is unchanging through time, but rather orbits that stably precess about the binary star system indefinitely. Doolin and Blundell [2011] reported that the rate of precession depends on both eccentricity and mass ratio in the sense that the more eccentric a system, or the more skewed the mass ratio is, the shorter the period of precession of the circumbinary orbital plane.

Given that such orbits are possible and stable, the natural assumption is that they probably exist somewhere in the vast universe. Blundell et al. [2008] presented the discovery of the circumbinary disc around the microquasar SS433, and Doolin and Blundell [2009] reported on its precession. While the ejecta of supernovae are launched at speeds far in excess of the escape velocity of the system, much lower speeds are observed in classical nova shells and accordingly not all the material is given enough energy to become completely unbound. With repeating cycles of classical novae, it is not unreasonable to imagine that there may be material from previous eruptions still orbiting the system on circumbinary trajectories. Chapter 9 presents the first direct evidence of such a pre-existing reservoir of circumbinary matter in emission around a classical nova, by fitting to an unusual and transient emission profile found in spectra of the 2018 outburst of V906 Car. As mentioned in Section 2.3, Williams et al. [2008] previously associated transient heavy element absorption in early spectra of classical novae with being formed in circumbinary material.

2.5 Jets and accretion

Chapter 10 presents the discovery of spectroscopic jets from the classical nova YZ Reticuli. Although jets are observed at a huge range of astronomical scales, from supermassive black holes to young stellar objects (YSOs), microquasars, planetary nebulae, pulsars and even recurrent novae, there are very few images of classical novae showing jets. The existence of jets however, as defined by Bridle and Perley [1984], does not inevitably imply the presence of synchrotron radio jets. For classical novae, the absence of strong magnetic fields like those measured around SS433 may provide the key to understanding why jets are not observed emanating from novae in the radio. A more in-depth discussion of the existing body of evidence surrounding jets in recurrent (as distinct from classical) novae is provided in Section 10.1.2, which sets the context for the discoveries in classical novae presented later in that chapter, and in Chapter 11.

The discovery of jets in classical novae should not come as a surprise — astrophysical jets are intrinsically tied to accretion, which is an essential part of the classical nova paradigm. Jets were previously thought not to occur following classical novae because computational models such as those presented in Figueira et al. [2018] suggested that accretion discs are sometimes disrupted in the immediate aftermath following an eruption. In fact, Retter [2006] proposed that it is only in the particular sub-case of intermediate polars, which make up only a small minority of classical novae systems, that the accretion process is disrupted. Considering the wide variety of jet-launching systems in astrophysics, and the probable survival of the accretion disc, jets should be keenly expected to occur in classical novae.

Chapter 3

Observatories and detectives: from Aluminium to Silicon

This chapter explains the capture and subsequent processing of the main data used throughout this thesis.

The Global Jet Watch comprises a set of five observatories, separated around the world in longitude so that one is always in darkness. These observatories are mostly located in residential schools, where students use the imaging instruments on the half-metre diameter telescopes to capture fascinating objects in the early night sky before local bedtime. For the rest of the night-time, the primary science time-lapse data are collected, using the 'Aquila' spectrographs (pictured in Figure 3.1). I was not involved in the building of the observatories or their various instruments, but I describe them briefly in Sections 3.1 and 3.2 respectively, to give context to the nature of the data they capture.

I have contributed significant development to each of a series of bespoke Python packages, which collectively operate the telescopes and process the data from them. Each of the packages in the Global Jet Watch ecosystem is named after a detective, and the hand-off between these independent packages is carefully managed so that they operate together smoothly.

Section 3.1 details the different observatories and explains their similarities and idiosyncrasies. The next sections give an overview of the different pieces of software that get the data from the sky into reduced form as follows: Section 3.3 describes GADGET, the code responsible for remote operation of the telescopes; Section 3.4 discusses RAMOTSWE, which is responsible for checking the metadata of files coming from the observatories and ensuring self-consistency; Section 3.5 provides only a brief first introduction to the main data-reduction code, ENDEAVOUR, which I contributed significant development to, explored fully in Chapters 4, 5, 6 and 7; Section 3.6

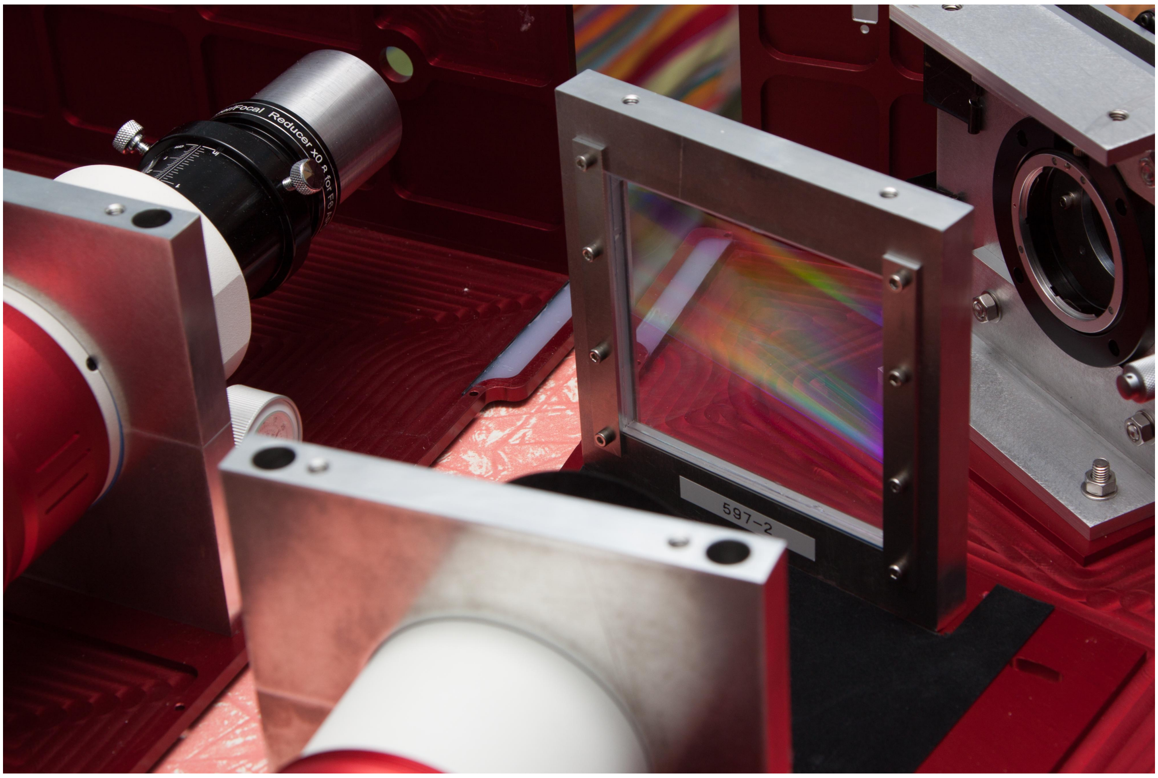

Figure 3.1: Inside Aquila spectrographs, showing the Volume Phase Holographic (VPH) grating on the right, and the collimator of another one on the left. This photo was taken by Steve Lee, who designed and built the spectrographs.

Figure 3.2: The GJW-WA observatory during evening twilight, co-starring the moon and Venus. This photograph was taken by K. M. Blundell.

covers the POIROT package, which is responsible for the post-processing and fitting of one-dimensional spectra, the main data product of astrophysical interest, as well as providing key visualisations.

3.1 Observatories

The five observatories each have a short code of two letters based on location, and a full code which is simply "GJW-" followed by the short code. The locations, and their short codes are as follows: South Africa (SA), Chile (CL), east Australia (OZ), Western Australia (WA) and India (IN). Throughout this thesis, these may be referred to by either their name, their full code, or their short code. The GJW-WA observatory is pictured in Figure 3.2. Figure 3.3 shows an example of the GJW-CL telescope within its observatory, and Figure 3.4 shows the view of the inside of the dome at GJW-IN.

The observatories, like all physical experiments, have their own peculiarities — the windiest observatory is GJW-SA, while GJW-WA has a relatively weak internet

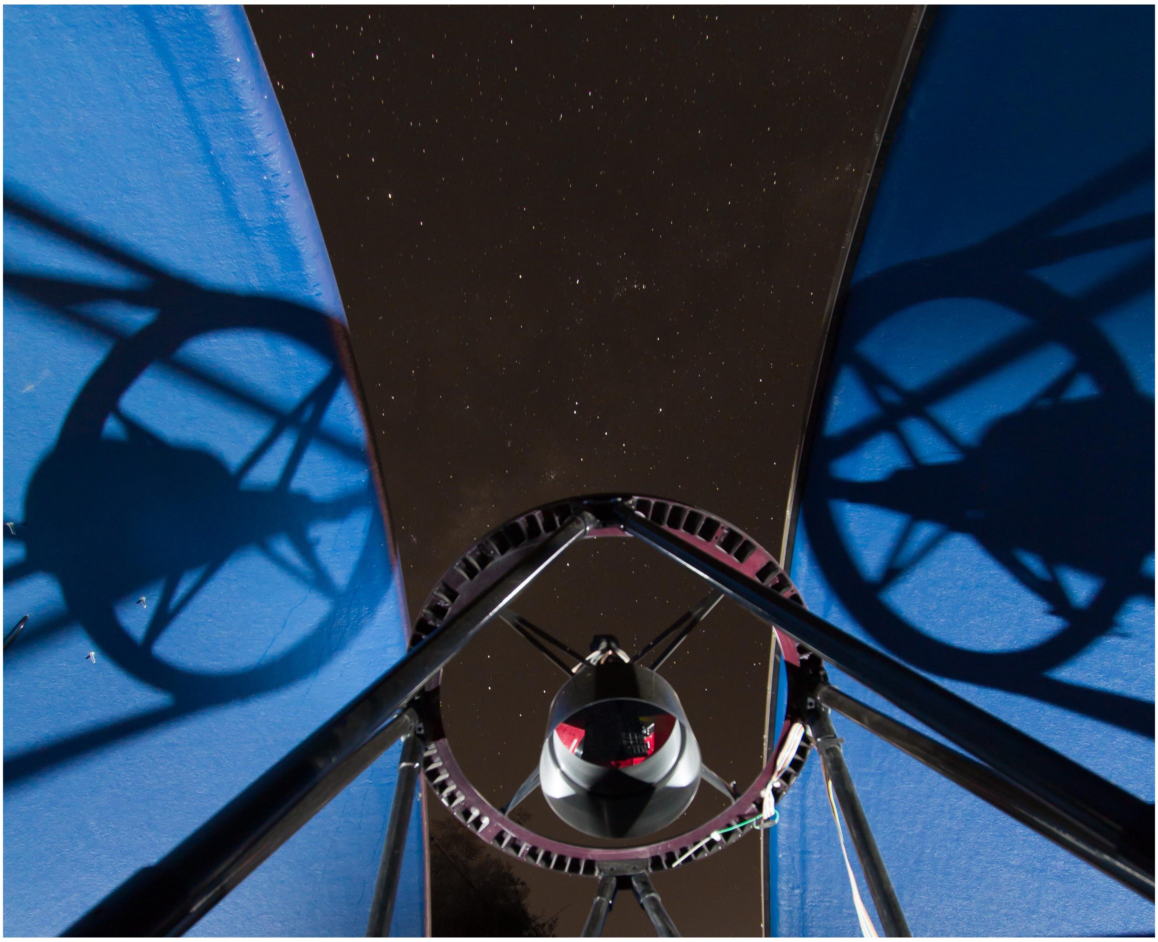

Figure 3.3: GJW-CL top end; the skyward view from within the observatory. This photograph was taken by K. M. Blundell.

Figure 3.4: View of the inside of the GJW-IN dome. The red light (switched off during observations!) helps to maintain dark adaption in the operator's eyes. This photo was taken by Steve Lee.

Figure 3.5: The solar array on the school roof next to the GJW-IN telescope. This photograph was taken by K. M. Blundell.

connection. The GJW-IN observatory is the only one situated in the northern hemisphere, and so is better placed for certain targets. Due to issues with electrical grid reliability in India during 2016 and 2017, a solar farm was built and commissioned, depicted in Figure 3.5, along with battery storage. The GJW-CL observatory is the most prone to earthquakes, and has the highest temperature swings, in part because it is so dry.

The geographical location of the telescopes is intrinsically linked to the main observational pattern; being spread around the world by longitude means that it is possible to achieve continuous and sometimes overlapping observations of rapidly varying astrophysical targets. This was particularly the case for classical nova V906 Car, which is the main subject of Chapter 9. The main data products from the Aquila spectrographs are densely time-sampled spectroscopic sequences of a selection of Galactic targets, some of which (e.g. V5668 Sgr, SS433, etc) have been observed for over five years so far. Sustained coverage of classical novae over different timescales unlocks an understanding of each target as it evolves through particular phase changes (the main topic of Chapter 11).

3.2 Instruments

There are several instruments in place at the Global Jet Watch observatories, with new capabilities being added. For this thesis, the two most important sets of observations come from the Aquila and LHires instruments.

3.2.1 Aquila

The Aquila spectra have a spectral resolution of $R \sim 4000$ according to the standard definition of $\lambda/\Delta\lambda$, and span the wavelength range from approximately 5000 Å to 8300 Å, although this differs slightly from one observatory to another and across time, with the advent of new, larger detectors. The camera in each Aquila spectrograph is either a FLI ML16200 (4500×3600 pixels, 6.0 µm pixel size) or a FLI ML8300 (3326×2504 pixels, 5.4 µm pixel size), and each is cooled to $-30\,^{\circ}$C by a Peltier cooler. In common with many modern spectrographs, the Aquila spectrograph is fibre-fed because it would be prohibitive to mount such a heavy instrument on the back of a 0.5-metre telescope. The area of the fibre projected on the sky is 3.3 arcseconds. Further details of the instrument will be presented in a forthcoming paper (Lee et al, in prep). The main purpose of the Aquila data is to provide sustained and densely time-sampled medium resolution spectroscopy.

3.2.2 LHires

As explained in more detail in Sections 8.1 and 8.2.1, the LHires instrument is a relatively new addition, on my recommendation. Due to global circumstances, it has been difficult to travel to the Global Jet Watch observatories and so it is currently housed at the Rainbow observatory, but will be permanently mounted at GJW-OZ. The main purpose of the LHires data is to supplement Aquila with detailed high-resolution ($R \sim 20,000$) observations, confirming the time-series analysis. In the fullness of time, it will also have robotic remote control of the central wavelength, enabling detailed checks of specific spectral complexes as these evolve in targets of interest.

3.3 Gadget

The automation and robotisation of the Global Jet Watch observatories represents a significant step, and replaces what was previously manual work. I built GADGET as a modular Python package when the Global Jet Watch observatories acquired Digital

Loggers IP Controlled High Power DC switches, allowing remote control even in the event of a catastrophic failure of the observatory PC (important in case of e.g. an electrical storm, during which the observatory must close). It handles a wide variety of set sequences of events, from setting up the observatory for the schoolchildren to conduct their imaging observations, to closing up the observatory at the end of each night, or setting up the spectrographs to record calibration lamps.

3.4 Ramotswe

I designed and wrote the Python package RAMOTSWE, which checks the metadata associated with each frame for self-consistency. Features of RAMOTSWE include checking that the observatory in the file path matches the name of the machine it was captured on, or using the coordinates of the observatory and the coordinates of the target to calculate the airmass of the observation and comparing with airmass calculated from telescope altitude and azimuth, or checking the CCD temperature for matching with darks.

3.5 Endeavour

ENDEAVOUR is the main Python processing pipeline for the reduction of the Aquila data stream, and implements the algorithm mapping two-dimensional CCD detector data to the final data product, a one-dimensional map from wavelength to intensity, known as a spectrum. ENDEAVOUR performs all the standard major corrections and calibrations, but some of these are implemented differently due to specific requirements of the observations.

An example of this is that the Aquila dispersion axes are specifically oriented to span the diagonal of the CCD detector, allowing a longer baseline and increasing the wavelength range, rather than horizontally, since precise parallel alignment is rarely achieved in practice anyway. This means that many of the steps of the various algorithms rely on first accurately detecting where the dispersion axis is, rather than just binning vertically down the columns for example.

3.5.1 My contribution to Endeavour

While two other people have contributed to the pipeline previously, I was initially responsible for the wavelength calibration and extraction originally, which I more recently modernised (Chapter 7). I also overhauled the main sequencing of events, making the pipeline modular to improve readability and enable quicker diagnostic reports of the middle stages (Chapter 4). I discovered a systematic error in some spectra, which was caused by poor selection of calibration arcs, and developed a new selection technique which minimises the discrepancy between calibration and science frames (Chapter 5). Finally, there was an error causing faint targets to show a beating effect, which I determined to be caused by a discretisation artefact during the flat-fielding step, and hence proposed a new method to counteract this (Chapter 6). ENDEAVOUR is designed to run in a scalable fashion on either a local desktop or a remote server. Since the Global Jet Watch Aquila data stream alone contains over 3 TB of observations to date, the heavy-duty running is nearly always done on the remote server, and desktop usage is now reserved for testing and development purposes only.

3.5.2 Types of frame

The starting point of ENDEAVOUR is two-dimensional raw data from the telescopes, including science frames for astrophysically interesting targets; LEDs and tungsten lamps used for flat-fielding; Thorium-Krypton (hereafter ThKr) bulbs for wavelength calibration; dark frames and bias frames for subtraction. Hot pixels are removed, and cosmic rays are excised.

3.5.3 Gradient method

The flat-fielding and extraction steps are somewhat involved. Instead of simply dividing out the flat-field, a straight-line regression is performed within each wavelength bin between the science counts and the flat-field values for exactly corresponding pixels. The gradient of this line is then more robust to the effects of outliers (such as cosmic rays or detector defects), although there are few of these, since pixels with counts beyond their linearity limit are excluded from the calculation. The gradient, representing the relative intensity of the science pixels compared to the flat-field, is the value extracted to the one-dimensional spectrum for that wavelength bin. This technique is improved upon for observations of faint targets in Chapter 6.

3.6 Poirot

POIROT has two main functions; fitting processed one-dimensional data, and visualising the results in various ways. It also performs the heliocentric correction to

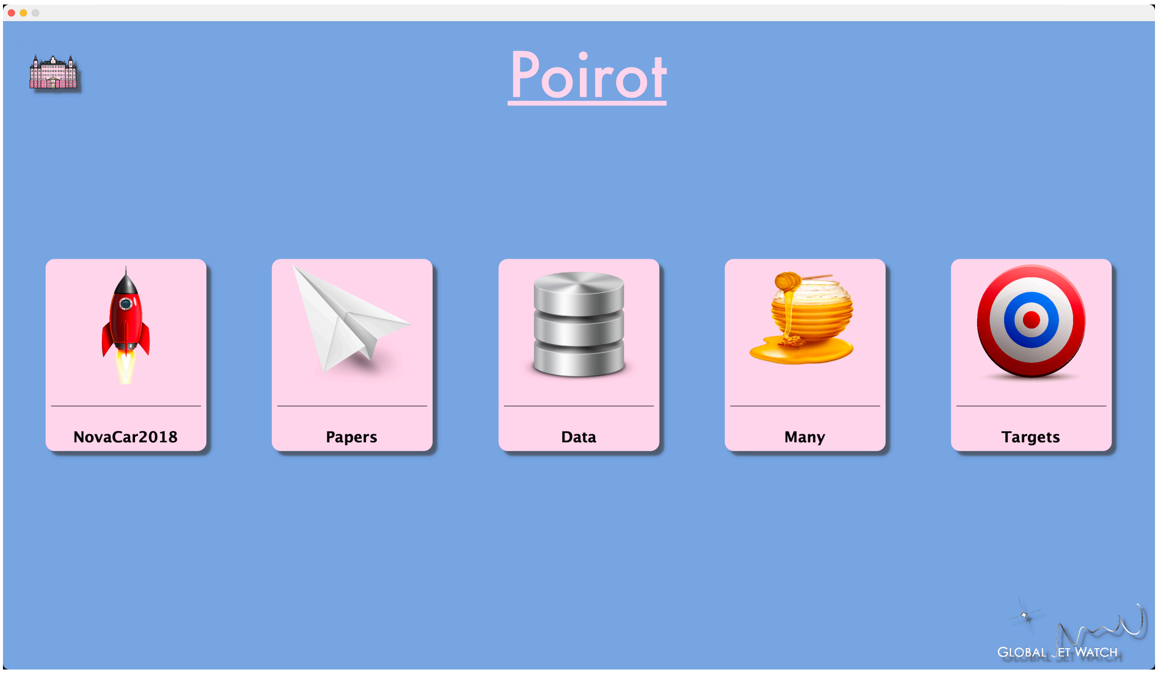

Figure 3.6: The POIROT desktop application for spectroscopic analysis. The first card allows various inspection and fitting tools for the selected target, here NovaCar2018 (also known as V906 Car). The target may be selected using the fifth card (Targets), out of those with data available on the system (see Figure 3.7). The second card (Papers) gathers together plots by publication, enabling automatical regeneration with new data. The third card (Data) gives a dashboard of various data diagnostics, and allows retrieval of data from the remote repository server, and conversions between Julian Day and days since discovery of each target. The fourth card (Many) contains plots which use data from many targets simultaneously for comparisons, rather than just the target currently selected.

the one-dimensional data, which is intentionally left out of ENDEAVOUR so that the telluric features are at their correct wavelengths for POIROT, which can use this as an additional check.

3.6.1 Poirot: fitting data

POIROT picks up exactly where ENDEAVOUR leaves off. It reads in the one-dimensional spectra, and compresses the observations for each target into a binary format, which is then very quick to read in subsequent runs. POIROT provides a set of Python objects to be used with spectroscopic data, such as the `Spectrum()`, `Sequence()`, and `Fit()` classes. These together enable rapid prototyping of templates to fit to the data, and rapid application of the templates to the full time-series data set for a given target. It is these objects which are serialised using PANDAS [Mckinney, 2010], a Python data analysis tool, with each spectrum uniquely indexed by its fractional Julian Day (JD).

As part of the computational side, POIROT allows for a rich variety of operations to be performed on the spectra as a set, including principal component analysis, computation of eigenspectra, composing nightly median spectra, or filtering spectra based on any of the FITS header metadata channels.

3.6.2 Poirot: displaying data

The second powerful aspect of POIROT is as a visualisation application. POIROT runs as a graphical user interface, providing an easy-to-use experience which carefully organises the user's plots, which automatically update and regenerate to incorporate new data as they are processed by ENDEAVOUR. POIROT is designed to be run on a local desktop environment. The home screen of the user interface is shown in Figure 3.6. The first card takes the user to the central page for the particular target (which in Figure 3.6 is set to V906 Car), where it is possible to view all the reduced Global Jet Watch spectra of a given target sequentially, and perform fits to spectral complexes for that target.

The target can be changed in the *Targets* interface, shown in Figure 3.7. The *Data* section manages the downloading of reduced data from the remote servers where ENDEAVOUR is run. The *Many* page contains plots which use data from multiple targets. All of the science plots from the publications used in this thesis are managed by POIROT, collected under *Papers*. I designed and wrote all of POIROT.

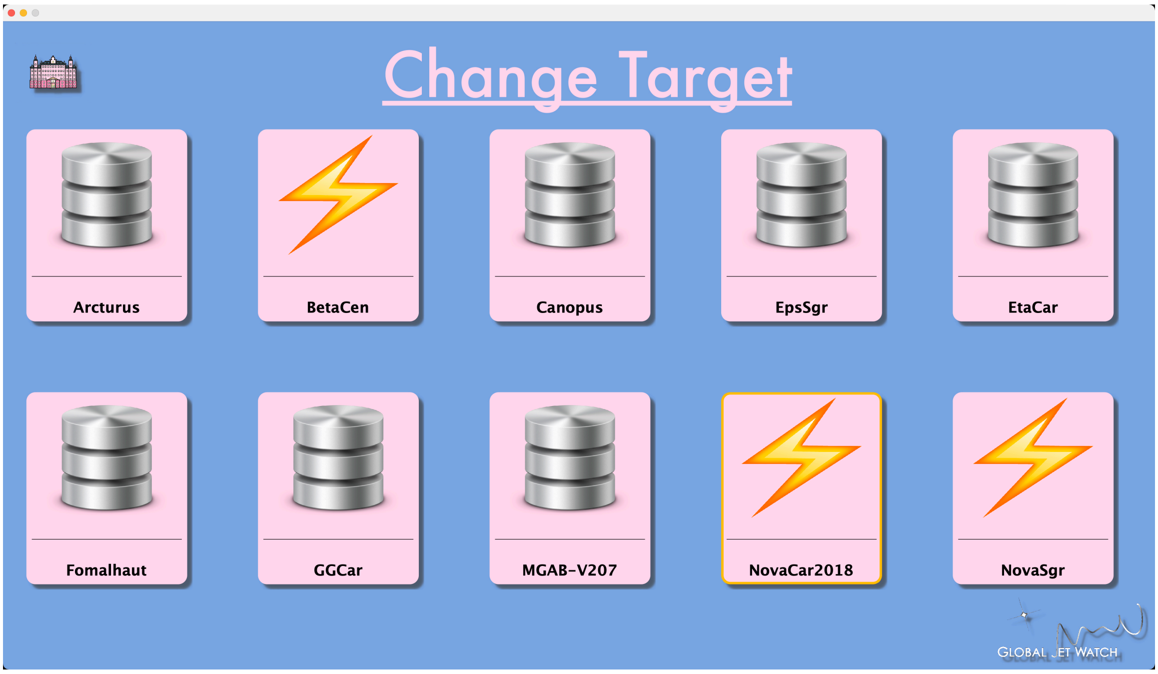

Figure 3.7: Example of POIROT change target page. V906 Car (here named NovaCar2018) is the currently selected target; V5668 Sgr (here named NovaSgr) and β Centauri, (here BetaCen) have also been loaded into memory.

Chapter 4

Endeavour: refactoring the pipeline

This Chapter details my refactoring of the main software pipeline responsible for reducing spectra captured at the Global Jet Watch observatories, namely ENDEAVOUR.

The complete re-writing of this package resulted in a revision known as AnotherOne. AnotherOne is partly a functional re-write to incorporate the upgraded techniques to be discussed in the following few Chapters, but also was an overhaul to allow the algorithm to be more robust, easier to maintain, and easier to use.

ENDEAVOUR was commissioned, although data reduction software is available online, because the Global Jet Watch needs its own bespoke system for the flexibility to fine-tune the reduction process to suit the specific needs of our unique and rich data sets, as well as to satisfy other operational constraints. This is explained further in Section 4.1, where I detail the role of the reduction pipeline, ENDEAVOUR. In Section 4.2, I outline the scope of the changes needing to be made as part of AnotherOne. In Section 4.3, I go through each individual step of the pipeline in more detail, explaining precisely what its function is, and how it achieves it. I also discuss the limitations and compromises that were made. In Section 4.4, I investigate the computational performance of the ENDEAVOUR algorithm, and detail how it runs in practice on different machines, both server and desktop.

4.1 Scope of Endeavour

This Section explores the scope of ENDEAVOUR. There are many constraints and quirks associated with the Global Jet Watch dataset, which the pipeline has to resolve robustly. The primary focus of the Aquila instruments at the Global Jet Watch observatories is to produce well time-sampled data, and as such, backwards-compatibility is particularly important so as to ensure we have the longest possible run of consistently processed data.

The ENDEAVOUR pipeline must be able to reduce all spectra taken at the 5 Global Jet Watch observatories. At the time of writing, there were over 180,000 FITS[1] files in the data repository, including all science observations, calibration frames, darks and biases, dating back to 2012-04-25. These files are all two-dimensional, and are the direct outputs of CCD detectors in the observatories. The job of ENDEAVOUR is to perform all necessary checks, calibrations and corrections, and to produce high-quality one-dimensional extracted spectra, ready to be used to perform science. It is important that it produces high-quality spectra regularly, so that the group has up-to-date observations of our targets of interest, while acknowledging that future iterations may bring enhanced signal-to-noise and other refinements.

The top priority for ENDEAVOUR is that it reduces data with the highest integrity, while also maximising the signal-to-noise. We want to faithfully reproduce the spectral distribution of light as it landed on the telescope, so correcting for the aberrations and other optical effects of the observing system is within the scope of ENDEAVOUR. An important part of this process is to optimise the choice of calibration arcs to be used in the reduction of a given science frame, which is the main topic of Chapter 5.

Running ENDEAVOUR can take some time, so there is also a need to optimise and economise which files are processed during this development process. In future, this will be run on a distributed cluster, allowing much faster processing of historical data. The Global Jet Watch telescopes are used to observe many varied targets, but with particular focus on certain objects which are the subject of current research provided of the group. For example, this thesis focuses on classical novae such as V906 Car and V5668 Sgr, whereas other members of the group work on SS433, Eta Carina, GG Carina and other targets. Thus it is pragmatic to have a mode of operation of the pipeline which can prioritise the observations pertaining to one of these targets for when rapid reduction is needed for a particular campaign. In steady-state operation, all of the data from every target will be reduced together in one data release, but currently it is adequate to reduce by target for speed of access and to help with testing new iterations of the evolving ENDEAVOUR pipeline.

Another constraint on ENDEAVOUR is that it requires the agility to handle different pixel dimensions of input data. Over time, due to failure of hardware components, the instruments used at Global Jet Watch observatories have changed. A result of this is that different CCD pixel dimensions are used over the history of the project.

---

[1] `.fits`, standing for Flexible Image Transport System, is the standard astronomical data file format representing an observation, and consists of the data itself, and a header containing metadata pertinent to the observation [Calabretta and Greisen, 2002]

One reason for having a bespoke pipeline is to robustly reduce the entire dataset, so that the resultant spectra are as well time-sampled and self-consistent as possible. ENDEAVOUR has to be able to reduce all spectra.

It is outside of the scope of ENDEAVOUR to decide whether a science spectrum is of high quality. The role of ENDEAVOUR is to reduce all science frames indiscriminately, and simply output the complete set of extracted one-dimensional data. The analysis codes (e.g. POIROT), together with visual inspection, are responsible for deciding what constitutes useful data for a particular experiment.

To ensure the long-term robustness of the Global Jet Watch data reduction software, it was important to make sure that it was easy to understand and update as novel techniques for data reduction were developed. To ensure readability and ease of access, it was decided that it should be written in the Python programming language.

4.2 Scope of AnotherOne

The major revision known as AnotherOne is required to address three major areas which would allow the pipeline to be more easily extended, parallelised and run, described in this section.

4.2.1 Readability and PEP8

Readability is of paramount importance for maintaining a large and complex pipeline like ENDEAVOUR. AnotherOne is designed to be as simple as possible (but no simpler). Central to Python's published style guidelines is a manual called PEP 8[2], a guide which outlines a defined set of style conventions so that future developers find it easier to contribute. AnotherOne follows this protocol where possible.

4.2.2 Modularisation

Modularisation is the principle that code should be made up of distinct blocks which perform an operation with a well-designed interface. Instead of having one long monolithic script, the processing should be split up into chunks which correspond to the thematically different operations being conducted. This makes it easier to identify where a problem occurs if something goes wrong. It also makes it easier to extend and improve the pipeline; changing the internal operation of one of the blocks, as long as you maintain the external interface, will not affect the other blocks.

---

[2] https://www.python.org/dev/peps/pep-0008/

4.2.3 Batch to continuous processing

The previous version of ENDEAVOUR relied heavily on the concept of the "observatory-night". This was a distinct directory containing observations and calibration frames from a 24-hour period at a given observatory. The original intention was that processing would happen within an observatory-night. However, in practice there are not always sufficient high-quality calibration frames for each data-night, so the pipeline was adjusted to copy calibration frames from another night into the same directory. One of the goals of AnotherOne was to eliminate this practice, by choosing the best calibration frames for each particular science frame instead of just the ones that were taken on the same day. Chapter 5 goes into more detail justifying why this was a necessary step to improve the overall integrity of the final data product, but essentially an investigation into the drift observed in calibration frames revealed it was due to diurnal temperature variations, superimposed on broader weather trends and seasonality, at the observatories.

4.2.4 Parallel code

A large obstacle facing the previous edition of ENDEAVOUR was that it ran in series. There are many steps to this process that are highly computational, and are limited by CPU speed. The parallelisation of ENDEAVOUR was a necessary improvement in order for the reduction time for the entire time sequence data for a given astrophysical target to complete in a practical amount of time. However, this needed to be done in a way that was accurate. Many steps of the previous version of ENDEAVOUR worked by writing interim data to a temporary file, which was then read later on in the pipeline. While this reduced the memory footprint of the program, it presented a risk that separate processes might attempt to read and write to the same temporary file. This will be especially important going forwards when attempting to run the pipeline on a distributed server.

4.3 AnotherOne: phases of Endeavour

The major revision AnotherOne splits the operation of ENDEAVOUR into five key phases, which handle the different parts of the reduction process. These are: Scan, Prepare, CAM, Identify, and Extract. This section explains the role of each of these phases, and their implementation in the new pipeline.

4.3.1 Scan

The first phase of ENDEAVOUR is called Scan. Scan iterates over all available FITS files in the *input* directory, reads them, and stores the header metadata. For ThKr frames, it also calculates and stores (to disk) the pixel coordinates of peaks in the frame, which are later used in a regression to determine a low-order polynomial fit to the dispersion axis. This saves time later, as Scan performs this operation once per frame. As new frames are added to the input directory, the pipeline may be re-run and it will skip frames which have already been processed by a previous run (although this may be overridden in the case of improved algorithms). This step is fully parallel, and results are hereafter known as the *scan data.* The scanning algorithm for ThKr frames was upgraded significantly, and the method is discussed fully in Chapter 7.

4.3.2 Prepare

The next phase is called Prepare. The role of Prepare is to perform sanity-checks and produce median correction frames. The dark and bias frames vary on a timescale of several months, so it is sufficient to have monthly medians. Prepare takes the Scan output as an input, and from this deduces which dark frames to median for each month. It also manages the flats, and biases in a similar way. If there are science frames which have an exposure time for which there isn't a corresponding dark that month, Prepare defaults to a master dark, which are maintained for each operational exposure time and are detector-specific.

4.3.3 Calibration Arc Matching

The calibration arc matching phase (CAM) uses a novel algorithm to optimally choose the best calibration frames for each science frame within the context of the unique set-up at Global Jet Watch observatories. As mentioned above, an investigation revealed that thermal variations at the observatories cause significant changes to the CCD illumination. Discussed thoroughly in Chapter 5, the conclusion of that work was that the best ThKr frame to use for a given science frame is not necessarily the one that is taken closest in time, but instead is the one for which the illuminated pixels in the two frames coincide most closely. ENDEAVOUR implements the CAM phase by evaluating the line integral of science pixel values that lie along the dispersion axis of the ThKr frame, using cubic interpolation to achieve sub-pixel accuracy. The score of a given science-calibration pairing is the line integral, normalised by the exposure time of the ThKr. The CAM phase takes the Scan output as its input, and returns a table of matching scores, which later phases can use to determine the best ThKr to choose for a particular science frame.

One outstanding issue with CAM is that under this implementation, to achieve the best possible matching, it is necessary to compare every science frame with every possible calibration frame taken with the same detector. The resultant algorithmic time complexity is therefore $\mathcal{O}(mn)$, where m and n are the number of science and calibration frames respectively. This is an issue, because calculating the integral of many thousands of interpolated frames, comprising well over a million pixels each, is currently fairly computationally expensive. Although this can be mitigated against by only comparing files from the same observatory, or the same year, this would be more robust to future increases in data repository size if it scaled more favourably. A potential improvement to this algorithm I have considered is to find a parameterisation which captures the location of the dispersion axis for each calibration frame and each science frame. Then sort the calibration frames, at a worst-case run-time of $\mathcal{O}(n \log n)$. Then, for each science frame, search the calibration frame parameters for the closest match, at $\mathcal{O}(n)$. The difficulty with this operation relates to subtle differences between the profile of light landing on the CCD during a faint observation (some science targets) and a bright observation (the ThKr calibration lamps).

4.3.4 Identify

The next phase is Identify, which takes the ThKr frames needed to process the current run's science frames, and produces corresponding two-dimensional wavelength maps. This phase previously required a human to manually identify a single ThKr frame for each observatory, which was done many years ago as a one-off task, and did not need to be renewed, although this became a problem when new instrumentation was added to the observatories. Chapter 7 supercedes this approach, which serves as a record of how this was done before.

The old Identify used a search-and-destroy algorithm. First, it scanned an off-axis region of the frame to determine the background mean and standard deviation noise pixel counts, filtering on this, as well as masking pixels which went above the saturation threshold of the CCD. Then, it iterated over the frame, picking out the island with highest pixel count and masking it. Next, it calculated and stored the pixel coordinates of the centre-of-mass of each island. It also attempted to fit a two-dimensional Gaussian profile to the pixel counts, not recording the peak if it failed to fit. In the intial manual run, these Gaussians were then assigned a wavelength by inspection, and then in subsequent automatic runs, Identify would compare the list of Gaussians with the manually identified set, matching the relative spacing ratios of Gaussians to form a best guess identification. Once the code matched these Gaussians, it fitted them with a polynomial, and used the manually identified list to predict where to search for subsequent Gaussians in the ThKr frame. This technique more than doubled the number of correctly identified Gaussians because the smaller peaks were previously often not found as they are comparable in size with the halo surrounding larger peaks, and knowing where to check gave the algorithm a greater chance of a successful fit.

Once Identify had a list of test peaks matched up to manually identified peaks, it used them to calculate two polynomials, the first describing the dispersion axis in x and y pixel coordinates, and the second mapping the x-coordinate of a pixel on the dispersion axis to the wavelength, λ. Next, Identify selected a stripe of pixels within a fixed distance on the chip from the dispersion axis, and determined the wavelength of each pixel. This involved solving a high order polynomial for around 100,000 pixels, and as such was one of the slower steps in the whole procedure. The end result of Identify was a two-dimensional map relating pixels to the corresponding wavelength, ready to be used in the next stage.

4.3.5 Extract

The final step is Extract, whereby each science frame is reduced from two-dimensional CCD data to a one-dimensional spectrum. The previous steps all record their output, making it available to Extract. Extract checks the CAM data to see which ThKr wavelength map to use. It uses the Prepare data to get the darks and flats. It then uses the gradient method to extract a 1D spectrum. This method entails iterating over bins cut perpendicular to the dispersion axis corresponding to a particular wavelength, and performing a regression between the science pixel value and the flat-field pixel value. It then uses the gradient of this relationship, considering only the pixels with enough signal but not too much to be saturated, as the intensity value for that wavelength bin. This was not my original idea, but it is essential to understanding how the data are reduced. Chapter 6 discusses an enhancement to the gradient method used during the extraction of observations of faint targets. Extract does not make any adjustments depending on the observing conditions (e.g. cloud cover or lunar phase); instead, the plotting and fitting package, POIROT, allows the user to mark spectra that are not high quality.

My contribution to the Extract step has been primarily to speed it up significantly, which is discussed in greater depth in Section 4.4. I have managed to reduce the processing time for a single science frame from 5 minutes down to 11 seconds, in optimal conditions. During a large run, the worst-case scenario time taken per science frame was 21.3 seconds. This improvement was especially necessary for planned observations, which have since quadrupled the total number of science frames in the Global Jet Watch data repository. Further extension of the observatories to incorporate a blue arm to the spectrographs will further increase the load on ENDEAVOUR. The key areas where this improvement happened were optimisations in the logical order of performing operations, moving some of the slower Python functions into efficient C code, and making it run in parallel. I also improved the technique for processing the flat-field frames, which Chapter 6 covers in more detail.

Parallelisation introduced a lot of space complexity, which is addressed in Section 4.4. Memory is not an issue for the serial version of Extract, since there is only ever one set of two-dimensional frames loaded in memory. It is worth noting that due to the excellent operation of NumPy [Harris et al., 2020] and SciPy [Virtanen et al., 2020], it is possible to get the best of both worlds, by running one science frame at a time to limit memory usage, while using vectorised operations optimised to efficiently spread load across several CPUs.

4.4 Performance

I discovered that the run-time performance of ENDEAVOUR showed marked differences between operating systems. This section explains the experiments I conducted to quantify the different behaviours of ENDEAVOUR.

For the remainder of this section, work was predominantly conducted on two machines - an Apple iMac Pro, and a Linux virtual machine called GJW-DM. While the development of the pipeline was conducted on the desktop iMac workstation, the pipeline is typically run on GJW-DM.. Most of the experiments run constituted full passes of the ENDEAVOUR pipeline. Sections 4.4.2 and 4.4.3 describe the workstations used to conduct the tests. The first and second set of experiments carried out are detailed in Sections 4.4.4 and 4.4.5. Sections 4.4.4.1 and 4.4.5.1 explain the results of the two sets of experiments. Discussion is provided in Sections 4.4.4.2 and 4.4.5.2.

4.4.1 Python

The two workstations were set up with identical versions of the same Python distribution (3.7.7 GCC 7.3.0 by Anaconda Inc), which was used for all of the experiments carried out in this Section.

4.4.2 iMac Pro

The iMac is a Apple iMac Pro (2017), running macOS Catalina, version 10.15.4. The processor is a 3.2 GHz 8-core Intel Xeon W. The machine has 32 GB RAM, comprised of four 2666 MHz DDR4 units, manufactured by Hynix. The graphics card is an 8 GB Radeon Pro Vega 56. Hyper-threading technology is enabled. The hard drive is an Apple SSD AP1024M with 1TB capacity.

4.4.3 GJW-DM

GJW-DM is a Linux machine, which at the time of the experiments was running Ubuntu 16.04.6 LTS. The machine hardware architecture is x86_64. The operating system is GNU/Linux, with Linux kernel release 4.15.0-96-generic. There are 8 CPUs with 2500 MHz, of specification Intel(R) Xeon(R) CPU E5-2680 v3. The machine has 108 GB RAM and 115 GB swap space.

4.4.4 Performance experiments E1

The first set of ENDEAVOUR experiments are called the E1 experiments. These consisted of five standard runs of the ENDEAVOUR pipeline across increasingly large data sets, recording both timing information and system information (such as the memory/core usage). These represent a way to quantify how robustly the pipeline scales in practice, and what computational resources would be needed to perform a full retrospective reduction of the entire Global Jet Watch archive of Aquila data.

Each experiment was run on both the iMac and GJW-DM, as detailed in Sections 4.4.2 and 4.4.3 respectively.

4.4.4.1 Results E1

Figure 4.1 shows the time taken per science frame to complete each of the steps in the pipeline, on the iMac and on GJW-DM. It is clear that the iMac runs faster for the smaller datasets, but slows down drastically when more science frames are input. This slowdown is predominantly affecting the Extract stage, but is also noticeable for the

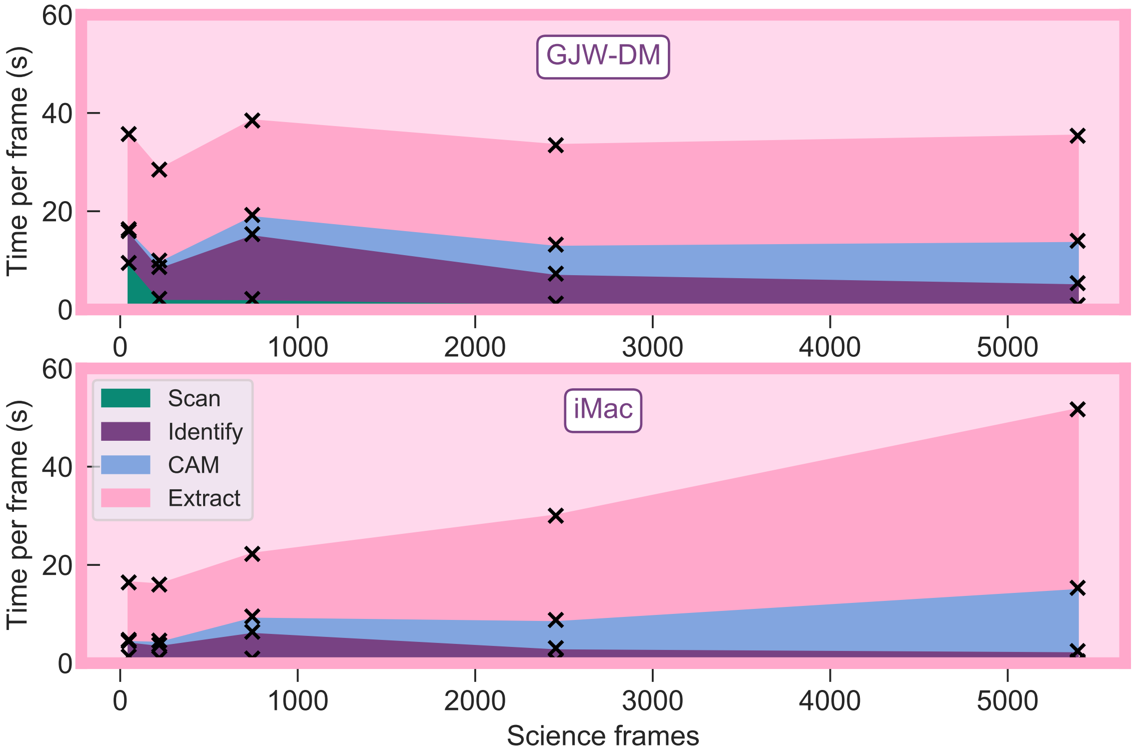

Figure 4.1: Timing information for the four most time-consuming phases of ENDEAVOUR, during the E1 experiments, on GJW-DM (upper panel), and the iMac (lower panel). For each run, I varied the number of science frames to be reduced, from tens of spectra in the smallest run, up to beyond 5000 in the largest. Each phase is shown in a different colour. On both machines, Extract takes the longest time, but on the iMac the duration per-frame increases with number of science frames, while GJW-DM does not show this behaviour.

CAM stage. On GJW-DM, every step of the pipeline except for CAM benefits from economy of scale. This is probably to do with the overhead associated with splitting up work into chunks which can be run in parallel. On smaller runs, the initial cost of forking into more threads costs a non-negligible amount of time, whereas for longer runs lasting several days, this is less important.

4.4.4.2 Discussion E1

The E1 experiments show that the algorithm efficiently scales up to at least 5,000 science frames (chosen as a typical number of observations for a given target) on GJW-DM, but that both CAM and Extract phases scale poorly on the iMac. This difference is because of the way that the two operating systems handle memory. The macOS system keeps the memory used at approximately 16GB, and swaps memory out of RAM onto disk. The longer science runs use more RAM, and this results in a slower operation on the iMac as the swap memory is slower to access compared with RAM. On GJW-DM, the Linux operating system simply allocates more RAM, and the whole operation remains fast. This may hit a hard maximum at some scale, when it may be necessary to break the data up, for example, into separate runs spanning a year of data each, separate observatories already being the case.

4.4.5 Performance experiments E2

The second set of ENDEAVOUR experiments are called E2. For the E2 experiments, I took the output of the largest data set E1 experiment (Run 5), and simply recalculated using smaller subsets of the science frames. This test aimed to determine which of the following two factors was causing the pipeline to run slower when extracting larger numbers of science frames:

1. Memory footprint of loading all the scan data
2. Number of Unix processes created via `fork`

To distinguish between these, every run in E2 was performed with the largest scan data set (E1 run 5) loaded in memory, but was set up to only process a variable number of science frames. This was to determine whether the limit on the memory is the size of the scan data being held in memory, or the cumulative effect of setting off many parallel jobs computing extraction.

4.4.5.1 Results E2

The average time taken to process a single science frame was measured, and I calculated the relative decrement when compared to the time it took under E1 conditions. The results are shown in Figure 4.2. The general trend is that it took far longer under E2 than E1 to process a given frame when there were fewer frames needing to be processed. As expected, when the full set of science frames were computed, the decrement is zero, as the only difference between E1 and E2 is the size of the dataframes.

4.4.5.2 Discussion E2

The downward trend shown in Figure 4.2 is concrete evidence that the size of scan data dictates the amount of time taken to process each frame. When processing only 48 frames, it takes 30 seconds longer per frame with the larger memory footprint

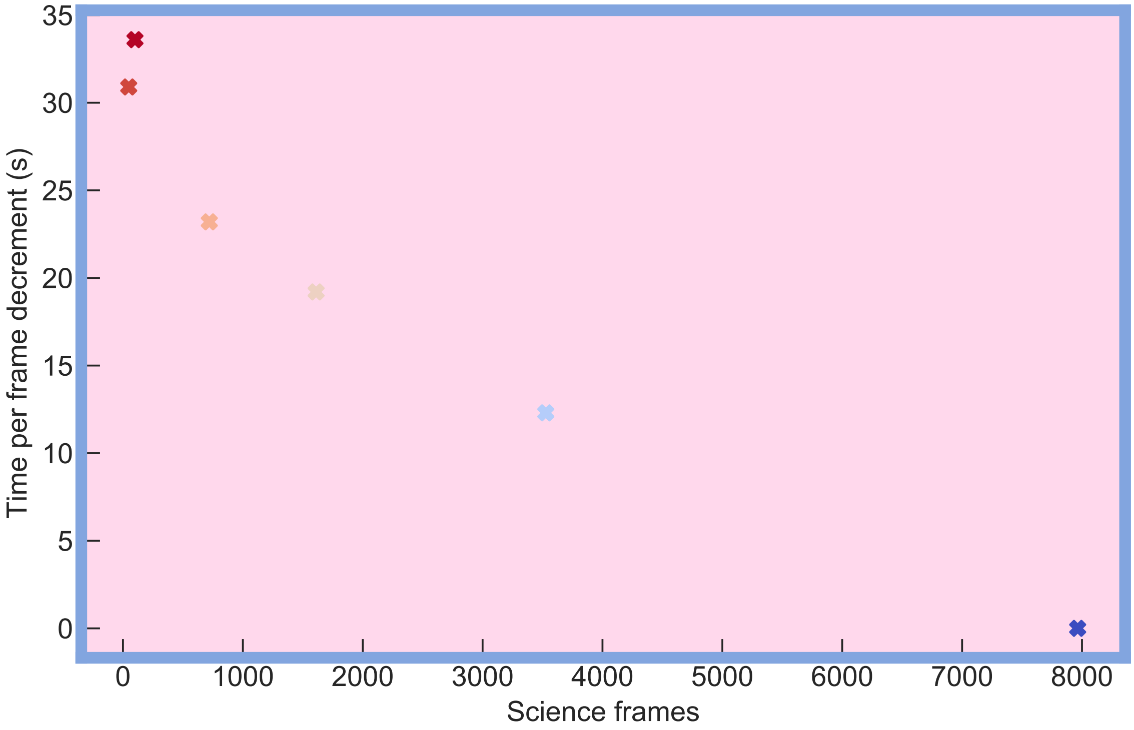

Figure 4.2: Time decrement when running with maximal sized scan data, compared to running with dataframes only the required size for that many science frames.

scan data compared to a fresh run. It is interesting to note that this is a problem which does not affect the Linux runs on GJW-DM. I therefore recommend that the ENDEAVOUR pipeline is run on GJW-DM going forward, rather than on local macOS workstations, unless the number of science frames is below 2000. Future work is also recommended on the ENDEAVOUR pipeline to ensure that the scan data (see Section 4.3.1) is kept to a minimum, and is only accessed as needed, to help limit the memory footprint when run on an iMac system.

4.4.6 Galactic supernova

A supernova has not been observed in the Milky Way since October 1604, although supernova 1987A erupting in the Large Magellanic Cloud was an important local example. If such an event were to happen again, it would be ideal to have a pipeline which can react promptly. While it is not essential to plan for a once-in-a-century event, preparedness for stochastic transient events as a principle is of particular importance for the Global Jet Watch observatories and the associated data processing infrastructure. The ability to quickly generate and visualise spectra is particularly important for objects with unusual and varying brightness, as it can be difficult to decide upon a plan of exposure times, and it would be a shame to lose precious minutes on-sky to saturation. Turnaround time is therefore important for ENDEAVOUR, and as such it is good to know which machine can give us the fastest answers.

Chapter 5

Characterising thermal breathing on CCD detectors

The changes outlined for the data processing pipeline (ENDEAVOUR) in the previous chapter pave the way for a re-assessment of each individual step of the process. This chapter is based around a sequence of Thorium-Krypton arcs (hereafter ThKr, used for wavelength calibration) taken back-to-back and spanning the course of five days, which I show to be caused by diurnal variations, weather patterns and gradual seasonal changes in the ambient temperature. Since I carried out this investigation two years ago, some of the techniques used have since been improved upon and superseded, and some of the physical hardware used for these tests has had thermal insulation added; these details will be mentioned as they come up.

Section 5.1 explains the original motivation for capturing this ThKr data set. Section 5.2 describes how I conducted this analysis. Section 5.3 presents some of the unique issues associated with frames in the data set. Section 5.5 covers the changes that I made to the ENDEAVOUR pipeline to compensate for the thermal breathing effect discovered. Chapter 5.6 concludes by offering a proposal for a follow-up investigation that would help to characterise the extent to which the new apparatus solves the problems identified here.

5.1 Motivation

Fitted centroids of emission lines can show large amounts of drift in a single direction within a single observing night, on the order of 1Å or $50\,\mathrm{km\,s^{-1}}$. This change is far greater than the wavelength resolution, which was 0.3 Å, determined by fitting residuals on a ThKr bulb calibration arc. Earth's annual and daily rotation velocities have already been corrected for, and are stable with respect to the barycentre of the Solar System.

I decided to investigate what caused this nightly drift, and whether it could be calibrated for, or otherwise mitigated. The Global Jet Watch telescopes were equipped with an in situ thermometer to measure the temperature of the spectrograph, in order to test the hypothesis that diurnal thermal variation (*thermal breathing*) is responsible for changing where exactly light lands on the CCD detector, known as the dispersion axis, as it is possible that changes in temperature cause differential deformation of the various components of the optical set-up, which may subtly alter where photons of a specific wavelength land on the detector. With this temperature gauge in place, in December 2018 a five-day continuous run of ThKr arcs were taken specifically for this investigation, during a stint of bad weather at the GJW-CL observatory. This data set comprises a set of arcs with exposure times: 0.3 s, 1 s, 3 s, 10 s, 20 s, each repeated every 30 minutes for 5 days.

A reason why thermal breathing would be particularly problematic for the Global Jet Watch project is its emphasis on time-domain astrophysics. A 0.2 Å effect in a single spectrum may not be catastrophic for line identification, as this would probably be mistaken for a systemic Doppler shift, and would not affect the relative spacing between lines, nor their widths. However, the Global Jet Watch is searching for signals in the time domain with sustained spectroscopic monitoring. Such signals might closely resemble those that thermal breathing would produce — small movements in wavelength, over astronomically short amounts of time (hours, days, months).

An additional motivation for this investigation was to establish an algorithm that detected the xy-coordinate relationship defining the dispersion axis in ThKr frames, which is then used multiple times throughout the extraction process to exclude significantly off-axis pixels from calculations. The algorithm discussed in the Chapter is mostly superseded by a new method, detailed in Chapter 7.

5.2 Method

This section describes how I quantified the locus of points where light landed on the detector, during the five-day sequence of GJW-CL ThKr frames. Within this dataset, I decided to fit the locations of emission line peaks with a quadratic, and compare the evolution of the polynomial coefficients. The process naturally splits into two steps; determining the location of peaks, and selecting good peaks to fit.

The method needed to exclude anomalous bright patches found vastly off-axis, such as cosmic rays. This can be achieved by either of the two steps, and must at least be achieved by the combination of them. During the first step, filtering and signal processing algorithms may reduce the off-axis noise enough that they are not detected as a peak in the first place. Bad peaks may also be rejected after detection, by imposing a tolerance distance from the fitted dispersion axis.

5.2.1 Finding the peaks

I describe three iterations of improvements to the method for finding peaks in two-dimensional ThKr spectra. While computational speed is important, accuracy is more so.

5.2.1.1 Top peaks

The first-iteration fitting technique for finding the peaks simply detected the 30 brightest peaks in each of the raw spectra, storing their pixel coordinates along with the rest of the information in the header, in a PANDAS McKinney [2010] `DataFrame`. This data set enabled plotting various attributes of the observation metadata against each other for easy visualisations of a cross-section of ThKr arcs in the Global Jet Watch. However, using this alone resulted in many off-axis identifications, partly because the halo surrounding some of the brightest emission lines in the spectrum contained pixels with a substantially higher count than the weaker peaks. The halo effect, along with other complications, is discussed in Section 5.3. This technique is very fast, and is dominated by the file read time.

5.2.1.2 Top Peaks Fit 2

To improve the accuracy of the fast peak-finding method, I changed the algorithm so it would split the arc into four subsections by x-coordinate, find the top 100 peaks in each subsection, then fit a quadratic to the total collection of peaks. Then, reject any peaks found which lie more than 0.3 times the standard deviation of the residuals from the fitted curve. This was far more effective at finding *true* (on-axis) peaks than the original algorithm, with limited speed penalty. However, in certain cases it would fail; if not enough of the original points were true peaks, the standard deviation would be large enough to still incorporate many off-axis peaks.

5.2.1.3 Threshold Filtering

This technique uses a modified version of the original peak-finding algorithm used by ENDEAVOUR, optimised to quickly find the peaks in a frame. The original method had a lower success rate on a particular frame and was much slower than the Top peaks method, but outperformed it in terms of finding the correct dispersion axis. It worked by going through eight levels of threshold value, finding connected islands of pixels all above that threshold, and fitting the centre of mass of these islands to use as the coordinates of peaks on the dispersion axis. I tweaked this algorithm initially to only threshold at the first level. This saved a considerable amount of time, taking only a few hours to run through the whole pipeline rather than three days. The method achieved a success on 8043 of the 9482 frames, an 85% success rate. I improved this method further by initially using a median filter to remove noise. This increased the pass rate to 9309/9482, or 98% on the ThKr arcs, and the selected peaks all lay on the dispersion axis by visual inspection.

5.2.2 Fitting the peaks

The result of finding the peaks was a set of metadata from 9300 ThKr frames, including all header information, along with the pixel coordinates of the centres of peaks. I then applied a second-order polynomial fit to the pixel coordinates of peaks in the five-day sequence of GJW-CL ThKr arcs, to find the three coefficients a_i for each frame, according to Equation 5.1. These three additional parameters were then used to characterise the frames, and are referred to as the curvature coefficient, the linear coefficient, and the constant offset. The largest variation was seen in the offset parameter, so I subsequently fixed the curvature and linear coefficients to the mean for that data set to avoid overfitting in those degrees of freedom, allowing only the offset to vary.

$$\lambda = a_0 + a_1 x + a_2 x^2 \tag{5.1}$$

5.3 Problems

There are some significant real-world challenges that complicate this process. This Section outlines particular difficulties associated with some of the frames in the data set, and how ENDEAVOUR addresses these.

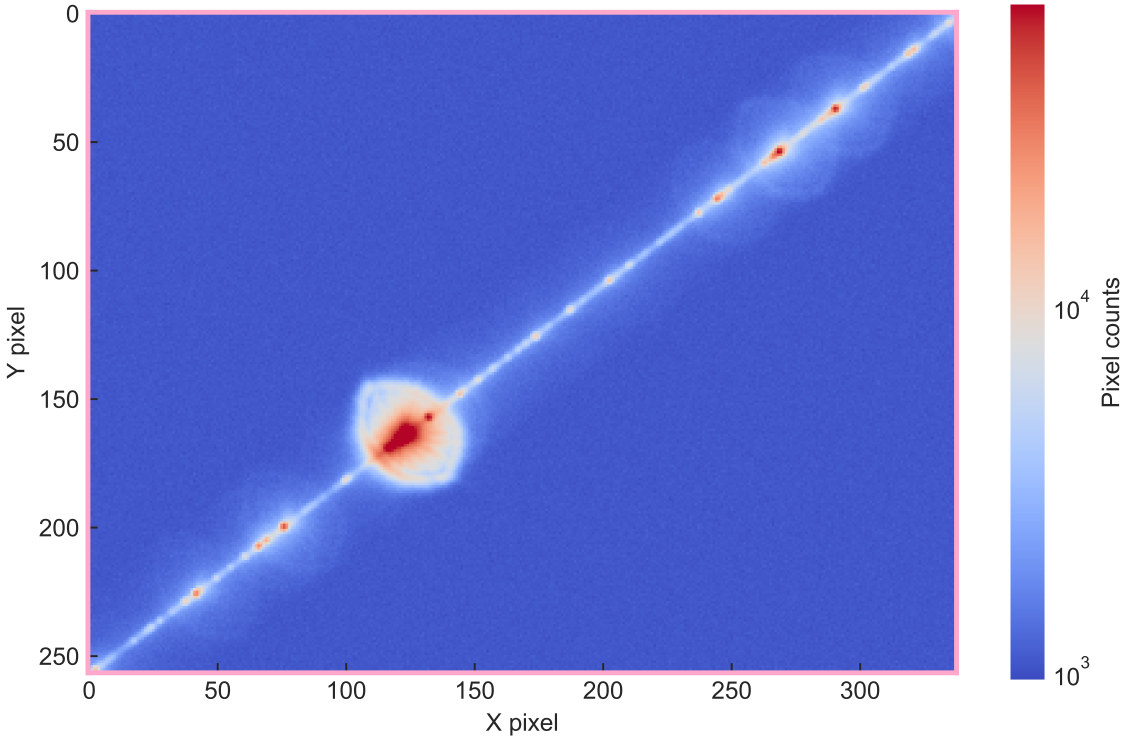

Figure 5.1: The brighter emission lines in the arcs are surrounded by large, bright halos which can be far from the dispersion axis. The final fitting algorithm must exclude peaks off axis, or they may bias the axis determination.

5.3.1 Halo effect

There are large halos around the brightest peaks on the longer exposure time arcs, as per Figure 5.1, arising from the Canon lenses (85 mm, F1.2) . These halos cause problems with peak finding, because they grow in brightness away from the dispersion axis, and have higher pixel counts than some of the weaker lines, which together lead to false positives. While it may be possible at some stage to deal with this problem *in vitro* at the telescope, at present the algorithm must be capable of handling this. The threshold filtering technique described in Section 5.2.1.3 solves this issue using “connected islands”.

5.3.2 Undispersed light

Figure 5.2 shows an example of a ThKr arc that was difficult to process, taking the threshold method 35 minutes to process alone! There is a bright light illuminating the detector away from the dispersion axis, and the brighter bins have comparable photon counts to the peaks used to find the dispersion axis. Potentially we could solve this problem using a high-pass filter, but since a very small fraction of frames are affected by this, it is simpler to just ignore them. This effect was caused by undispersed light entering the CCD and has since been fixed the old-fashioned way: with black tape.

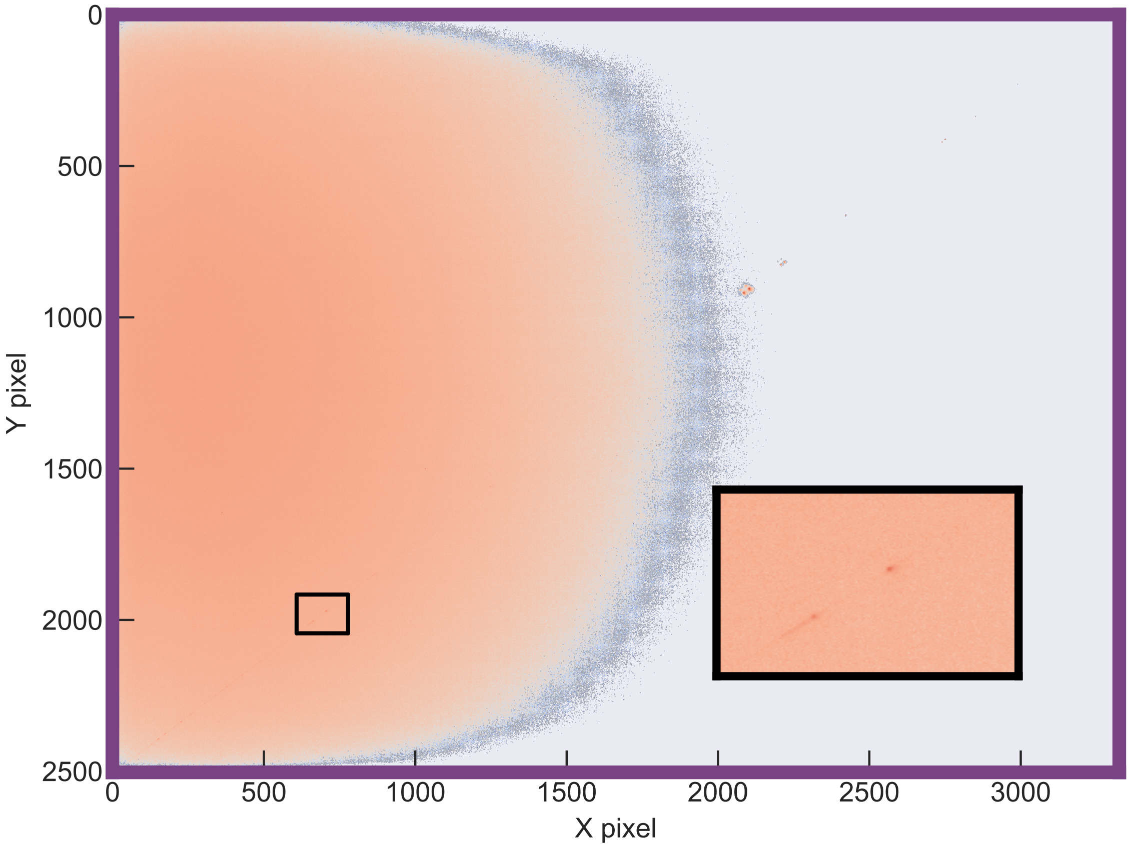

Figure 5.2: An example of a difficult to process ThKr arc because of a light leak. This 10 s exposure was captured at GJW-IN on 2018-12-02.

5.3.3 Left side illumination

In a small minority of frames there is a higher count in pixels on the left hand side of the detector. This could be to do with the way the CCD is read, or additional light leaking into the camera, or a heating effect increasing the dark electron count at one edge. Since this issue has now been resolved with the upgraded detectors, I simply exclude these historical frames.

5.3.4 Median filter offsets

One of the issues I encountered was that the fitted peaks were systematically offset on the order of half a pixel in both x- and y-coordinates from the true location of the

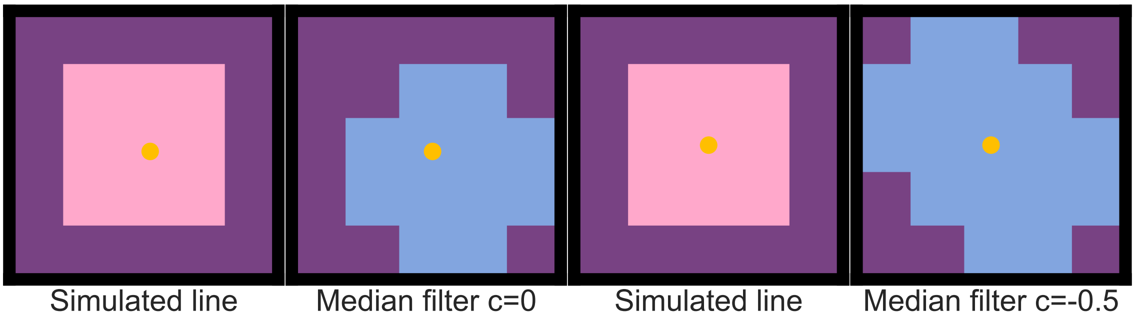

Figure 5.3: Panels one and three show a simulated line in a two-dimensional frame. This mock line has a value of zero around the edge (purple) and a value of one in the central square (pink). Panel two shows the result of passing a standard median filter with centre set to zero over the simulated line, in blue. Panel four shows the same as panel two, but with the centre set to -0.5 — the mean of two medians taken with centre equal to zero and one. The gold dots in panels one and two show the calculated centroid of the pixels in the simulated line which coincide with the median filter, and there is a visible offset from the actual centre. The gold dots in panels three and four show the correct centre, because the improved median function preserves the centre of mass, i.e. true centre of light distribution.

peak by visual inspection. I tested this, and found the cause was the SCIPY [Virtanen et al., 2020] median filter. Passing an even-sized kernel over the data in the Python package resulted in a constant shift of half a pixel. The index 1 refers to the start of the cell, whereas the coordinate 1 refers to the centre of a cell, which is where the data is taken to have been measured. When the centre is set to zero, the median filter effectively acts forwards. By taking the mean of a centre=0 and centre=-1, I was able to create a median filter with no bias, effectively centre -0.5, as shown in Figure 5.3. The plotted centre point is calculated as being the centre of mass, using the median filtered image as a mask on the original data. So in this case, having the median filtered image in advance of the original data means that the right and lower sides of the mask lie on zero-valued pixels (shown in purple). This has the net effect of biasing the centroid calculation towards the upper left. However, the corrected symmetric version of the median filter does not suffer from this problem, and so the centroid was in the correct location for masked files.

5.3.5 Defining the dispersion axis

There are at least two distinct and potentially valid definitions of a dispersion axis:

1. The best-fit line through the pixels with the highest counts.

2. The line such that half the photons incident on the optics are recorded as landing on one side, and half on the other.

While all previous efforts had been centred around the first definition, I decided to try to quantify to what extent the second definition agrees. I decided to measure the asymmetry across the fitted dispersion axis, and define the balance metric as the ratio of counts above/below the dispersion axis. However, this took a long time to calculate, and developments at the telescopes mean that it does not make sense for the most recent observations, so I did not continue with this line of thought. It does, however, make explicit the question of what the dispersion axis actually is.

5.4 Results

This section details the results of the thermal breathing experiments.

5.4.1 Variation of curvature coefficient

When fitting the dispersion axis with Equation 5.1, I explored the time variability of the coefficients a_i (introduced in Section 5.2.2) across historic data from several of the Global Jet Watch observatories to see if this was a useful way to characterise the axis. Within each observatory there are long-time drifts of the curvature coefficient a_2, but these are on a smaller scale than the variation within any given observing night. There was no correlation between time and curvature within the five day sequence at GJW-CL. The resolution of this issue was to assume that the curvature was not varying in a tractable way, and was in fact simply correcting for variation in the constant offset term. The resolution was to fix the linear and curvature terms to their respective means, allowing only the constant offset to vary.

5.4.2 Number of peaks

This algorithm found the coordinates of a large number of peaks, shown in Figure 5.5. The number of peaks found by this algorithm was typically higher than the number that the original ENDEAVOUR used in the Identify step, because for that stage it is also important to have an accurate estimate of the wavelength of the line, rather than simply the pixel coordinates of a peak known to be located on the dispersion axis.

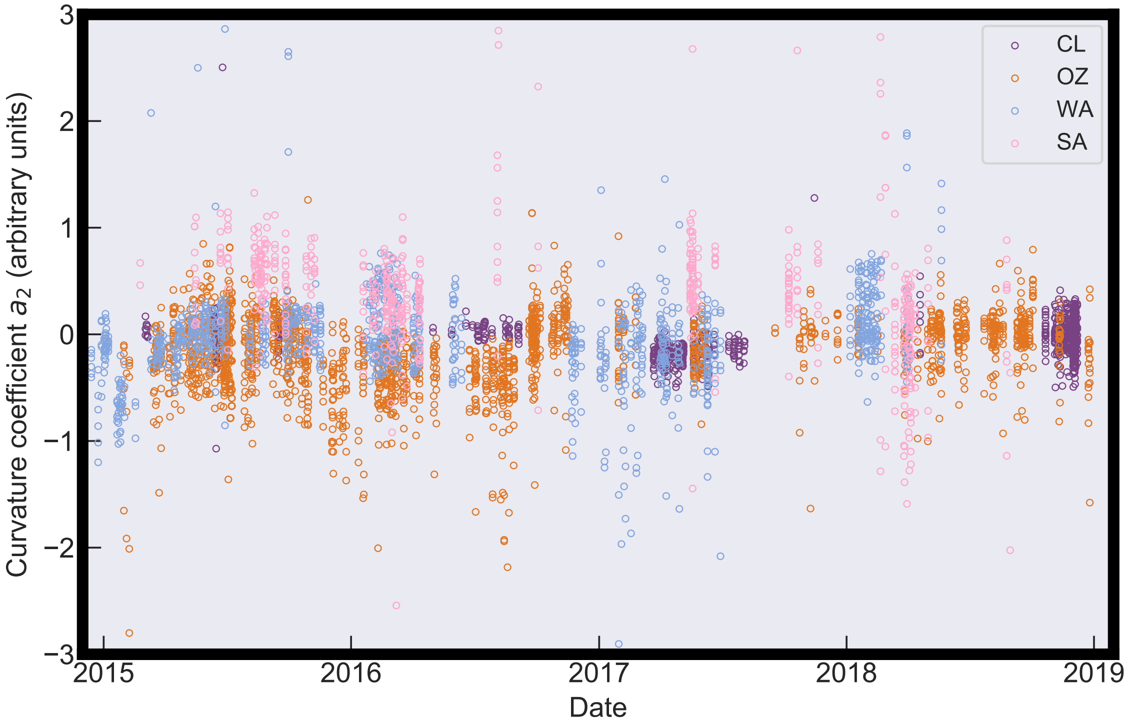

Figure 5.4: Long-term relative variations in the a_2 coefficient, normalised by dividing each observatory series by its mean.

5.4.3 Daily temperature profile

Figure 5.6 shows how the offset parameter varied with time over the five day sequence in GJW-CL. It shows clear diurnal variations — presumably a thermal breathing effect whereby the constant offset parameter of the dispersion axis quadratic fit pulses with time, on a daily cycle. This would seem to imply that there is a stretching of some component in the telescope due to something that varies on a daily cycle, most probably a thermal expansion (we haven't yet ruled out tidal effects!).

Figure 5.7 shows the effects of various physical parameters (determined from the metadata for each observation) on the variations in the offset parameter. While the offset seems to vary with changing atmospheric pressure, it is far more tightly coupled to the temperature profile during the day, and so pressure is most likely a confounding variable. There are several rings in the tracks of temperature with offset, confirming the presence of hysteresis. The implication is that it is not only temperature, but the history of temperature over the day which is significant. Otherwise, the loops would be flattened. This could be due to the different rates of cooling or heating of different parts of the apparatus causing differential stretching and tensions. The lower left

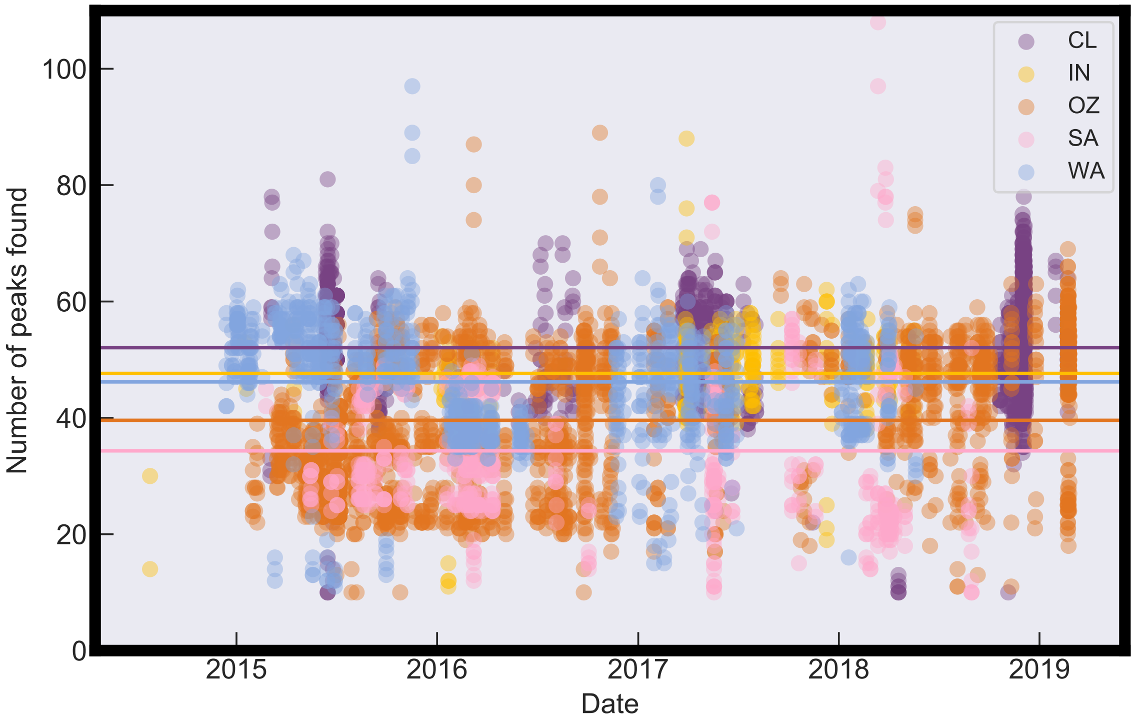

Figure 5.5: The number of peaks found in the dispersion axis of each ThKr frame in the Global Jet Watch. No observatory has consistently fewer than 20 peaks in a ThKr frame, easily sufficient for quadratic regression.

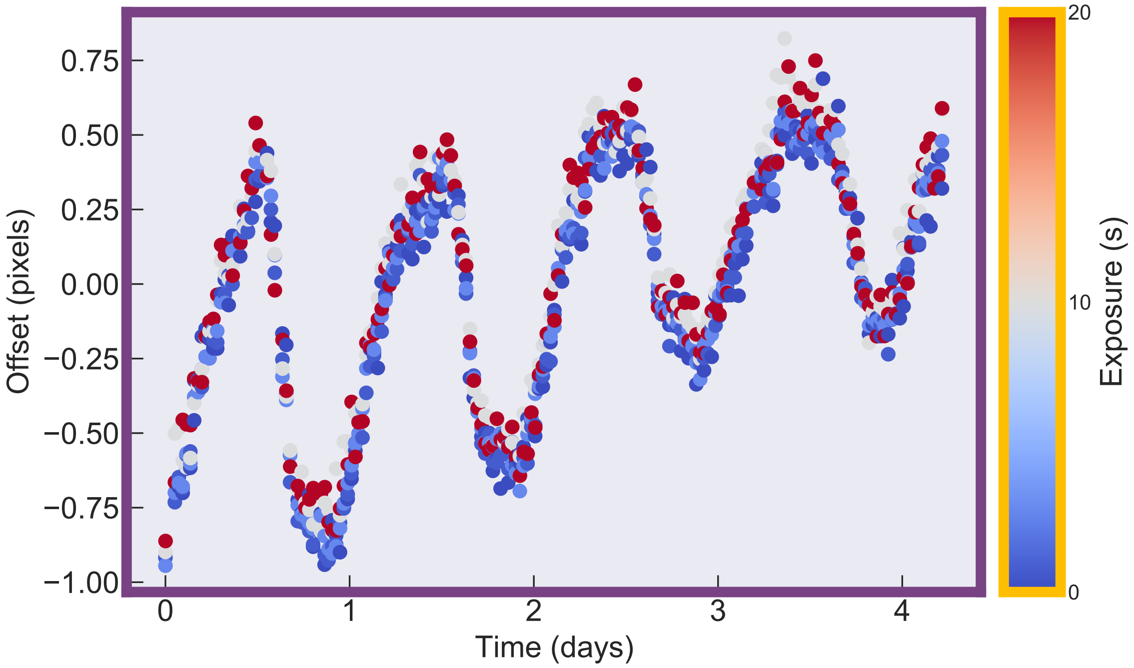

Figure 5.6: Diurnal variations in the offset parameter (with a minor drift over the days imposed by changing weather), plotted as a function of time of observation. The colour of each point represents the exposure time for that observation. While there is little discrepancy between the different exposure times, we see that lower exposures typically have marginally lower values for the offset parameter.

panel, which plots offset against time of day in hours, shows that the parameter rises and fades every day. It was informative to colour-code this also with exposure, shown in the lower right panel. Taken together, this figure shows that ambient temperature is indeed the most important factor in determining the offset, and that exposure time does not greatly influence the determination of dispersion axis.

Note that for all of these observations, the temperature of the CCD detector itself was measured to be −30.0 °C, kept constant by a triple-stage Peltier cooler; it is the ambient temperature, not the CCD temperature that is the source of the thermal breathing effect.

A key consequence of this is that arcs taken in the morning (between 9am and noon, local time) are clearly unsuitable to be used to calibrate the science frames taken that night. There is a fluctuation of 1.5 pixels between different times of day, meaning a different set of pixels are illuminated in the science frame compared to the calibration arc. It is not therefore unreasonable to assume that throughout the entire collection of arcs, there exists a better ThKr for calibrating a given science frame than one which was captured on the same day.

5.4.4 Exposure time

While clearly a secondary effect, close inspection of the upper right panel of Figure 5.7 shows that points making up the rings are still spread vertically in groups of five. This clearly corresponds to scatter between the five different exposures, which were captured back-to-back every half hour. It would be useful to know whether there is a dependence of offset parameter on exposure time. To assess this properly, I have requested more data, in which a new series of Thorium-Krypton arcs, taken consecutively to reduce impact of typically slower thermal changes, with a set of 30 different exposure times from 1 s to 30 s. This would enable a direct comparison between exposure time and how the light lands on the detector. This data set would also clearly show what the saturation limit was for different lines in the ThKr lamp, enabling accurate determination of the optimal balance between over-exposure and signal-to-noise.

5.5 Pipeline upgrade

In Chapter 4, I introduced AnotherOne, a rework of the ENDEAVOUR package to allow for individual components of the pipeline to be rewritten. The effects of thermal breathing mean that the best arcs to reduce a given science observation are not

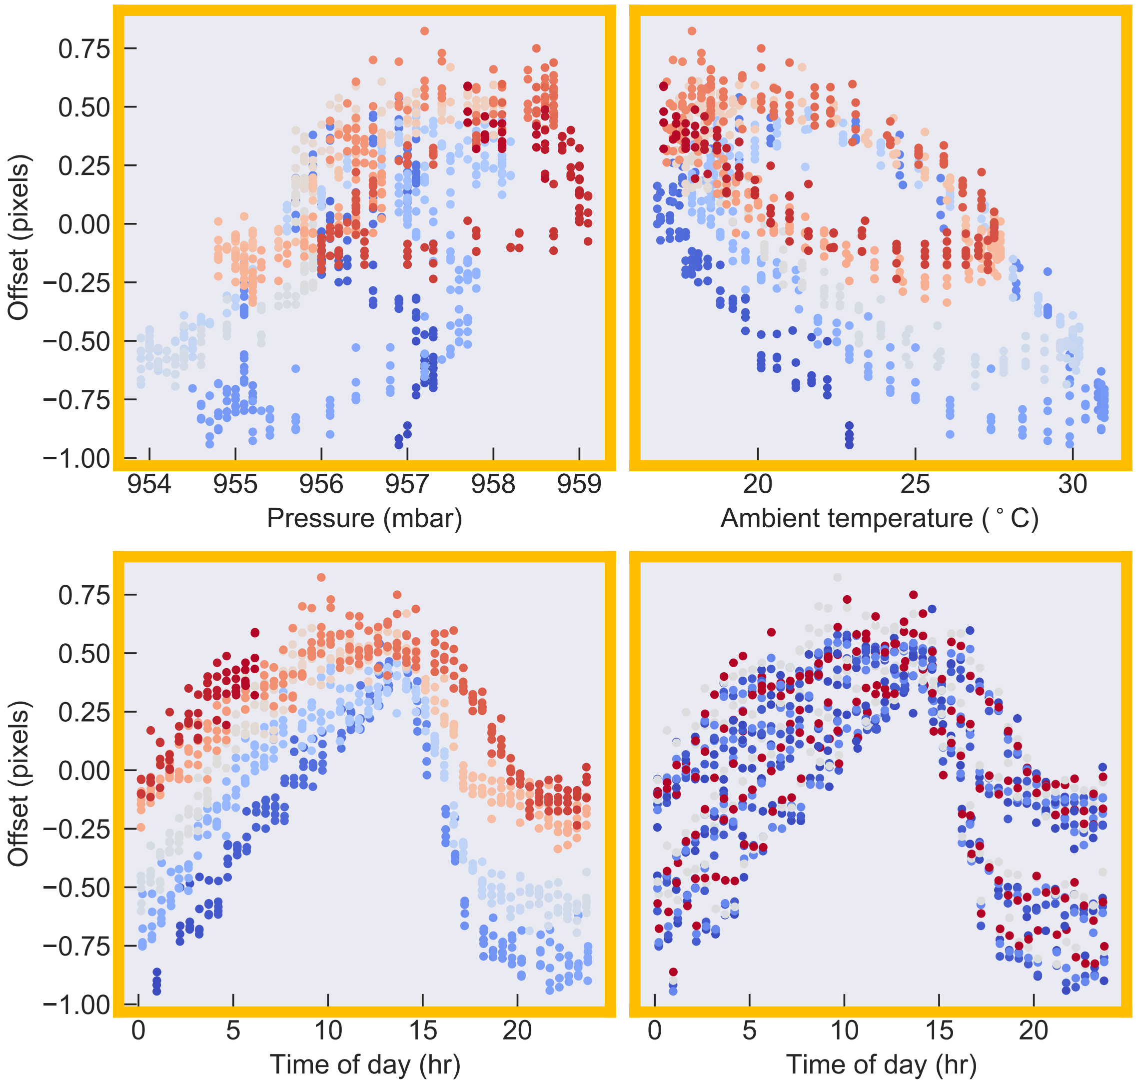

Figure 5.7: The upper left panel shows how the offset parameter depends on ambient pressure, colour-coded by the day of the observation. It clearly traces variations during each day, although each day has a different profile. The upper right panel shows the corresponding dependence on temperature, this time tracing clear rings, which change shape each day. The third panel shows how the offset parameter varies with time of day, showing a hump pattern which is different between days. The lower right panel shows the same data but coloured by exposure time instead, confirming that the exposure time is not a factor in determining the offset.

necessarily always those that are closest in time. This section explores the changes that I made to the calibration arc matching phase, hereafter CAM.

5.5.1 Current technique for CAM

CAM is the process of selecting which calibration frames to use to reduce a given science frame. The CAM method previously in use for the Global Jet Watch was to take a series of ThKr arcs during the daytime, and then use a median of these to reduce each science frame captured by that observatory during that night. In cases where there were no calibration arcs taken that day, the data reduction algorithm would select the arc closest in time to the observation set, and use that to calibrate the data. A significant problem with this approach was that the telescopes are located in schools, where they are used for imaging by the pupils, and can be physically knocked, disturbing the alignment. Also, as alluded to in Section 5.4.3, thermal breathing means that there is no guarantee that the light illuminates the same pixels for arcs taken during local morning as the ones illuminated in a science observation taken at night. Together, these represent a serious issue for the method of selecting arcs based on time proximity alone.

5.5.2 Novel technique for CAM

I proposed a novel technique of CAM for a given science target; select the arc whose measured dispersion axis most closely matches that of the particular science frame. This represents a significant shift in mindset; instead of choosing calibration arcs that were taken close in time to the science target, we choose ones which were taken close in conditions. While this may sound unusual, it means that the actual pixels illuminated in the calibration arc are more likely to be the same as those illuminated during the science exposure. A successful CAM system must take a science frame and the set of ThKr frames, and determine which ThKr is best.

I have considered three main approaches to optimising the CAM for a particular science observation. The first is minimising the chi-squared between fits to the dispersion axes of the two frames. The second is to measure the offset parameter and compare that number between frames. The third approach is to sample the science frame along the dispersion axis as calculated using each Thorium Krypton frame, and maximising the total pixel count. The third approach is immune to the fact that we may not have a good fit for the dispersion axis of a science frame if it has failed for some reason. While the first is more robust, the second may be more time-efficient, which is important given this process will happen on the order of 10^4 times for every single observation in the pipeline. However, the third approach is the most robust as it will work for any science frame conceivable, so this is the one that was selected, and is now used in the reduction of Global Jet Watch data.

5.5.2.1 Sampling dispersion axis

To test the best ThKr arc to be used given a particular science frame, the new algorithm takes the predetermined fitted peaks of each ThKr arc, and fits them to a polynomial. It then evaluates the polynomial at 2×10^3 points, producing a set of estimated y-coordinates for each x, and estimates the value of the integral of science frame pixels which lie under these points, using linear interpolation for speed. The ThKr arc which maximises the value of this integral is then selected as the one which has peaks that lie closest to the true dispersion axis of the science frame, and is selected as the calibration arc to be used. This uses the first definition of the dispersion axis as set out in Section 5.3.5. I tested this method using the V5668 Sgr dataset, which consists of 2791 raw frames at the time of writing, and spanning 335 distinct observatory nights. I found that 306 unique ThKr arcs were selected by this method. It took approximately two seconds to process each frame using this method.

5.6 Moving forwards

The thermal breathing described in this chapter was an issue for early Global Jet Watch data. However, it has since been successfully mitigated by a series of distinct factors. The novel CAM technique described here means that ThKr frames are now implicitly tested for their alignment with science frames, rather than hoping that same-day observations are ideal. Second, the addition of several layers of insulation around the spectrograph should have helped to palliate the thermal breathing. Finally, when developing the new technique for processing ThKr arcs, which forms the basis of Chapter 7, it was found for the new FLI ML16200 detectors that the dispersion axis by a straight line, rather than a quadratic, although it is unclear whether this is down to the new algorithm or the new physical hardware.

I therefore propose a further set of ThKr observations is carried out, using the recently-installed upgraded detectors and with proper thermal insulation of the spectrographs. I conjecture that a five-day series, similar to that captured in December 2018 at GJW-CL, would under the new conditions reveal a greatly reduced diurnal thermal breathing pattern. I further suggest that a suitable acceptance threshold would be variations at or below the quarter-pixel level, at which level it is unlikely that the discrepancy between ThKr and science would be impactful.

Chapter 6

Refinement of the gradient method in the low signal-to-noise regime

In July 2020, classical nova YZ Reticuli erupted, and the Global Jet Watch Aquila instruments set about recording data. During the first few hours of observing this exciting new target, a range of different exposure times were used, since it is *a priori* unclear how many seconds it will take for Hα to saturate. The result of this was that several spectra were taken with a short exposure time, and are thus relatively faint, with lower signal-to-noise. Keen not to squander these valuable early data, I used the existing ENDEAVOUR pipeline to reduce them anyway. However, as demonstrated in Figure 6.1, I discovered a strange artefact in the output spectra. This brief chapter describes the investigation into this phenomenon, and its remedy. Since the existing ENDEAVOUR algorithm is described in greater detail in Augustus Porter's DPhil thesis in the Chapter on endeavour, here I only describe how it previously operated so that the changes I have made to it are clear.

Section 6.1 highlights the problem that was encountered when studying low signal-to-noise observations of classical nova YZ Reticuli. Section 6.2 explains the role of pixel gain maps in ENDEAVOUR, and how flat frames are smoothed in the parallel and perpendicular directions with respect to the dispersion axis. Sections 6.3 and 6.4 detail the existing methods to calculate the dependence of intensity on the parallel and perpendicular directions respectively. Section 6.5 presents a new method I developed which improves the calculation of the pixel gain map, and Section 6.6 shows the improvement in spectral quality from the new method.

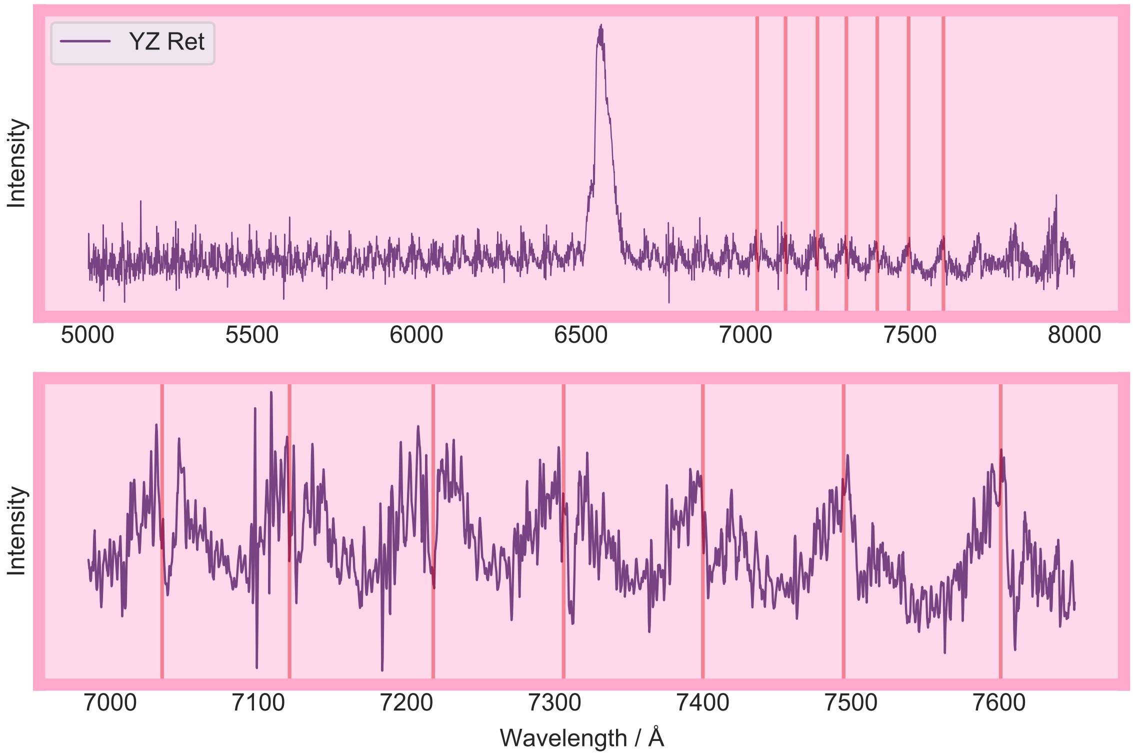

Figure 6.1: Spectrum showing repeating blocky pattern. This spectrum is of YZ Reticuli taken on 2020-07-16 at the GJW-CL observatory. It is a three-second exposure.

6.1 Understanding the problem

6.1.1 Periodic blockiness in faint spectra

In some of the spectra produced via the old algorithm, I observed a periodic signal. As demonstrated in Figure 6.1, approximately every 85 pixels there is a bulge in the spectrum, although the gap between bulges increases with wavelength.

6.1.2 Interference?

I suspected this was caused by a beat frequency interference pattern, an interaction between the discrete pixels and wavelength bins, perhaps related to the diagonal nature of the dispersion axis (which may also be slightly curved). In Figure 6.2, I plot the residuals when subtracting the Y-coordinate along the dispersion axis from the polynomial fit to the Y-coordinate along the dispersion axis, giving the fractional height of each pixel where the fitted polynomial intercepts it. I observe intervals which closely match those observed in the example spectrum (Figure 6.1), which constitutes strong evidence that this is what causes the problem in the final spectra. I therefore decided to check the existing pipeline algorithm for any steps which compute a value directly from discrete pixel values and replace these references with sub-pixel coordinates fitted through interpolation.

6.2 The role of the pixel gain map

Not all pixels are created equal. When photons land on a pixel, they trigger a variable response, with some pixels recording more counts per photon than others. To correct for this effect, ENDEAVOUR calculates a *pixel gain map*, which describes the relative intensity each pixel would show if they were all subjected to an identical number of photons — a *flat-field*. A standard approximation to a flat-field involves using a continuum source; at the Global Jet Watch observatories, the Aquila instruments are pointed at a combination of tungsten lamps and LEDs. These observations are subsequently summed together, creating a frame which effectively has been illuminated across the full wavelength range, shown in the upper panel of Figure 6.3 as *calibration*.

While the lamps do partially illuminate each pixel, they do not illuminate each one equally. There are three major factors which contribute to the pixel value within the flat-field: the wavelength dependence of incident light (parallel to dispersion axis); variations in intensity perpendicular to the dispersion axis known as *the illumination profile*, and the pixel gains themselves. To recover the gain map, it is necessary to

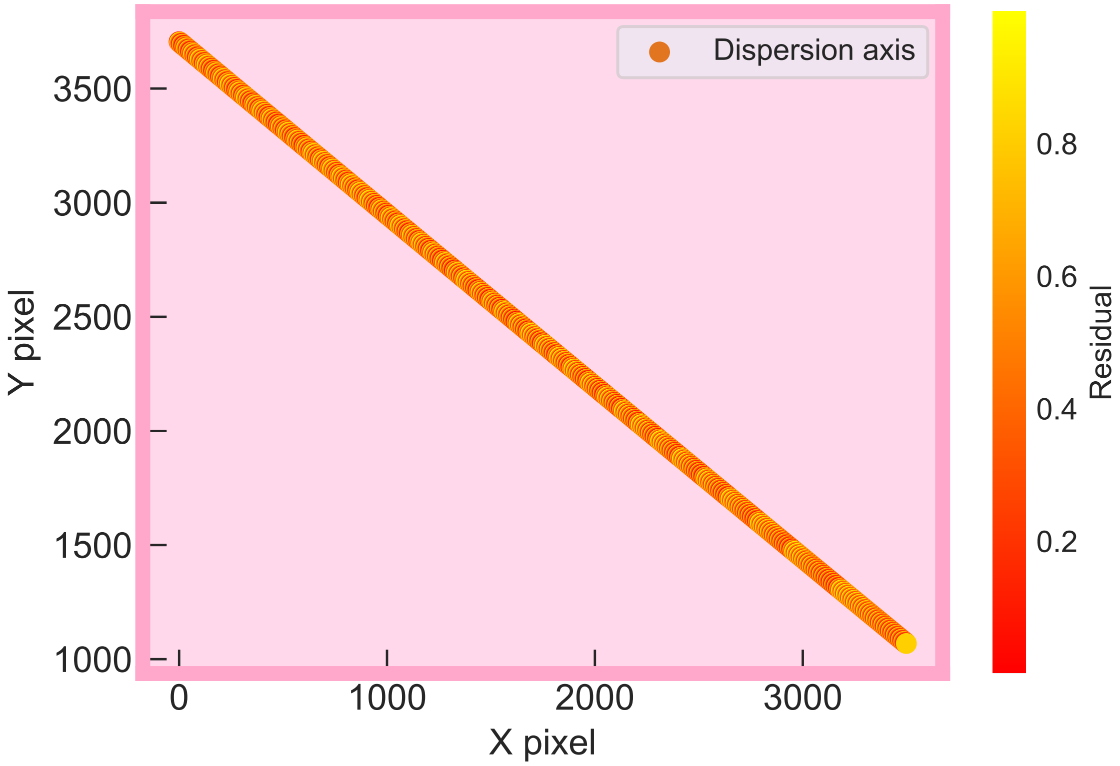

Figure 6.2: Polynomial fit X-Y minus the actual y-coordinate of pixels on the dispersion axis. Colour is mapped to the residual error induced by discretisation of pixels. There is an interference effect apparent in the colours, with a steadily increasing length scale along the dispersion axis from the top left to the bottom right. This is strikingly similar to the relative spacing between successive bulges in Figure 6.1.

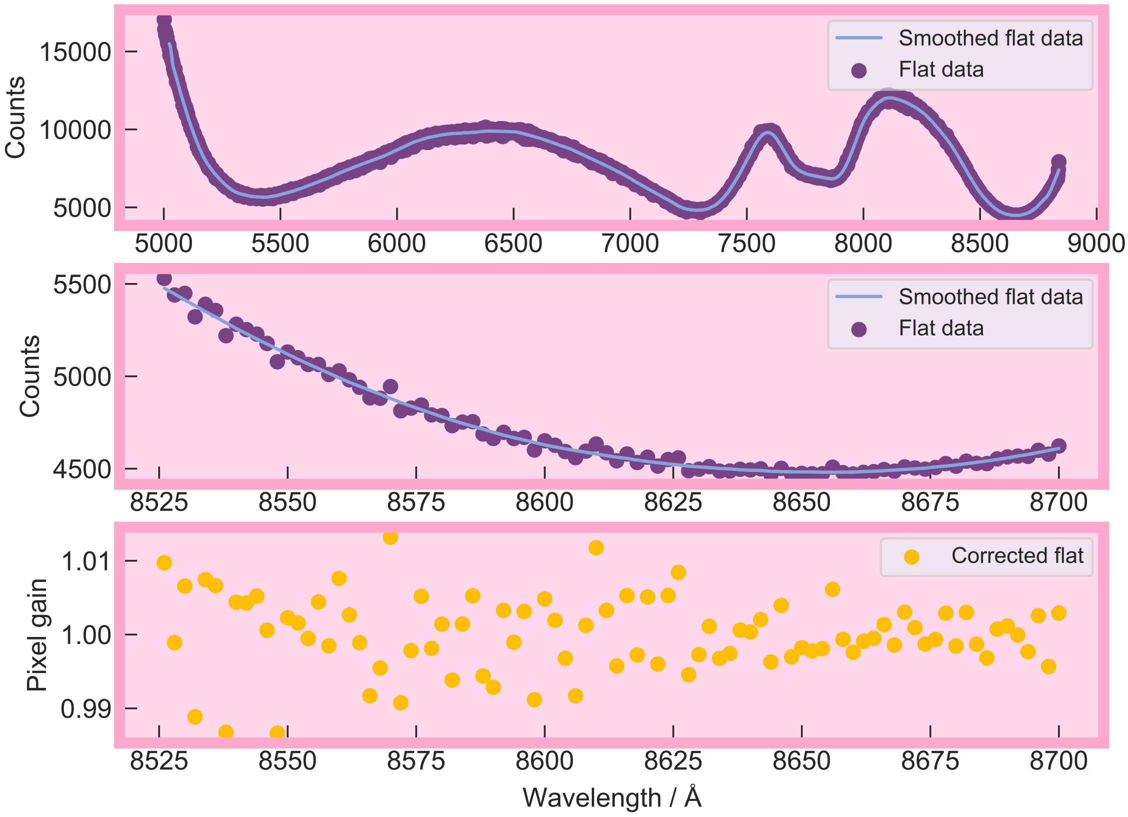

Figure 6.3: The upper panel shows the flat-field light, traced along the dispersion axis, and binned by wavelength. The flat-field light at 5500 Å to 7000 Å is from the tungsten lamp, while the relatively narrower peaks at 7600 Å and 8200 Å are caused by the LEDs. The smoothed flat data is computed using a cubic spline, and is used to divide out the overall wavelength dependence (variations parallel to the dispersion axis) from the flat data. The middle panel shows a close-up view of a subsection of the curves shown in the upper panel. The lower panel shows the same subsection, but of the resultant flat field, in gold.

divide out the other two effects, the parallel and perpendicular dependence. The pixel gain map is a useful diagnostic tool for understanding the behaviour of the detectors, and future enhancements to ENDEAVOUR will include alerts triggered upon detection of hot pixels (pixels which give a high count independently of their illumination).

6.3 Old method to correct parallel dependence

To correct for the large-scale wavelength dependence of the calibration lamps, these values are binned parallel to the dispersion axis, and smoothed using the Savitsky-Golay filter [Press and Teukolsky, 1990]. The algorithm then interpolates this smooth function, assigning an estimate of the light intensity at each pixel using its exact wavelength. This is then divided out of the flat field pixels, shown in the lower panel of Figure 6.3, eliminating the wavelength dependence of the flat. Next, the cross-axis, or perpendicular, variations need to be accounted for.

6.4 Old method to correct perpendicular dependence

In the absence of a flat-field frame, it would be possible to extract a science spectrum by simply binning the number of counts at each wavelength. However, this would weight the near-dispersion-axis pixels far more heavily than is necessary. The ENDEAVOUR pipeline employs a particular method of extracting information from the relationship between flat data (calibration) and science data, known as the *gradient method*. It iterates along the dispersion axis, considering each bin of approximately 2 Å parallel to it and several pixels wide in the perpendicular sense. Within each bin, it fits a linear regression model between the set of corresponding calibration and science values, one pair for each pixel.

The gradient of this straight-line fit is a value that represents the relative illumination of the pixels in a given bin in the science exposure compared to the flat field. In other words, the gradient is proportional to the intensity of light arriving from the science target at this wavelength.

A further complication is that the dispersion axis may have shifted slightly due to temperature fluctuations (Chapter 5) between the observations of flat-field and science frames, so for calibration frames, ENDEAVOUR re-weights values in the science frame compared to the calibration frame, to account for any misalignment. The existing algorithm computed a histogram measuring counts as a function of distance from the dispersion axis in bins containing 2-3 pixels. It would then smooth these values using the Savitsky-Golay algorithm, and divide through by the smooth values, in much the same way as the overall spectrum shape. This is the part which caused interference.

A limitation of the Savitsky-Golay algorithm is that it requires its input array to be evenly-spaced. In any given wavelength bin, the distances from the pixels to the dispersion axis are not integers, because the dispersion axis is diagonal. The previous algorithm binned pixels according to their distances from the dispersion axis, and used these regularised bins as the input to the Savitsky-Golay filter for smoothing. Calculating the cross-axis illumination profile by binning meant that the algorithm was determining the smoothed cross-axis illumination shape without taking into account the actual distance of each pixel. There is no need to discretise the distances, destroying this information, and I describe in the next section my method to overcome this.

6.5 New pixel gain map algorithm

Instead of histogramming the illumination profile, it is preferable to use a zeroth-order spline interpolation to obtain the pixel illuminations as a function of distance. My new algorithm calculates the illumination values on a finely-sampled and equally-spaced set of points, and smooths these using a Savitsky-Golay filter. It then calculates another zeroth-order spline interpolation, which maps these smoothed illuminations back to the original points. For each pixel we now have a smoothed value for illumination, which represents the amount of light that originally landed on that pixel. Dividing through the illumination by its smoothed version for each pixel gives the complete pixel gain map, without the errors incurred by unnecessarily binning into a histogram. This process is shown in Figure 6.4.

Ideally, the pixel gains should be randomly scattered around the *neutral gain* line, which is the pixel gain for an ideal pixel, with the value unity. The plots in Figure 6.5 show these steps. The left plot of Figure 6.5 also demonstrates that the illumination profile perpendicular to the dispersion axis has its own wavelength dependence, demonstrating the need to correct for this separately in each wavelength bin.

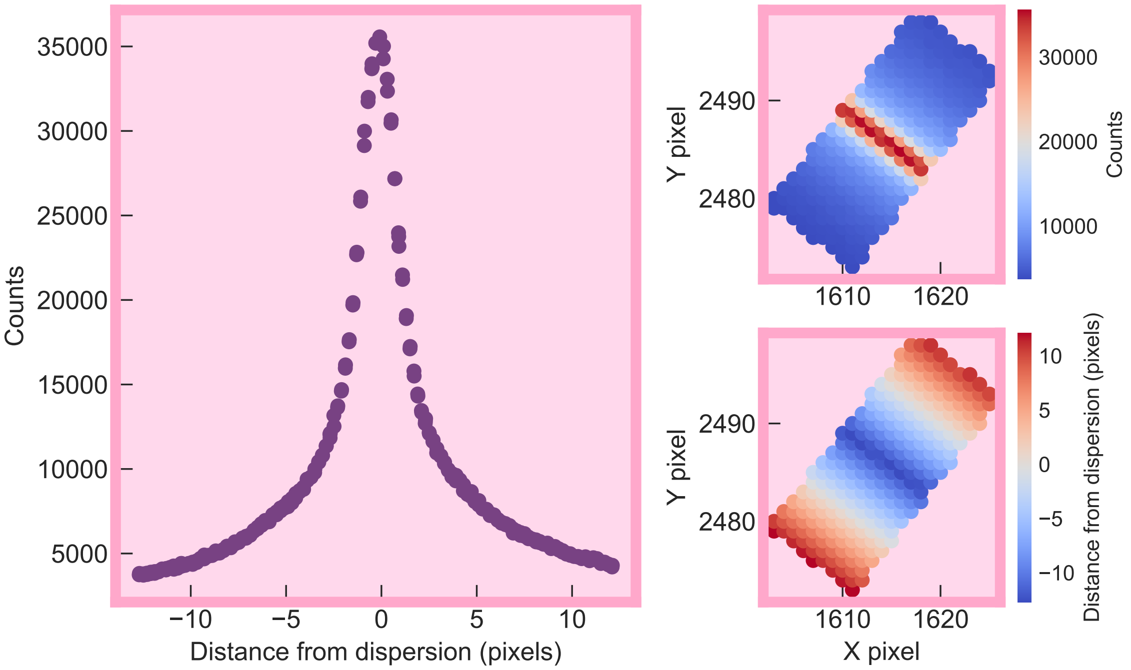

Figure 6.4: Illumination and distance from dispersion axis for pixels within a 2 Å bin. Left shows binned illumination as a function of distance from dispersion axis, top right shows distance of each raw pixel from the dispersion axis by colour, and bottom right shows illumination of each pixel by colour. As expected, the pixels with the lowest distance from dispersion axis have higher illumination, and vice versa — the light is focused predominantly on-axis.

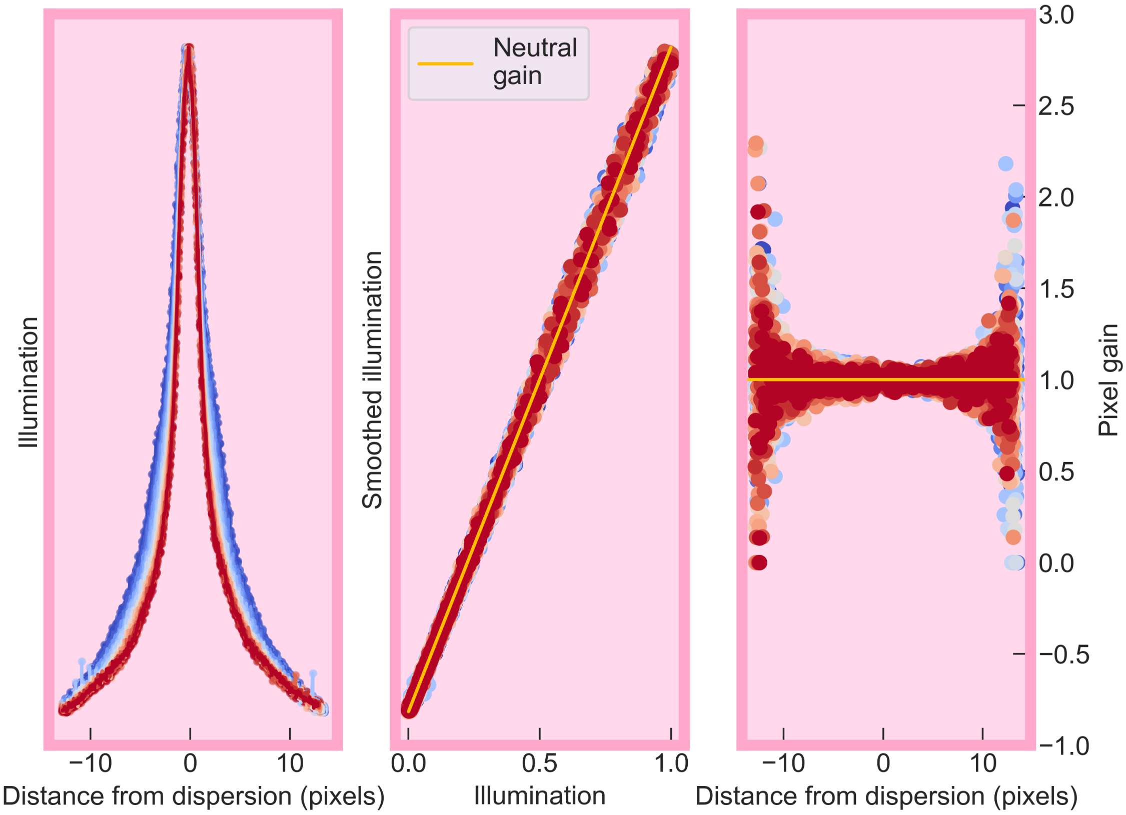

Figure 6.5: In these plots, each 2 Å wavelength bin is given a different colour, from red at the redder wavelength to blue at the bluer wavelengths. The left plot shows the cross-axis shape, which has a clear wavelength dependence due to optical distortions. The centre plot shows the illumination against smoothed values, which are distributed around the neutral gain line. The right plot is the resultant gain for each pixel as a function of distance from dispersion. The peak (zero distance from dispersion) is better sampled, which explains why there is a tightening around the neutral gain line at small (perpendicular) distances.

6.6 Improvement in spectral quality

This novel technique enhances the usability of output spectra by eliminating the source of interference previously produced in low signal-to-noise bins. Figure 6.6 shows the same observation as that in Figure 6.1, but using the new algorithm. There are striking improvements to the spectrum, while it remains unchanged near Hα. This is key to understanding the regime in which the problem was manifesting itself — low signal-to-noise bins. This subtle point had previously gone unnoticed in the ENDEAVOUR pipeline because it is very rare that the Aquila spectrographs observe a relatively faint target for such a short amount of time as this present example. This was in fact the very first spectrum of YZ Reticuli, and was used to gauge the exposure time needed to faithfully recover Hα, and that needed for the other, fainter lines in the spectrum. The very slight issue caused by discretisation of the cross-axis correction of the pixel gain map could only manifest itself in the absence of significant flux.

There is something truly remarkable about the improved spectrum shown in gold in Figure 6.6. Inspection of the upper panel reveals that while the main peaks and troughs of the purple spectrum have been eliminated, there is still a dip around the 7600 Å mark. Initially, I thought this was a further error, but then recognised it as a very strained detection of the telluric absorption band which stretched from approximately 7590 Å to 7680 Å. It serves as a strong validation of the new method, that this is visible in a three-second exposure of a distant target.

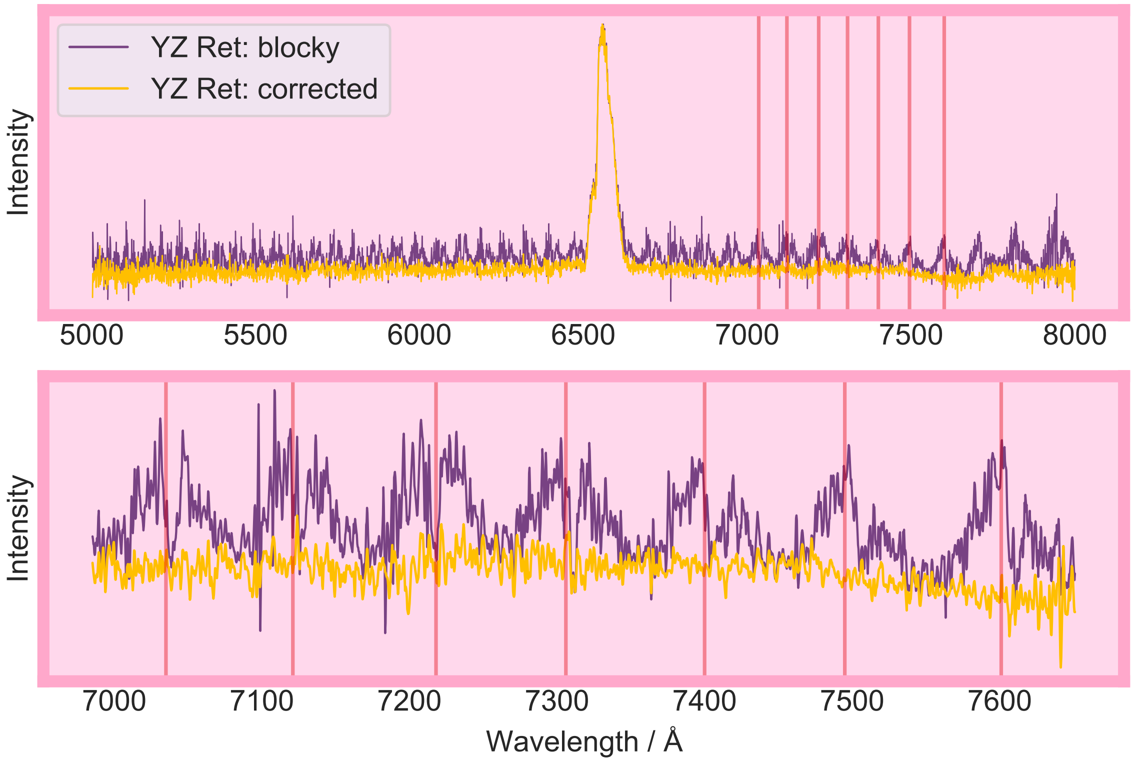

Figure 6.6: Corrected spectrum using the technique I developed in this chapter (gold) and using the existing method (purple). The upper panel shows the entire spectrum, while the lower panel shows a close-up on the blocky pattern as well as the improved and smoother version. The spectrum is of YZ Reticuli taken on 2020-07-16 at the GJW-CL observatory, is a 3 s exposure, and is the same for both methods of reduction.

Chapter 7

Machine learning for enhanced wavelength calibration

Wavelength calibration is the process through which ENDEAVOUR learns the wavelength of each pixel. It works by analysing a two-dimensional spectrum of a calibration lamp with a well-known emission line spectrum. In the case of the Global Jet Watch, this is a Thorium-Krypton bulb, hereafter ThKr. The end result is known as a wavelength map, which is a two-dimensional virtual frame with the same dimensions as the science frame, where the value stored at each pixel describes the wavelength of photons which contribute counts to that bin. It is important that the error on this is minimised where possible; under the Doppler shift 1 Å roughly corresponds to $45\,\mathrm{km\,s^{-1}}$ at the rest wavelength of Hα. It is also important that this process is automated, since the Global Jet Watch data repository contains over 25,000 ThKr calibration frames.

When moisture got into the CCD chamber of the ML8300 cameras, they were replaced with ML16200 ones, with many more pixels. Under the old paradigm for processing ThKr frames, the new chips needed to be calibrated manually so that the pipeline could be run again, because there were lines of ThKr in the chip beyond the wavelength range of the original calibration arcs. At this point, I suggested an improvement on the dated identification algorithm, so that it would become independent of chip size and geometry, instead using properties of the emission line spectra of Thorium and Krypton, since this is more future-proof. It is worth noting here that, as mentioned in Chapter 3, the Global Jet Watch Aquila spectroscopy frames are observed such that the dispersion axis lies roughly diagonally across the CCD detector, rather than parallel to the edges. This gives a larger wavelength range, and roughly ensures that the pixels are well-sampled.

Section 7.2 details the pre-existing algorithm for wavelength calibration, and Section 7.3 describes the limitations associated with that. Section 7.4 introduces my new algorithm to perform wavelength calibration using modern machine learning and image-processing techniques. Section 7.5 covers how the algorithm resolves the issues it was intended to solve, while Section 7.6 describes the limitations of the new method.

7.1 Scope of algorithm: understanding the problem

In Figure 7.1, I reproduce one of the calibration frames, which is the well-defined input to the algorithm. The frame is very sparse in terms of significant ThKr lines, which only appear on a tight band tracking from the lower-left to upper-right corner of the image, which is known as the *dispersion axis.* Close-up boxes show the structure of some emission lines as they appear in the raw image. The light is very sharply peaked at the centre of the lines, the pertinent quality which makes such frames ideal for calibration purposes. There are however some issues that need to be treated by any algorithm aiming to determine the wavelength calibration from these data. The left box in Figure 7.1 shows one of the lines from the longer wavelength end of the dispersion axis with a significant *halo* effect; there is a ring surrounding the peak. While it is clear where the centre of this peak is, consider the very faint peaks also enclosed in the box that also lie on the dispersion axis; these are of comparable brightness to features contained within the halo of the bright peak. For this reason, selecting the brightest points in the two-dimensional frame does not successfully find the maximum number of true peaks.

The right box shows a close-up of a pair of doublet lines, which can be used as an anchor, since they are the brightest peaks in most frames, and the ratio of the distances between peaks within a doublet to the distance between the two pairs is well-known, and physically constant. The emission line contained within the left box is also bright, and provides a long baseline across which to evaluate the wavelengths. However, the different observatories have had different detectors since being commissioned, and some of the detectors were set up not to include this wavelength range. This is a design feature of the Global Jet Watch network, to enable a slightly wider baseline of wavelengths, by having them offset relative to each other, while overlapping for the most part to enable cross-validation. Also, for very short exposure times, some weaker peaks become undetectable. The impact of these issues is that the algorithm cannot rely on always having particular lines present.

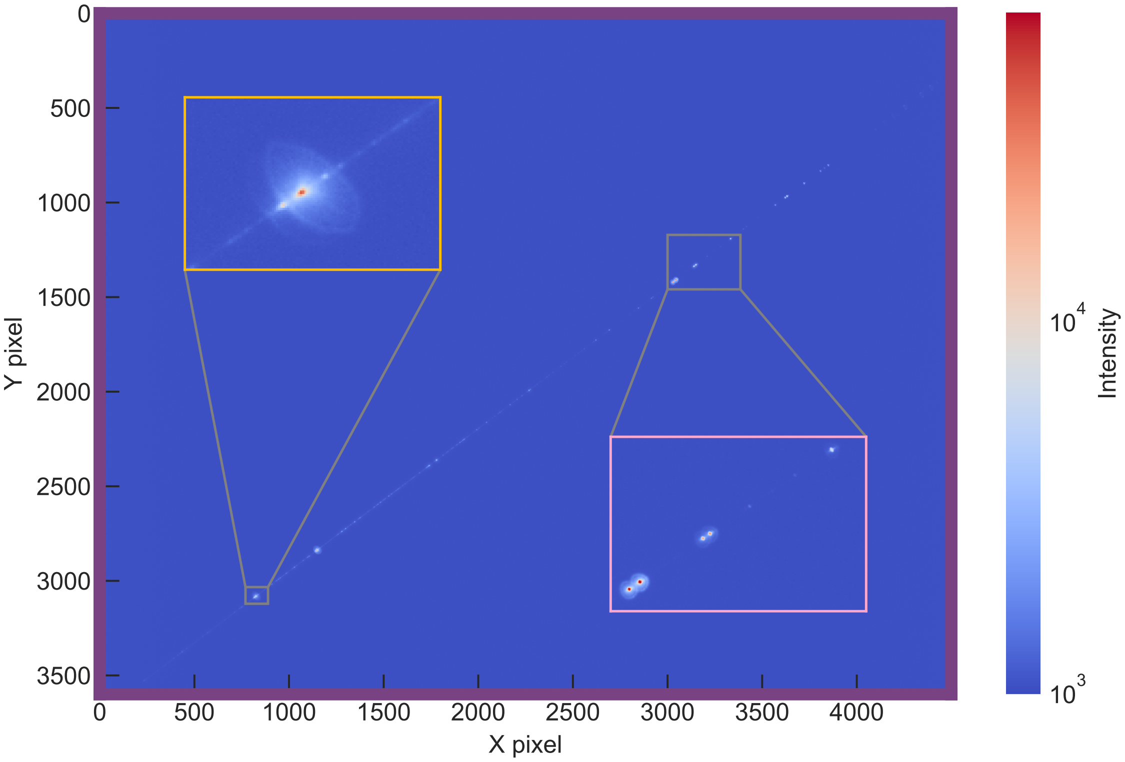

Figure 7.1: A raw ThKr frame, taken at GJW-CL with the Aquila instrument with an exposure time of 3 s. The left box (gold) shows a close-up of a very strong emission line with a particularly large halo. The right box (pink) shows a close-up of a pair of doublets, which are useful for locking onto as they have a very well defined wavelength spacing ratio.

There is also the issue of *saturation*. The CCDs used in the Aquila spectrographs are set to saturate at around 65,000 counts, just below the limit of 65,535 imposed by the analogue-to-digital converters, such that the number of counts in a pixel no longer increases no matter how many further photons arrive prior to the next readout. This leaves entire regions on the two-dimensional image where pixel counts are around 65,000 and information is lost about where the exact centre of a peak was.

The final piece of information available is the standard list of laboratory wavelengths available from the NIST database [Kramida et al., 2014]. This contains wavelengths and relative intensities, although some of the relative intensities are apparently derived from inconsistent sets, and as such they are not always reliable. The wavelength spacing between peaks correspond very precisely with those in the two-dimensional Global Jet Watch ThKr frames, and so those are deemed always trustworthy.

While the endpoint of the wavelength calibration phase is a full frame populated with wavelength values for each pixel, due to the sharp peaking in flux near the dispersion axis, in practice the wavelengths are only calculated for pixels within a predefined distance above or below the dispersion axis. For this chapter, I will refer to two polynomial mappings, one from x-coordinate to y-coordinate, and one from x-coordinate to wavelength (often abbreviated as λ), hereafter $P_{xy}(x)$ and $P_{x\lambda}(x)$ respectively, for points on the dispersion axis. Any off-axis pixel (i.e. all real pixels) has the same wavelength as the point on the dispersion axis from which a normal to the curve intercepts it. The modification I made to the algorithm only concerns the methodology up to the point of determining these polynomials, and so this chapter skips the final step which goes from the two polynomials to the full wavelength map. The general method of both the old and new algorithms is to find peaks in the two-dimensional frame, determine the coordinates of the centre of that peak and its wavelength, then fit for the polynomials.

7.2 Old algorithm

This section details the wavelength calibration which I wrote originally.

7.2.1 Search-and-Destroy

The previous algorithm used a recursive search-and-destroy algorithm on the two-dimensional ThKr data. This would *search* for the brightest point in the spectrum, which is most likely the brightest emission line, and define an island around it, large enough to contain the whole halo. Then, it would fit a two-dimensional Gaussian to the counts in the selected pixels, and record the coordinates of the centre. Next, it sets the pixels in the island to zero, thus *destroying* that peak. It would then iterate recursively until a peak either did not fall within a certain tolerance of a straight line, or a Gaussian could not be well fit to the island.

7.2.2 Identification

In order to find the true wavelength for each peak, the old algorithm used the assumption that while the pixel distances between peaks are unknown, the relative spacing between them should be reliable. It would produce a set of triplets of peaks, and compare the ratio of the distances between the points, $r = x_{\mathrm{ab}}/x_{\mathrm{bc}}$ where x_{ab} is the pixel distance between the first two peaks in the trio and x_{bc} is the pixel distance between the second two. Then, it compares the r values between all the trios in a reference identified frame to determine which peaks map to which wavelengths. This step relies on a human to manually populate the reference frame with known identifications of all the lines in the reference frame ahead of time.

7.2.3 Feedback search

The next step would take the identified set of ~ 10 peaks, and fit a low-order polynomial to them, between their x- and y-pixel centres, and between the x-pixel centre and the wavelength. This resulted in first iterations of the two mapping polynomials, $P_{xy}(x)$ and $P_{x\lambda}(x)$. The algorithm then searches the reference list of wavelengths, and solves for the xy-coordinates that each emission line should have, and checks the spectrum at that location. Each time it finds one, it repeats the process, thus refining the polynomial mappings. It concluded upon failure to add a peak. This typically allowed the algorithm to find at least double the number of peaks found in the original spectrum.

7.3 Limitations of the old algorithm

While the old algorithm represented a significant step forward compared to manually identifying each spectrum, there were still a host of reasons to update it. This section identifies some of those.

7.3.1 Time complexity

Computing ratios between locations within each set of three Gaussians scales with $\mathcal{O}(n!)$ where n is the number of peaks found. As the peak-finding algorithm improved over time and more peaks became available, this became very slow.

7.3.2 Slow operations

This algorithm used for-loops and repeated inefficient Gaussian fitting, which made it very slow. It would iterate over the whole frame and individually select pixels to be considered in the process, by determining its distance from the dispersion axis, rather than using filters to do things in a vectorised fashion. Since the pipeline is written in Python, where possible it is recommended to use standard image-processing tools, as these hook into low-level C code which is significantly faster.

7.3.3 Unnecessary repetition

The algorithm performed every operation twice as many times as necessary - once on the current ThKr and once on the master copy (manually identified) ThKr, and it did this every time there was a new frame to process.

7.3.4 Failure rate

This algorithm was also frequently unsuccessful, because it relied on fitting a Gaussian to the brightest peaks first to find a good initial fit to the dispersion axis. If the brightest peaks were saturated, the algorithm would break, and fail the arc frame in its entirety.

7.3.5 Accuracy

In order to compensate for saturation, the code had many arbitrary rejection criteria which would consider all the peaks and discard ones which did not fit well. While this has some merit, it meant that often too many peaks were discarded, which only weakly constrained the polynomial fit to the dispersion axis.

7.3.6 Manual identification

The old algorithm required a user to run it in a special mode as a one-off. This was difficult for two reasons - it required maintaining a piece of code that has two jobs, and it also was difficult to remember how to use because it only needed to be run again when a new detector was added to an observatory. The manual record was stored as a serialised Python object, which meant that it could not be easily read or edited, instead requiring the script to be run.

7.4 New algorithm

In this section, I explain the algorithm I developed to perform automated and fast wavelength calibration. I refer to pixel coordinates as X, Y and wavelength as W, to maximise clarity. The steps as stated are independent of any specific implementation; in the ENDEAVOUR pipeline, they are implemented in Python using standard packages including NumPy [Harris et al., 2020], SciPy [Virtanen et al., 2020], Pandas [Mckinney, 2010], scikit-learn [Pedregosa et al., 2011] and Astropy [Robitaille et al., 2013, Price-Whelan et al., 2018].

The general outline is to find all possible peaks (based on a prominence criterion), then use a clustering algorithm to select peaks which are grouped together; this selects those which are on the dispersion axis. At this stage, it is only necessary to have ten peaks in order to get a good determination of the determination axis, but off-axis noise is problematic. The algorithm then checks the brightest ten peaks to see if the spacings satisfy those expected of the brightest known peaks. If a match is found, it adds more peaks from the reference. A crucial difference here is that the new algorithm does not make any comparison to a manually identified master ThKr, but instead uses only the ASCII reference files (stored in .csv format). It also uses pixel-count weighted centre of mass to determine accurate centroids for each peak rather than fitting a Gaussian, which is not a good description of the halos, and is slow to compute.

Sections 7.4.1 to 7.4.4 detail the algorithm:

7.4.1 Find peaks

The first step is to find peaks caused by emission lines in the given two-dimensional ThKr spectrum. This is performed by first flattening the spectrum to one dimension, then removing background.

1. Calculate the mean and the trimmed mean (mean after discarding the highest 15 % and lowest 15 % of values) vertically
2. Divide the mean by the trimmed mean. The difference is due to the brightest pixels. (These first two steps are shown in Figure 7.2.)

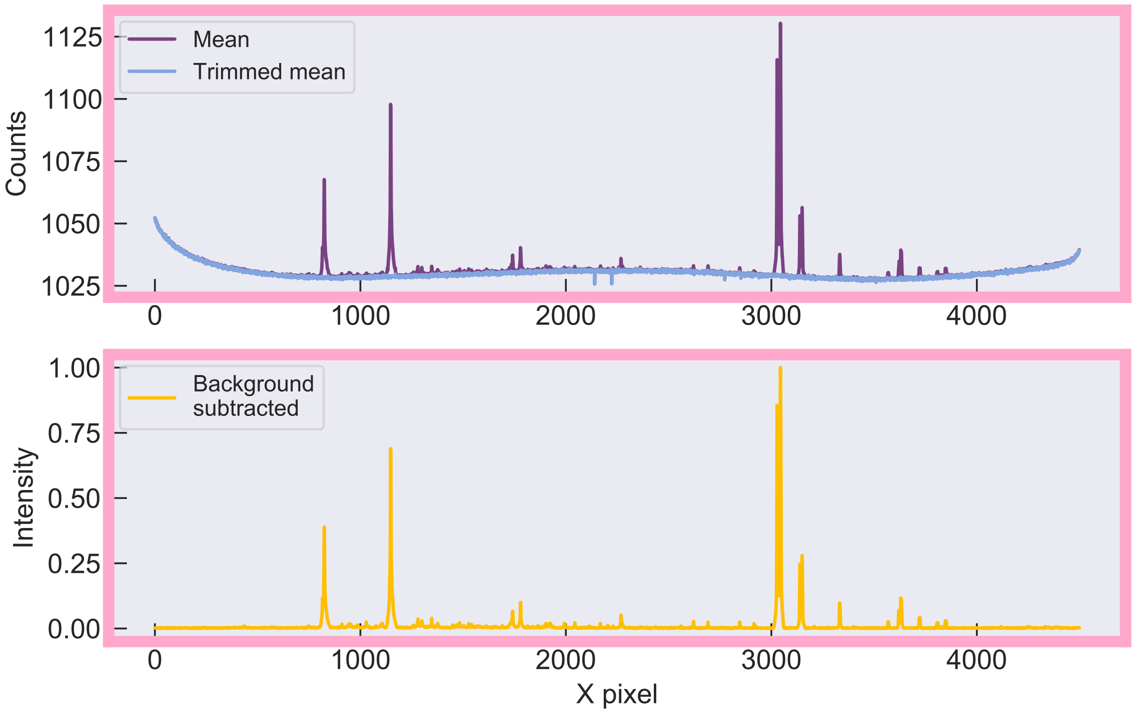

Figure 7.2: The background subtraction process. The upper panel shows a ThKr spectrum projected onto x-coordinate by mean and trimmed mean. The trimmed mean does not capture the peaks, and thus serves as an accurate estimate of the background. The lower panel shows the flattened one-dimensional ThKr spectrum, as the mean spectrum divided by the trimmed mean.

3. Locate the peaks in 1D, using a prominence test, record X for each.
4. Compute the associated Y for each peak, as the index of the maximum in the X column in the original 2D frame. (These two steps are shown in Figure 7.3).

7.4.2 Screen peaks

This step removes peaks using a clustering algorithm to ensure those selected lie on the dispersion axis.

1. Cluster the data (I use the DBSCAN algorithm provided by SCIKIT-LEARN [Ester et al., 1996]) with high acceptance criteria.
2. Remove any peaks marked as noise. These can happen if there is a bright off-axis pixel in the same column as a low peak value, or if a peak is isolated from other peaks. Record the removed peaks, as the isolated peaks lie on the

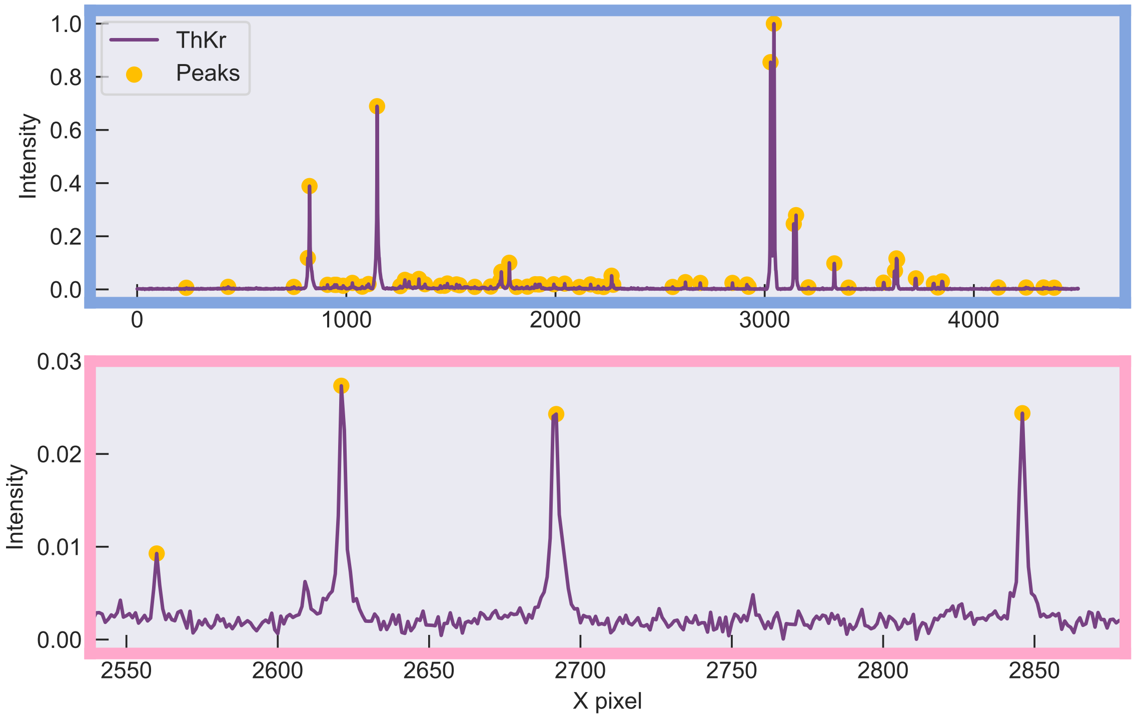

Figure 7.3: Automatically determined peaks, shown as gold markers, for the ThKr background-subtracted one-dimensional spectrum. The upper panel shows the entire dispersion axis, projected onto x-coordinate, while the lower panel shows an enlarged scale for some lower signal-to-noise peaks.

dispersion axis and therefore represent usable information. These first two steps are shown in Figure 7.4.

3. Apply Random Sample Consensus (RANSAC [Fischler and Bolles, 1987]) algorithm on remaining peaks to ensure any peaks that are not strictly near the dispersion axis do not affect the fit.

7.4.3 Refine

This step calculates the peak coordinates to sub-pixel accuracy. Figure 7.5 shows how, during this step, boxes are formed around peaks and specific pixels with values above the defined threshold are used to determine an improved estimate of the location of the peak centroid.

1. Select a box around each peak
2. Regrid the box to higher density by smooth interpolation (a numerical scale factor of 11 was used here)
3. Select pixels within the box which exceed a threshold level of 0.9 relative to the peak
4. Calculate weighted centroid ("centre of mass") of these pixels and return sub-pixel coordinates

7.4.4 Determine polynomial mappings

This step calculates fitted polynomials $P_{xy}(x)$ and $P_{x\lambda}(x)$ that describe the dispersion axis and the wavelength dependence respectively.

1. Assign W to the 5 brightest peaks using their X separations to form five pairs of X-W
2. Use first-order linear regression to determine an initial estimate for $P_{x\lambda}(x)$ from the five pairs
3. Find the next not-yet-found brightest peak in the reference data, solve $P_{x\lambda}(x) = W$ to compute its predicted X
4. Find the brightest peak in the 1D spectrum within approximately 50 pixels of this prediction, and assign it the wavelength W

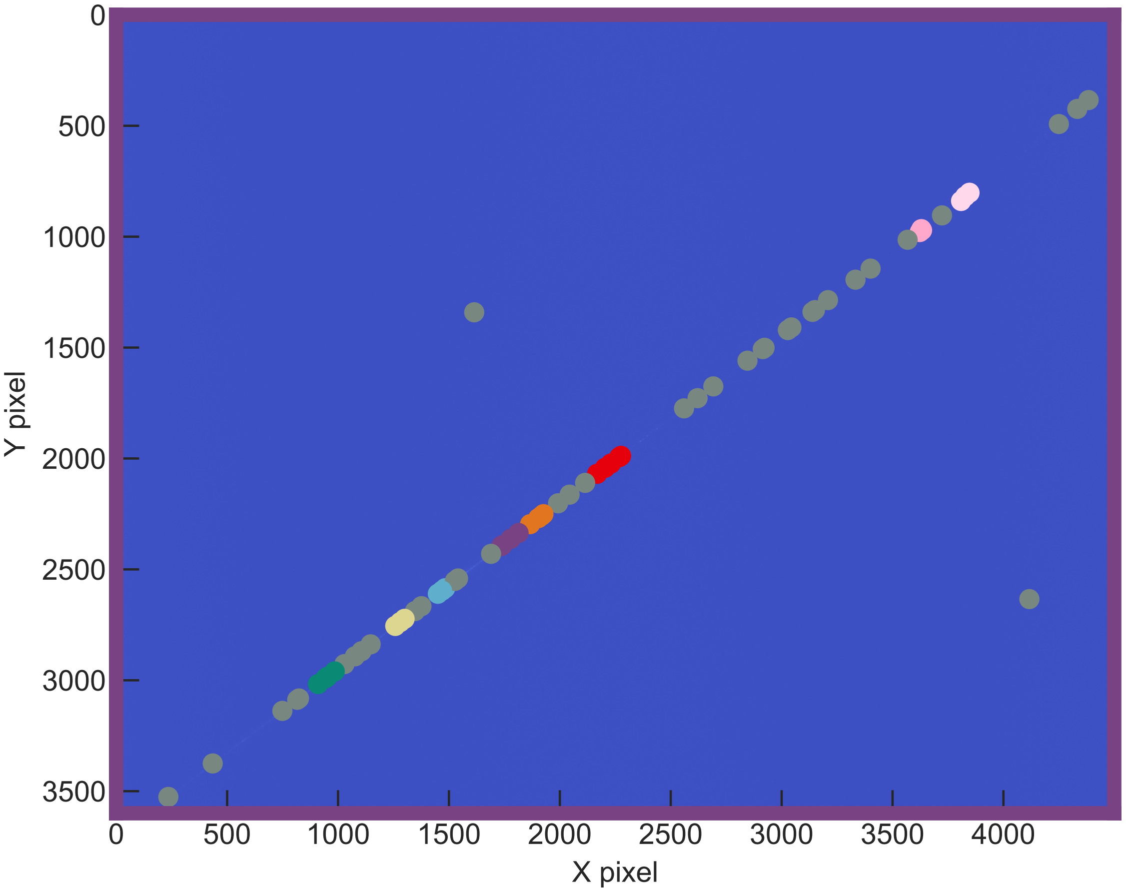

Figure 7.4: Clustered peaks are shown in bright colours, with rejected points (anything remotely noise-like) shown in grey. The present implementation uses the DBSCAN algorithm with the epsilon radius parameter set to 50 pixels and the minimum number of samples set to 3. This means for two points to be considered neighbours, they must be no further than 50 pixels apart. A point must have at least 3 neighbours to qualify as a core point. If this does not successfully identify the dispersion axis, the run is re-tried with different values for these parameters.

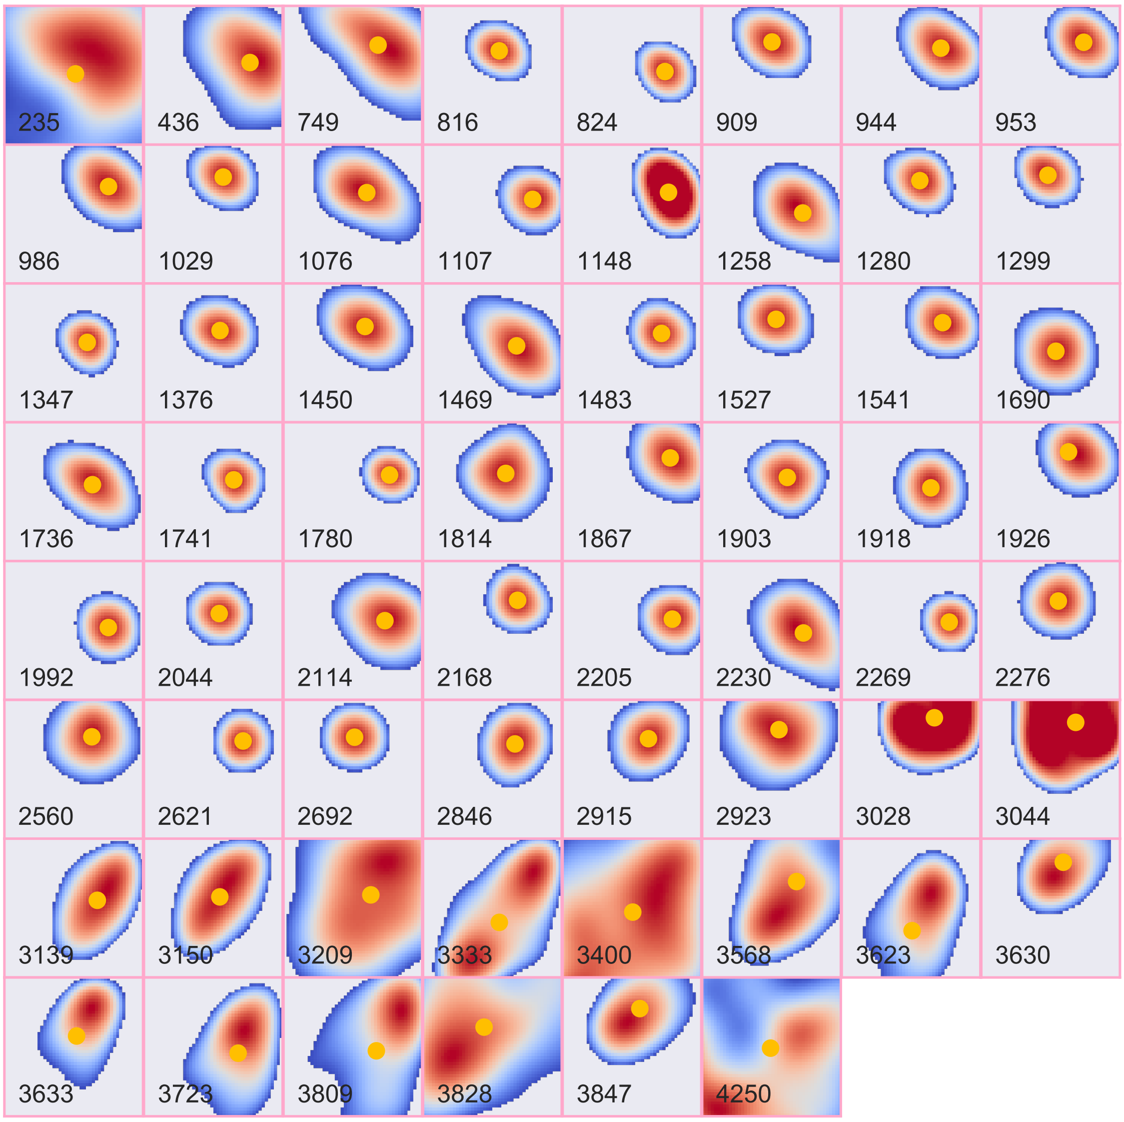

Figure 7.5: Centre of mass of local islands to refine the central coordinates for each peak. Central x-coordinate (X) is marked inside each box, which measure seven pixels by seven pixels each and are interpolated by a factor of three to increase resolution. Note that a lower threshold value was used to generate this figure compared with that actually used in the pipeline, to better demonstrate the principle.

5. Repeat the previous 3 steps until there are 15 peaks with known wavelengths (output shown in Figure 7.6)
6. Fit a linear regression model from the set of X to the set of W, to improve estimate of $P_{x\lambda}(x)$
7. For each of the remaining X coordinates of peaks in our spectrum, predict a value for their wavelength W
8. For each peak, select the line in the reference with the nearest wavelength, and (tentatively) assign that wavelength to the peak
9. Exclude peaks which are above a threshold wavelength (as these are known to be out-of-focus)
10. Fit for $P_{x\lambda}(x)$ and $P_{xy}(x)$ and reject outliers, then fit again for the final versions.
11. Return final $P_{xy}(x)$ and $P_{x\lambda}(x)$ as output

7.5 Advantages

Overall, the new algorithm outperforms the old one in all of the areas outlined earlier, and I discuss the improvements in this section.

7.5.1 Accuracy

From an astrophysical perspective, the most important property of the wavelength calibration algorithm is the accuracy. In this case, it is measured as the standard deviation of the discrepancy between the wavelength value assigned to each peak and the wavelength value predicted by the algorithm for that peak. We reject any frame below a certain tolerance accuracy level.

7.5.2 Time complexity

Every step of the new algorithm scales at worst $\mathcal{O}(n)$ where n is the number of peaks, which is physically capped by the number of emission lines of ThKr in the wavelength range of the detector.

7.5.3 Automatic

The algorithm requires no manual identification.

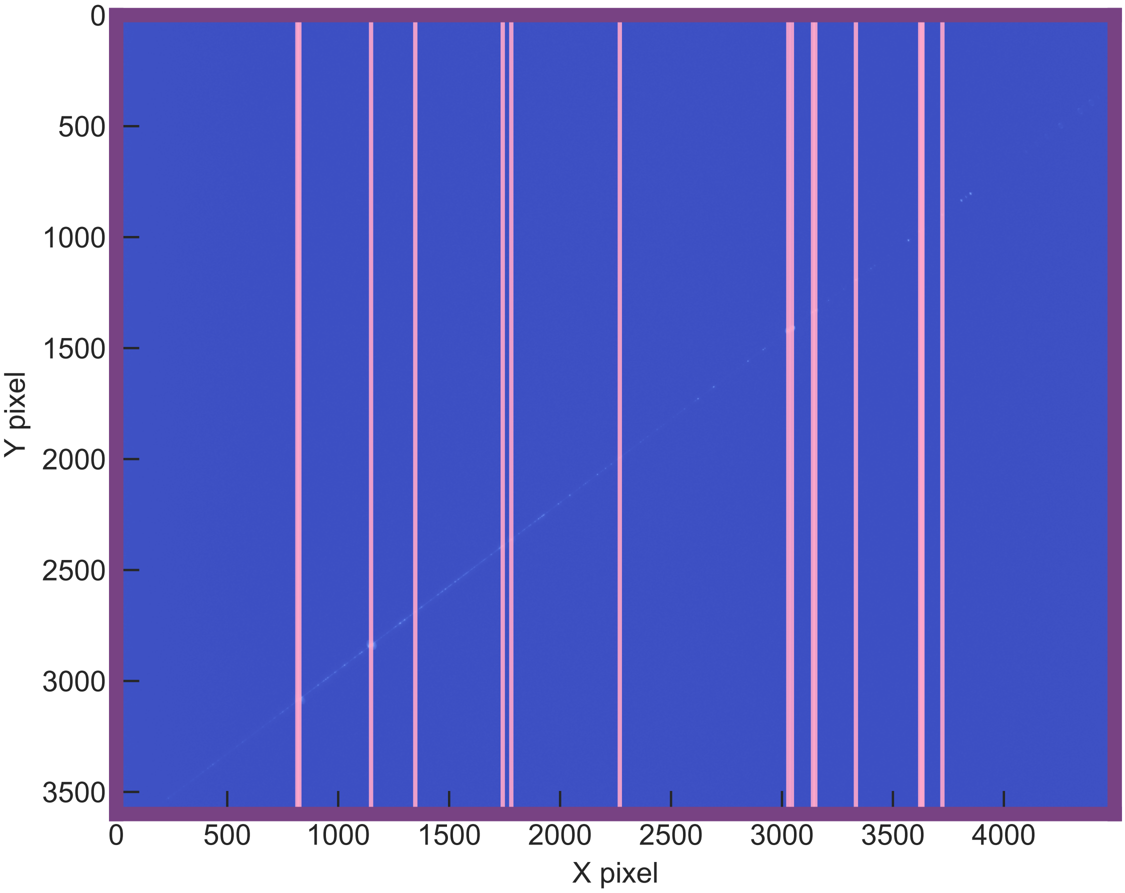

Figure 7.6: This is the two-dimensional ThKr data, with the x-coordinate of each of the 15 peaks added from the ThKr line reference file indicated by a vertical line.

7.5.4 Speed

As the project grows, the number of spectra needing to be processed will increase, and the number of calibration frames with it. The current requirement is that the average time taken for a single spectrum should be on the order of a few seconds. A typical observing month might contain around 10^3 ThKr frames, and a few hours is reasonable. This algorithm takes on average 0.68 s to process a single frame, which is inconsequential compared to the overall running time of ENDEAVOUR.

7.5.5 Reliability

To pass, the algorithm must find at least 10 identified peaks, while the standard deviation of residuals between identified wavelengths and the $P_{x\lambda}(x)$ prediction for that peak should be below 0.5 Å. It also must successfully reject frames which are not ThKr frames. For example, a dark frame should fail. This standard is achieved, and exceeded for $> 99\,\%$ of ThKr frames of sufficient exposure time. It is, by design, detector independent, which should increase reliability in the future also.

7.5.6 Robust fitting

Before fitting, the algorithm excludes peaks in regions of the chip where the optical focus is known to be poor. Using either robust least squares fitting or the RANSAC [Fischler and Bolles, 1987] algorithm, which separates points into 'inliers' and 'outliers' and then rejects outliers, offers little to no improvement over the base algorithm (detailed in Section 7.7), which implies it already successfully treats outliers.

7.6 Limitations of the new algorithm

If two peaks are near each other, the extra flux from the second peak in the box can lie above the threshold in step (see above). This is not too problematic, since the centres of mass of these pixels will still lie on the dispersion axis, and thus $P_{xy}(x)$ is not affected. However, the $P_{x\lambda}(x)$ would be (minimally) affected if the wavelength difference between the doublet centre of mass and the actual assigned wavelength was below the threshold for error (if it was above the error threshold, then it would not count). Given there are usually many more isolated peaks than paired peaks, this is not a problem in practice.

7.7 Results of a run

The key outputs to be understood are the two polynomial mappings, $P_{xy}(x)$ and $P_{x\lambda}(x)$ to y-coordinate (measured in pixels) and wavelength (λ, measured in Å). As is apparent upon examination of Figure 7.1, the dispersion axis is well-fit by a straight line. However, the same is not necessarily true for the wavelength mapping, since the optical set-up may incur wavelength-dependent aberrations. In this section, I provide quantified analysis of these mappings.

7.7.1 Shape of dispersion axis

The shape of the dispersion axis on the detector is described by the $P_{xy}(x)$ mapping, relating the x-coordinate of a pixel on the dispersion axis to its corresponding y-coordinate. Given the set of $\{x, y\}$ pairs, there are many techniques for retrieving the 'best' estimate for the underlying relationship. 'Best' here is subjective, as it depends on how the outliers are treated. Figure 7.7 details an experiment to determine the most reliable fitting routine to use for this context. The first panel simply shows the $\{x, y\}$ pairs, and a rough estimate of a straight-line fit. The subsequent panels in turn consider three different algorithms for this low-order polynomial fitting. The second panel of Figure 7.7 uses a standard least-squares approach to fit first- and second-order polynomials to the data, achieving a standard deviation of 0.31 pixels and 0.30 pixels respectively. Since the residuals show no trend or structures in either order, I conclude that a straight-line fit is sufficient to describe the dispersion axis.

To see if there was any improvement by treating outliers differently, I then confronted the data with two further techniques of low-order polynomial fitting. The third panel uses soft l1 loss (which measures the mean absolute difference between the data and prediction, rather than standard least-squares), which shows no deviation from the second panel, while the RANSAC algorithm here performs worse than the second panel. This is most likely because all the outliers have been successfully removed already by the iterations within the new calibration algorithm presented in this chapter. I therefore selected the technique shown in the second panel to be used to compute $P_{xy}(x)$.

7.7.2 Wavelength dependence on x

The results of a typical run are shown in Figure 7.8. The top panel shows the refined peaks, showing excellent coverage of the detector; beyond x values of around 3000, the focus is somewhat blurry so these peaks are excluded to maximise accuracy. The

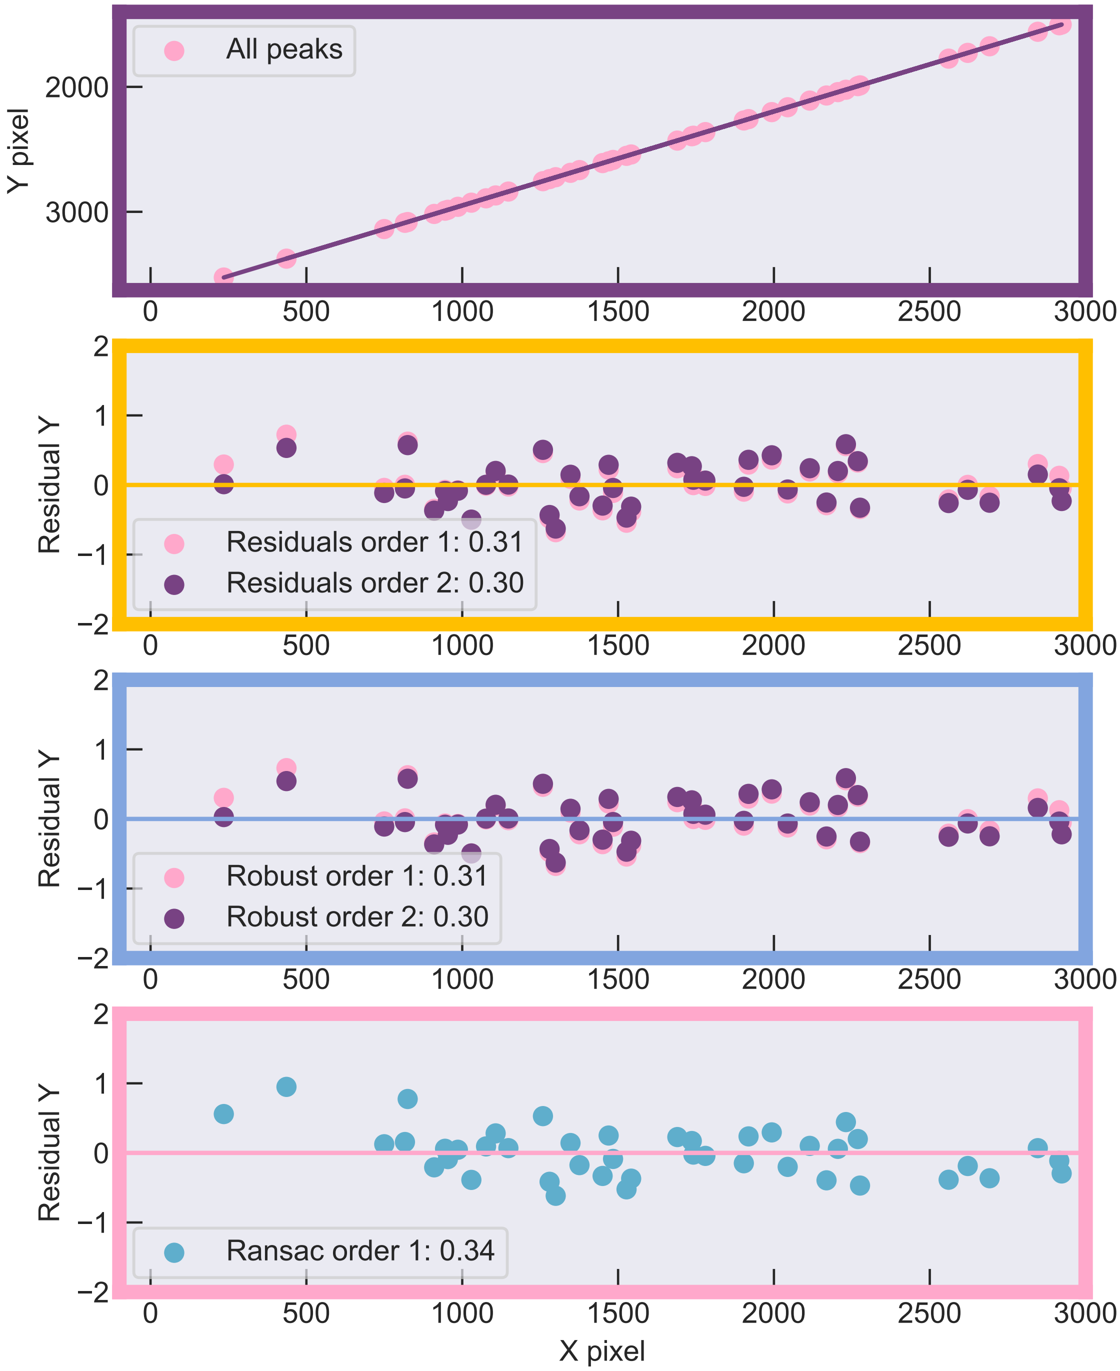

Figure 7.7: Results of several different methods of polynomial fit to the xy-mapping. The upper panel shows the locations of peaks used, along with a first-pass estimate of the mapping. The second panel shows standard least-squares fitting using linear loss. The third panel shows standard least-squares fitting with soft l1 loss, described in the text. The third panel shows the result of using the RANSAC [Fischler and Bolles, 1987] algorithm.

second panel shows the first-pass linear fit to $P_{x\lambda}(x)$, with the peaks marked. Despite this appearing like a good fit at first glance, the standard deviation of residuals (the discrepancy between reference wavelength and predicted wavelength ($=P_{x\lambda}(x)$) was very large, at 6.3 Å. I use sequentially increasing orders of this polynomial to fit the wavelength distribution, and the results are presented in the lower panel of Figure 7.8. The systematic error in residuals is corrected for by using a third-order (cubic) fit, and so this is the final functional form that the new algorithm returns. This method achieves an accuracy of below 0.15 Å, as measured by the discrepancy between the known wavelengths of ThKr emission lines and their values determined through this approach.

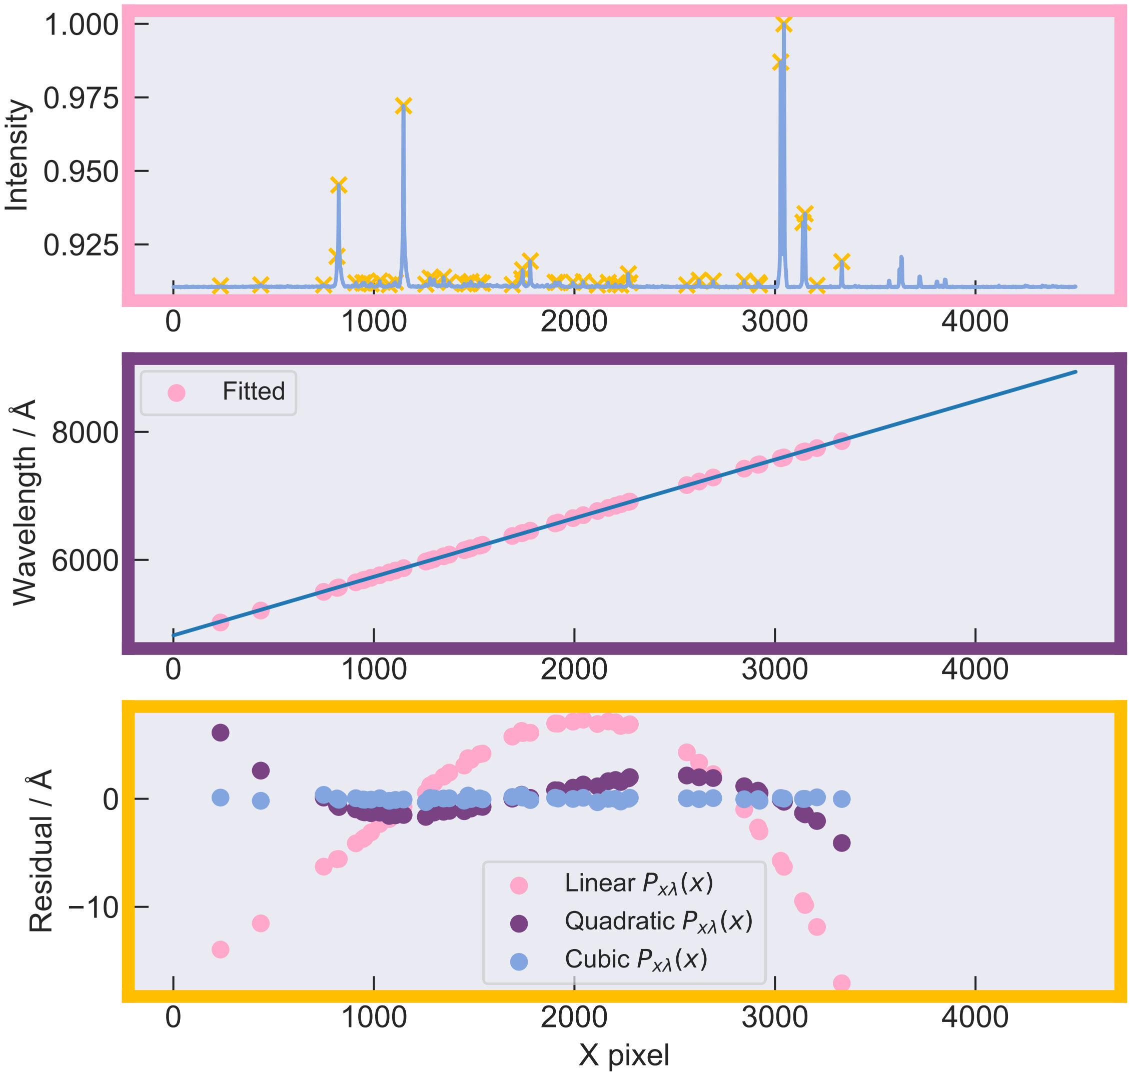

Figure 7.8: The upper panel shows the flattened spectrum with the final accepted peaks marked. The middle panel shows a linear first-order approximation to $P_{x\lambda}(x)$, with matched peaks marked. Note that even with a first-order fit, the discrepancy is minimal. The third panel shows the reduction in residuals (difference between actual wavelength and fitted wavelength for each peak) attained by using a higher-order polynomial to fit $P_{x\lambda}(x)$. The standard deviations of these sets of residuals are 6.3 Å, 1.6 Å and 0.12 Å when using a first-, second- and third-order polynomial respectively. There are no significant improvements beyond a cubic, so this is selected as the final calibration polynomial, with an accuracy of 0.12 Å in this case. Generally, the algorithm achieves an accuracy of 0.15 Å.

Chapter 8

High Resolution

The Global Jet Watch observatories were originally designed to perform highly time-resolved spectroscopy at moderate resolution. In this Chapter, I discuss the installation of a new high-resolution instrument, which is used for targeted one-off observations to supplement the main time-series data. For testing, it was installed at the Rainbow observatory during July 2020, and will presently be permanently housed at one of the Global Jet Watch observatories. In Section 8.1, I discuss the science case for the new instrument, focusing on three main use-cases. Section 8.2 covers the installation of the physical apparatus, and explains how such an addition is possible. In Section 8.3, I discuss the design of a new Python package within ENDEAVOUR, which controls the reduction of high-resolution observations, and in Section 8.4, I describe the algorithm I developed.

8.1 The science case for high resolution

The *modus operandi* of the Global Jet Watch observatories and the Aquila spectrographs is to produce successive sequences of time-resolved spectral data, so the benefit of performing relatively sparsely time-sampled spectroscopy is the extra resolution provided by the new instrument. The typical resolution of the spectra the Aquila instruments at Global Jet Watch observatories were designed to have is $R \sim 4000$ (in order to give the required wavelength coverage), while the Shelyak LHires instrument has $R \sim 20,000$. A particular requirement was to be able to significantly increase the spectral resolution of the Hα complexes seen in the targets. This could be achieved with this off-the-shelf spectrograph used together with a moderately priced camera having a large enough detector to give a wavelength range of around 130 Å, which equates to nearly $6000\,\mathrm{km\,s^{-1}}$ at Hα. A specific wish was to be able to measure this complex at higher spectral resolution to confirm the fidelity of the fits to the moderate-resolution Aquila data. This section discusses the ways that the new data stream can benefit the Global Jet Watch ecosystem.

8.1.1 Confidence

Detailed high-resolution snapshots are useful to complement densely time-sampled data at moderate resolution. A spectral feature of particular interest in the study of classical novae is the Hα line, and as such it is highly instructive to consider its evolution throughout the post-eruption development of a new system. In Chapters 10 and 11, I demonstrate how Gaussian fitting of the Hα line over time provides insights into the dynamics of classical nova systems.

8.1.2 Curiosity / digging deeper

It is helpful to have the ability to explore different regions of the spectrum as they emerge and develop following the eruption of a classical nova. The early absorption line systems of classical novae around their optical maximum are complex, blended, and varied. The addition of the LHires instrument will enable Global Jet Watch observatories to facilitate the study of such absorption systems, which can help provide insights to the nature of classical nova ejecta.

8.1.3 Galactic Supernova

In the eventuality of a highly anticipated Galactic supernova eruption, agility and responsiveness are key. The ability to quickly process (within minutes) and plot high-resolution spectroscopy would enable faster adjustments to the rapidly-varying observational parameters, and allow close inspection of specific spectral regions of interest. A future objective for the LHires when installed at Global Jet Watch observatories is to have fully remote-controlled wavelength agility.

8.1.4 Slit-spectroscopy advantage

The Global Jet Watch has traditionally focused on Galactic targets situated outside the Solar System. However, the agile network of telescopes are also used for a multitude of other purposes - education, outreach, exploration (discussed further in Section 8.2). The new LHires instrument is a long-slit spectrograph, which means that for certain spatially resolved targets it is possible to extract multiple spectra for the different physical regions of the target. Many modern spectrographs are fibre-fed (including the Global Jet Watch Aquila ones); a single optical fibre cannot give spatially resolved information.

8.2 Implementation

This section describes the physical deployment of the instrument, and the plan for its imminent installation at the Global Jet Watch observatory GJW-OZ in eastern Australia.

8.2.1 Mechanical details

The Shelyak LHires instrument is a slit spectrograph, which uses a neon lamp for wavelength calibration arcs. It uses a ZWO ASI camera (model ASI294MM) cooled to $-15°$ C. While the slit width is selectable between several preset widths, it was set to 23 µm for the observations presented in Chapters 10 and 11. The central wavelength for an observation with it is tunable, and is recorded in the FITS header under the keyword CENTWAVE.

8.2.2 Empty slot

A fundamental principle of the Global Jet Watch observatories is that they are mostly housed in residential schools, where the pupils can use them for imaging before local bedtime. This is made possible using an Optec Perseus 4-slot Instrument Selector (shown in Figure 8.1) and IP switches that control the supply of DC power to all the various devices. Section 3.3 describes how the code I wrote, GADGET, is used to enable remote control of which port the light from the telescope is directed to. The LHires instrument will fit on the available port three of the instrument selector as shown in Figure 8.2. And so, given the strong science case for LHires in the study of classical novae, I suggested the addition of such an instrument to an empty slot of the 4-slot selector. After observing a target throughout the night, and without adjusting the telescope acquisition, it will be possible to remotely switch to the LHires instrument to capture a detailed snapshot.

8.3 Software design

In this section, I discuss the key design principles which guided the development of a new pipeline to reduce LHires data from Global Jet Watch observations.

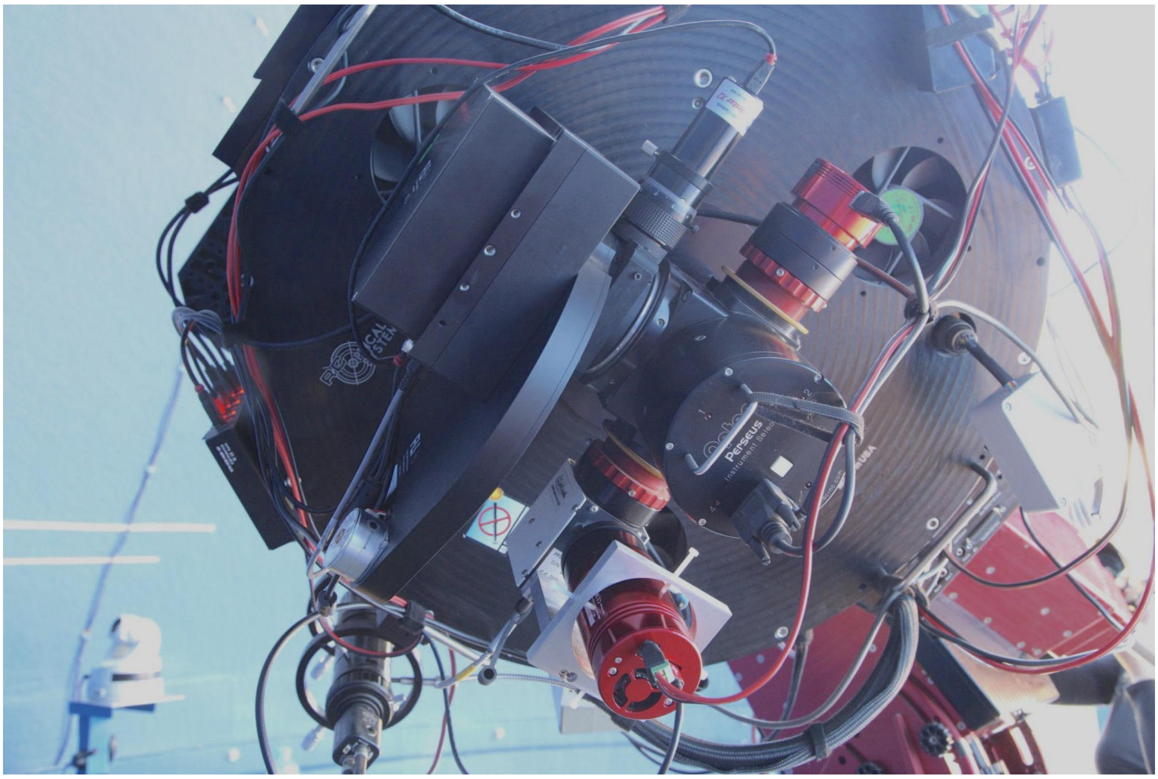

Figure 8.1: Perseus instrument selector, which can be rotated to direct the light path from the telescope to one of four ports, upon each of which is mounted a different instrument. This is controlled remotely, with imaging and spectroscopy configurations available.

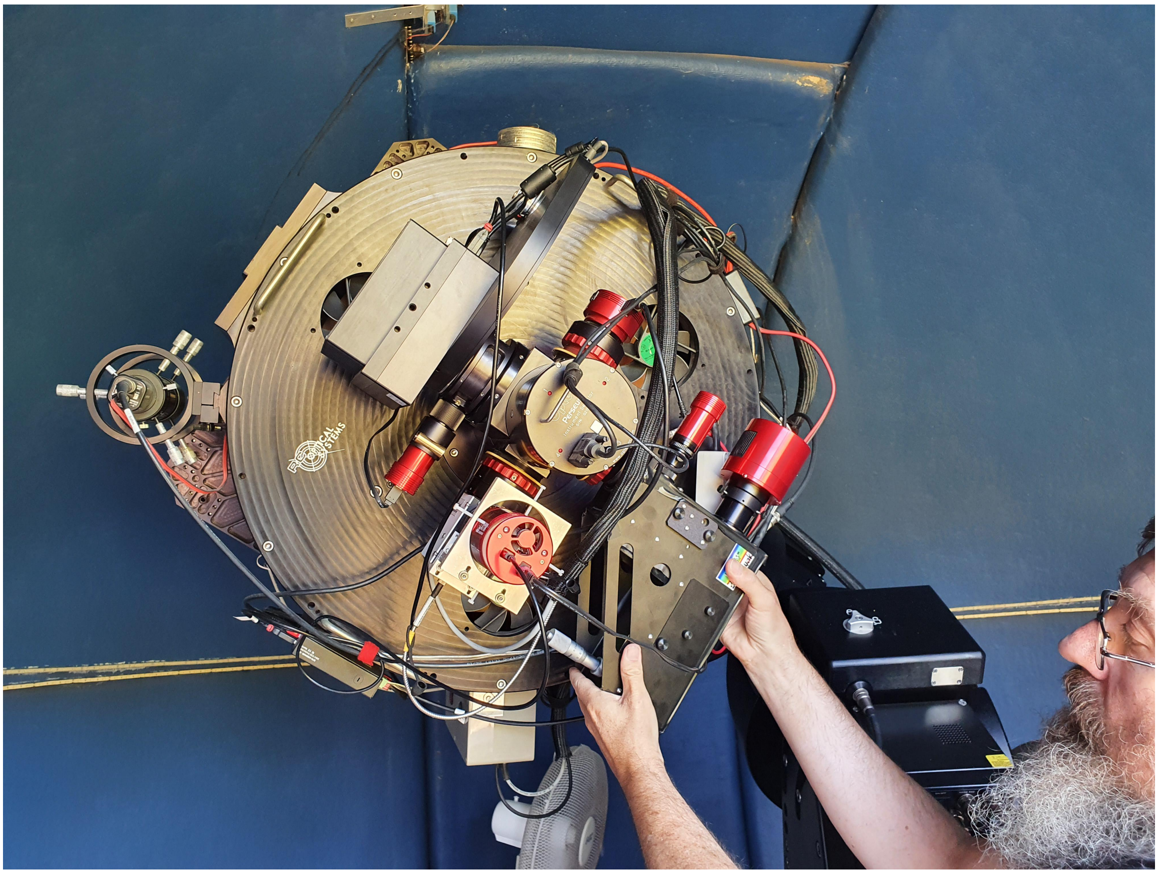

Figure 8.2: LHires is a much smaller and more lightweight instrument than Aquila, and it is manufactured to be mounted on the telescope. While flexure is more of a problem at higher resolving power, Global Jet Watch LHires science frames are always calibrated with an adjacent-in-time arc in order to match the flexure.

8.3.1 PEP8

The entire LHires data pipeline is compliant with the PEP 8 guidelines for Python. This is the style guide which, if adhered to, ensures internal consistency within a codebase and external consistency between projects.

8.3.2 Dependency injection

Dependency injection is a programming pattern, whereby code exposes an interface which allows dependencies to be passed into it wherever it is called from rather than creating them. This allows for code which does not create its dependencies within itself, but instead assumes that it will be given a dependency with a certain predefined interface. This is particularly useful for implementing the strategy pattern, explained below.

8.3.3 Strategy Pattern

Astronomical observations, and the accompanying calibration frames, vary considerably. While a simple algorithm can often perform well over large subsets of the entire data set, there are often idiosyncrasies of particular observations that require some tweaking of parameters involved in the reduction process. As alluded to in the previous subsection, to allow for this eventuality, I use dependency injection so that the main calling script can create a different algorithm controller to pass into the pipeline for different observations. For example, Chapter 6 describes an alternative treatment necessitated by science frames in the particularly faint (i.e. low signal-to-noise) regime for the Aquila data stream. The various algorithm controllers must expose a common interface, in this case so that the pipeline can call an EXTRACT method on each. Crucially, the pipeline does not need to know which particular controller it is operating on. This allows for a flexible and clearly structured approach which greatly simplifies the pipeline.

8.3.4 Interface directly to Poirot

POIROT, introduced in Section 3.6, is the code package responsible for plotting the outputs of the data pipelines. The LHires pipeline was designed with this in mind, and writes to a definable directory which POIROT can read from. As soon as the pipeline has completed, POIROT can instantly display the high-resolution 1D output spectra. The class representing a spectrum in POIROT has a method which describes which data stream it originated from, and the class representing a sequence of such spectra can instantly filter them based on this property. This agility is extremely useful, allowing insights on sequence composition, such as the relative number of spectra from each observatory, and instrument, as well as facilitating simple checks on wavelength calibration, and comparison of spectra. It also makes it straightforward to apply the same template models found to work for one data stream to others, satisfying the goal of Section 8.1.1. The fact that the same template used for the time-sequence Aquila data also closely fits the LHires observations represents an important orthogonal sanity-check. A future useful test could be to compare Hα with Hβ, to separate the hydrogen behaviour from [N II] 6584 Å emission, which is often blended with broad Hα emission during the nebular phase of classical novae.

8.4 High-resolution algorithm

This section details the technique used to reduce data from the LHires instrument. It detects new two-dimensional raw frames in the user-defined input repository and writes calibrated one-dimensional spectra to the user-defined output repository. The procedure is broken up into three main sections, which bear some resemblance (albeit with differences in implementation) to those in the main ENDEAVOUR pipeline, discussed in Chapter 4. These are: 'identify', to determine the wavelength calibration using lamp arcs; 'dispersion', to estimate the location of the main dispersion axis; and 'extract' to estimate the relative detected intensity in wavelength bins along the spectrum.

8.4.1 Identify

The first step of the algorithm is to identify which calibration arcs should be used for which science frame. In the case of ENDEAVOUR, this represents a challenge, because the Aquila instruments prioritise high time-sampling during the night, and takes calibration frames during the day described in Chapter 5. However, for targeted and sparsely time-sampled LHires spectra, the observing plan includes taking specific calibration frames before and after each observation, rendering this issue trivial. This is important for the LHires instrument, because there is significant flexing according to the altitude and azimuth angles of the telescope.

Figure 8.3 shows the CCD pixel values of a typical neon lamp arc used for wavelength calibration. The wavelengths of these strong emission lines are known from lab measurements available from NIST [Kramida et al., 2014]. The dispersion axis is horizontal, and variation in the vertical is a slit-spectroscopy effect giving spatial

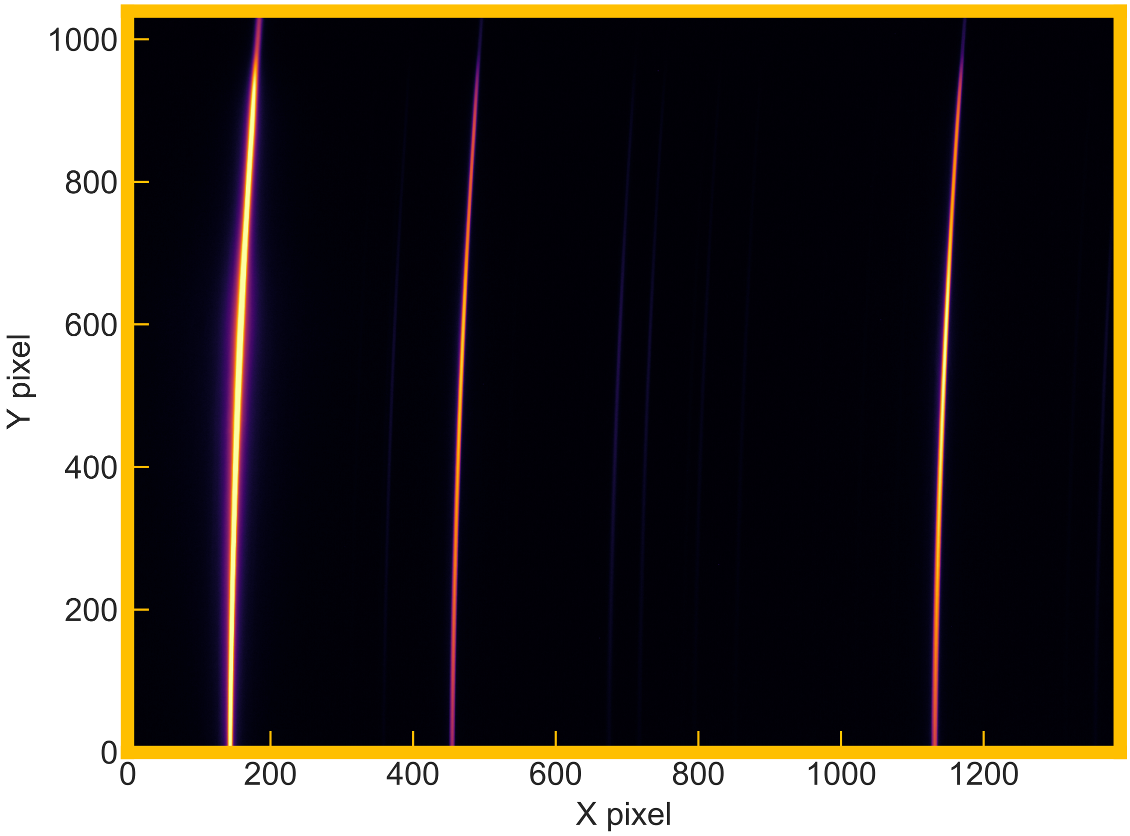

Figure 8.3: Spectrum of light taken from a neon arc captured with the LHires instrument, with an exposure time of 10 s on 2020-07-15.

information. Section 8.5 discusses manipulating this fact to extract several spectra from one observation.

To determine the wavelength calibration, the algorithm first selects the brightest points in the frame, and then clusters them using the SCIKIT-LEARN [Pedregosa et al., 2011] implementation of the DBSCAN algorithm [Ester et al., 1996]. As depicted in Figure 8.4, this results in a set of groups of pixel coordinates, which are separated by wavelength. There are multiple controllers with different configurations, and the main algorithm iterates through controllers until it finds one which gives the expected number (and wavelength separation ratio) of clusters. The target wavelength of the set-up is predetermined to match a science goal, in this case Hα with a wavelength of 6563 Å, and this target wavelength is recorded in the metadata associated with the observation. The algorithm knows, given the target wavelength, which neon lines to expect given the grating and the detector information in the original FITS header, and this is what it uses to assign a label to each of the clusters. Figure 8.4 presents this successful identification in the legend.

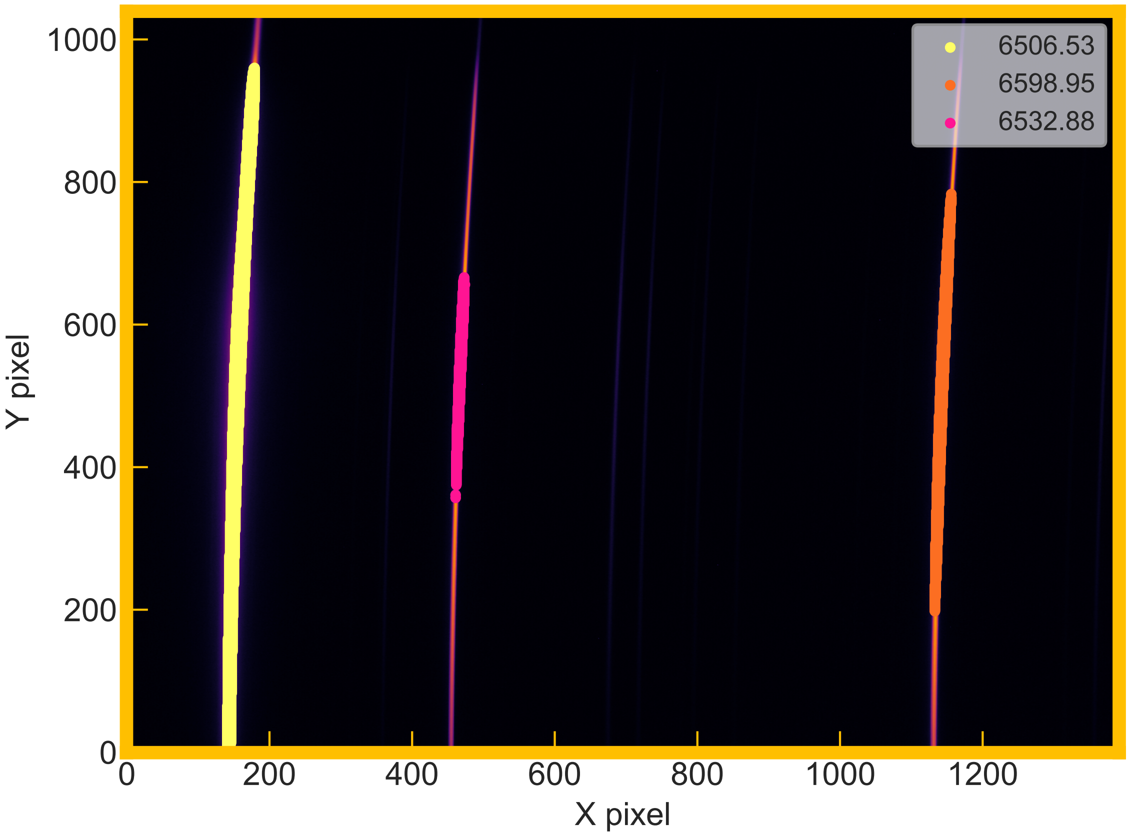

Figure 8.4: Neon data with arc emission lines identified, showing wavelengths ordered by the number of points included in the relevant cluster.

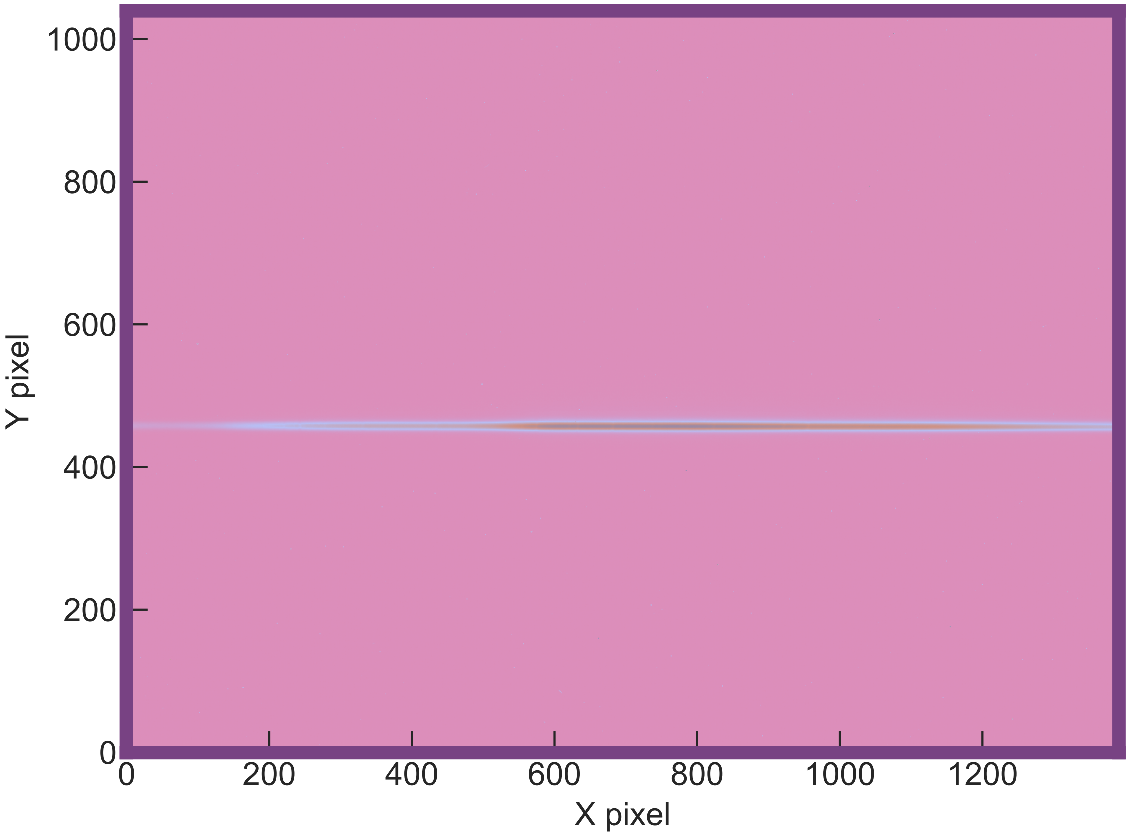

Figure 8.5: LHires two-dimensional spectrum of spatially unresolved science target, classical nova YZ Reticuli. Taken on 2020-07-17 with an exposure time of 1000 s.

8.4.2 Dispersion

The algorithm now proceeds to process the science frame needing to be reduced. Figure 8.5 represents an observation of YZ Reticuli taken using the LHires instrument on 2020-07-17, and shown in a different colour map from the arc to avoid confusion. This target is not spatially resolved, and as such has only limited scatter in the vertical direction, caused by the point-spread function of the optics.

To determine the dispersion axis in an automated fashion, the algorithm simply histograms the entire frame horizontally, shown as blue circles in Figure 8.6. There is a sharp, narrow peak in this histogram, which the algorithm fits a Gaussian to, shown in gold in the figure, and then passes the fitted parameters of this on to the extraction step.

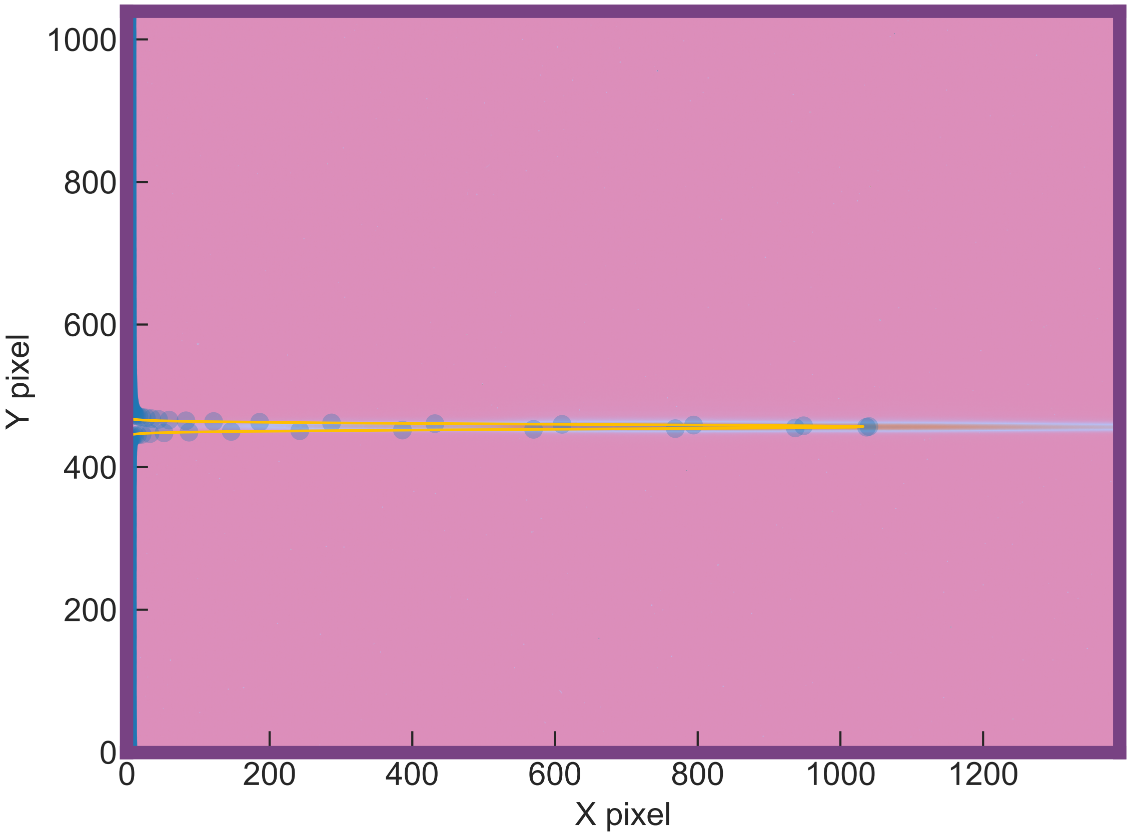

Figure 8.6: Science dispersion histogram. The blue circles represent a histogram of the data horizontally, with a sharp peak at the location of the science spectrum needing to be extracted. A Gaussian fit to this peak is plotted in gold, showing the automated estimation of the location of the dispersion axis.

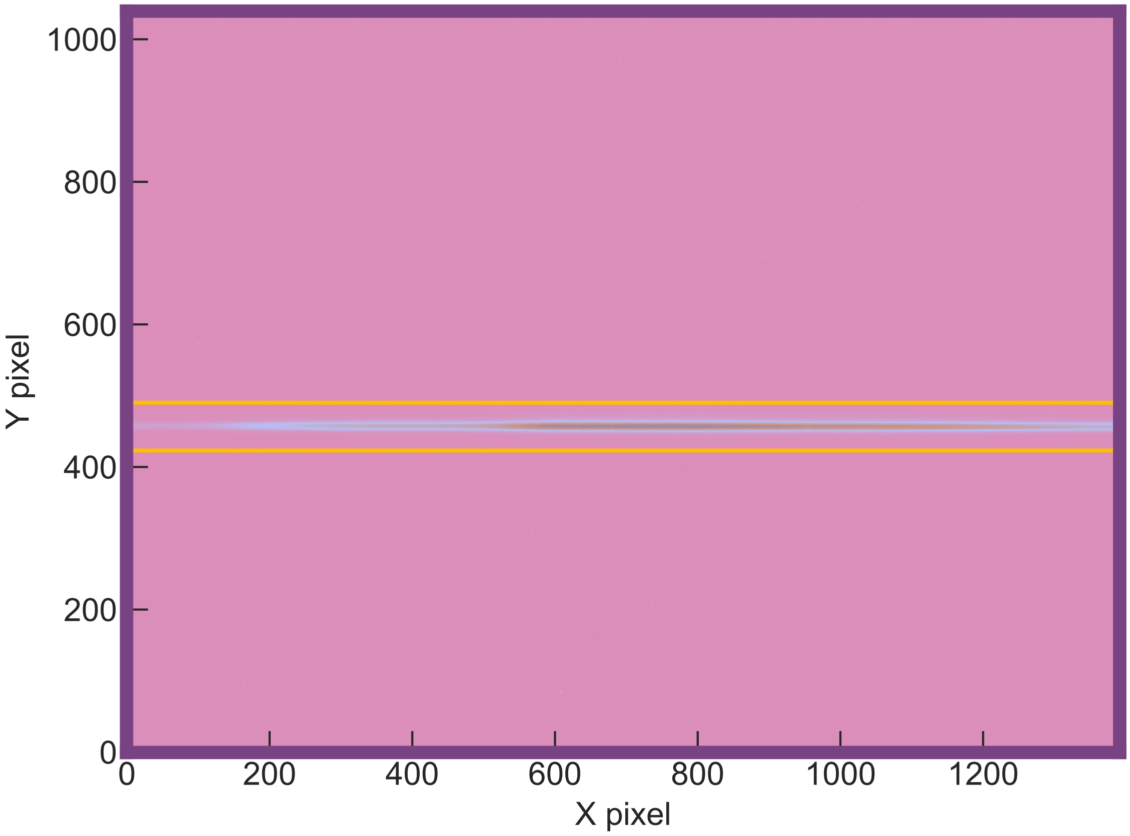

Figure 8.7: Science frame, with two horizontal gold markers enclosing the pixels which are to be used by this extraction.

8.4.3 Extraction strategy

The mean and standard deviation of the Gaussian returned by the dispersion step describe the vertical location and extent of the dispersion axis respectively. The two gold horizontal lines in Figure 8.7 mark the set of pixels to be used for the extraction of this particular frame. They are centred on the mean of the Gaussian depicted in Figure 8.6, and are separated by a (configurable) multiple of standard deviations. There are three main extraction methods currently in place for handling different science observations, but a benefit of the strategy pattern (Section 8.3) is that it is trivial to write a new controller with a different extraction strategy. All the strategies start by binning along the wavelength axis, and only differ in the method of extracting an intensity value for each bin.

8.4.3.1 Summation method

This method is simple - sum all pixels in the bin, and use this value for the intensity at that wavelength.

8.4.3.2 Trimmed mean method

Slightly more sophisticated than the summation method but still fast, the trimmed mean method first divides out the flat-field from the science data for the pixels in a bin. Then, it computes the mean of these ratios, after trimming the extremes. This limits the weighting given to saturated pixels and low-illumination (off-axis) pixels, as these will both increase the noise. The trimmed mean is then used as the intensity for the bin.

8.4.3.3 Gradient method

The gradient method (repurposed from the main ENDEAVOUR pipeline) takes the pair of science and de-trended flat-field values for each pixel in each bin, and performs robust linear regression between these two sets. The gradient of this linear relationship then represents the best estimate of the factor by which the science frame was brighter (or dimmer) than the flat frame. The gradient for each bin is then a measure of the relative brightness of that bin compared to others, and as such is used as the final value for intensity. While it is important to check that the pixel counts are in the linear regime of pixel response, this is less likely to be an issue for most targets compared with the main ENDEAVOUR pipeline processing of Aquila data, because the relatively higher resolution of this instrument ensures that the light is spread over many pixels.

8.4.4 Final output

The final extracted spectrum for the example used in this Chapter is presented in Figure 8.8, with tiny absorption lines due to the Earth's atmosphere clearly visible. Chapter 10 contains an example of this data stream being used in practice to supplement the Aquila campaign for YZ Reticuli.

8.5 Slit spectroscopy - a bonus

As a demonstration of the versatility of the new toy, I extracted a spectrum from various positions along the slit from a spatially-resolved spectrum of Jupiter, aligned

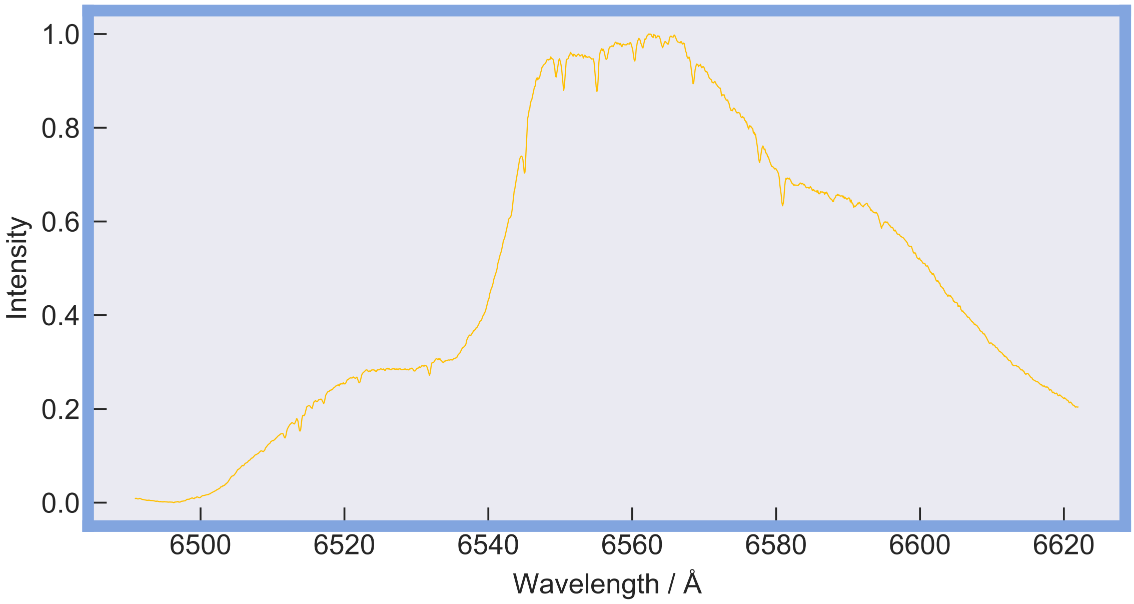

Figure 8.8: One-dimensional LHires extracted spectrum of classical nova YZ Reticuli, showing broad structure in Hα, but also very finely structured (telluric) absorption lines from the atmosphere of the Earth.

across its equator. The two-dimensional input image, shown in the upper panel of Figure 8.9, shows a deviation in wavelength from top to bottom of the slit for some lines of Jupiter, due to the differential speed caused by its rotation. It also has some lines which do not exhibit this slope, because they arise due to the Earth's atmosphere, rather than at Jupiter.

The lower panel of Figure 8.9 shows the set of spectra extracted from the bottom to the top of the slit, with those from the blue-shifted side of Jupiter (rotating towards Earth) shown in blue, and the red-shifted side of Jupiter (rotating away from Earth) shown in red. The telluric lines are easily identifiable as they show far less variation in wavelength for different heights along the slit axis, although they still show drift due to slight misalignment of the axis.

This instantaneous evidence for the rotation of Jupiter is, while not novel, quite fun; and totally orthogonal to the usual research operation of the Global Jet Watch observatories, but is also a useful educational tool for the host schools.

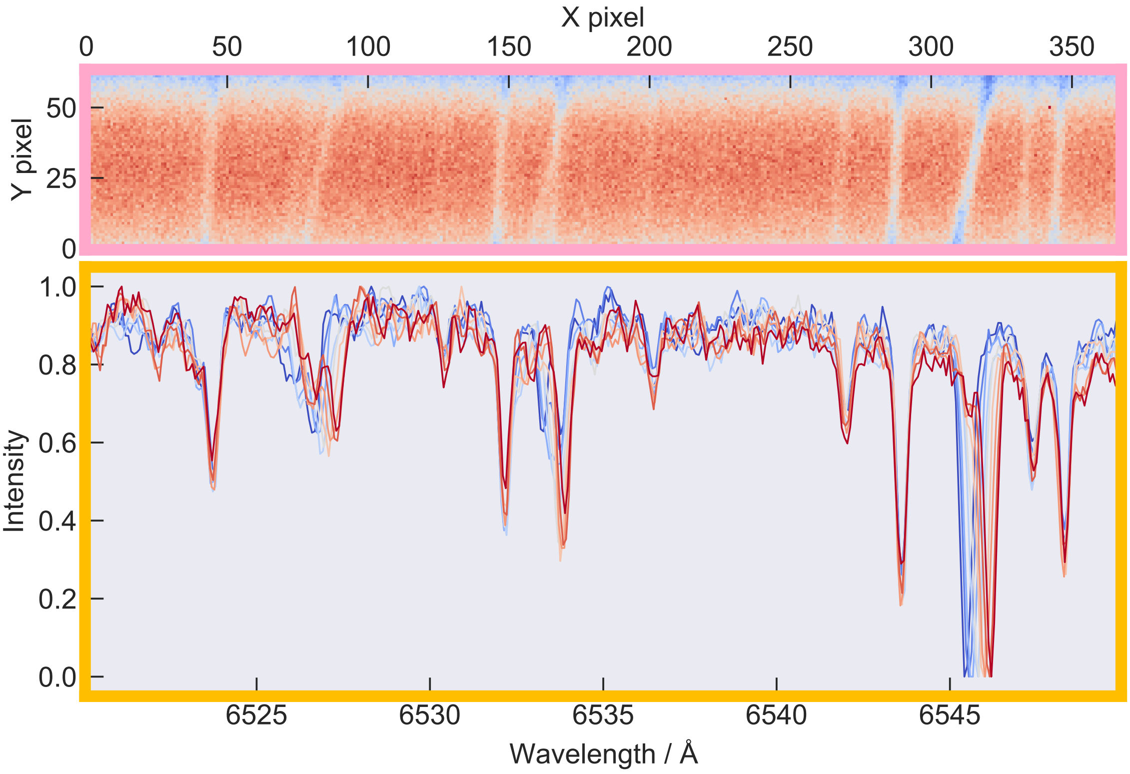

Figure 8.9: Spectrum of Jupiter from the LHires instrument of exposure time 30 s, showing sloped absorption lines since the relative Doppler shift of Jupiter's absorption lines at the top of the slit differ compared to those at the bottom. The same slope is not apparent for the Earth's telluric absorption lines, which remain fairly constant over the whole slit.

Chapter 9

Classical Nova V906 Car: Discovery of circumbinary iron and oxygen

The work presented in this Chapter has been published as McLoughlin et al. [2020].

Abstract

This chapter presents time-lapse spectroscopy of a classical nova explosion which commenced nine days after discovery. The data reveal the appearance of a transient feature in Fe II and [O I]. Different models are explored for this feature, which is best explained by a pre-existing circumbinary disc shock-heated following the classical nova event. This is significant because circumbinary discs may play an important role in novae in accounting for the transient heavy element absorption (THEA) systems, the transfer of angular momentum, and the possible triggering of the nova event itself.

9.1 Introduction

A classical nova is an explosion on the surface of a white dwarf which is in orbit in a close binary with a less-evolved companion. The companion is usually a main sequence star, or a red giant. Hydrogen gas accretes from the companion onto the surface of the white dwarf via an accretion disc (in the absence of strong magnetic fields), until the high temperatures and pressures at the bottom of the accreted layer on the white dwarf's surface pass a critical threshold and nuclear fusion commences in the hydrogen gas. This powers a very rapid thermonuclear runaway reaction, and the resultant explosion is thought to eject a shell of hot matter, on average of mass $2 \times 10^{-4}\,M_\odot$ [Gehrz et al., 1998].

Classical nova explosions are therefore rare opportunities to study what happens when a binary star system, consisting of a white dwarf and a donor star, undergo such a perturbation — dynamically, thermally and chemically. Spectroscopic study of the immediate post-explosive behaviour following a nova is a powerful tool to reveal, via the Doppler effect, dynamical changes, and via the appearance and subsequent disappearance of various emission lines, the changes that arise following nucleosynthesis and shock heating.

In March 2018, a classical nova detonated in Carina, a circumpolar location in the Southern Hemisphere sky very convenient for continual observation. It was discovered by the All Sky Automated Survey for SuperNovae[1] (hereafter ASASSN) and is named V906 Car, as well as Nova Carinae 2018, and ASASSN-18fv. The Global Jet Watch observatories followed this up spectroscopically with dense coverage (an average of 13.8 hr per nychthemeron) from days 9-20, with sparse coverage after that. Much rich behaviour was observed giving insights into different aspects of nova evolution in the first few weeks after explosion. This chapter focuses on remarkable transient spectral features associated with Fe II and [O I] lines.

As detailed in Arai et al. [2016], and described in Section 2.3, there has been much debate surrounding the origin of multiple absorption systems in classical novae. These systems were initially identified by McLaughlin [1950] as the 'principal' and 'diffuse-enhanced' absorption systems. Since then, Williams et al. [2008] has suggested another type of absorption component, the transient heavy element absorption (THEA). They suggested that a reservoir of circumbinary gas is the origin of the low radial velocity THEA system. Circumbinary discs around cataclysmic variables were theoretically investigated by Spruit and Taam [2001], who believe such

---

[1] http://www.astronomy.ohio-state.edu/∼assassin/

discs could be important in explaining angular momentum transfer in classical novae. However, this chapter presents a direct observation of a circumbinary disc surrounding a classical nova, which in McLoughlin et al. [2020] was the first in the literature to date, although Monnier et al. [2006] does hypothesise the existence of a circumbinary reservoir around the recurrent nova RS Ophiuchi.

Section 9.2 describes the round-the-clock spectroscopic coverage of the nova V906 Car. In Section 9.3, I confront these data with a circumbinary disc model. In Section 9.4, I describe the context and theoretical background for interpreting our observations. Section 9.5 discusses alternative possible models including a pinwheel geometry as reported by Tuthill et al. [2007]. Section 9.6 provides the conclusions.

9.2 Observations

The earliest detection of the nova V906 Car was at JD 2458193.8194 by the Evryscope-South observatory, according to Corbett et al. [2018], which I subsequently refer to as the discovery date for the system. It was independently discovered after this by Stanek et al. [2018] using AS-ASSN[2]. As published in McLoughlin et al. [2020], times are referenced as days since commencement of brightening (JD 2458193.8194). This will be important for the timing data, which is given in this format unless indicated otherwise. I make use of the unit mAU, for milli-Astronomical Units, as this is a convenient length scale to discuss the orbital separations and circumbinary radii of such a classical nova system.

9.2.1 Global Jet Watch observatories

Spectra of V906 Car spanning a wavelength range of approximately 5800 Å to 8400 Å and with a spectral resolution of $R \sim 4000$ were observed starting from +9.514d. These were carried out with the multi-longitude Global Jet Watch telescopes, each of which is equipped with an Aquila spectrograph; the design and testing of these high-throughput spectrographs are described by S. Lee et al. (in preparation).

A range of exposure times were taken throughout the observations of V906 Car since while unsaturated Balmer Hα had high signal-to-noise in the shorter exposure times, weaker signals in other parts of the wavelength range could be detected using the longer exposure times. The exposure times varied as the light curve changed, as

---

[2] The paper [McLoughlin et al., 2020] refers to this nova as ASASSN-18fv, its name from the AS-ASSN designation. In this thesis, names of the classical novae are standardised to use their variable star designation.

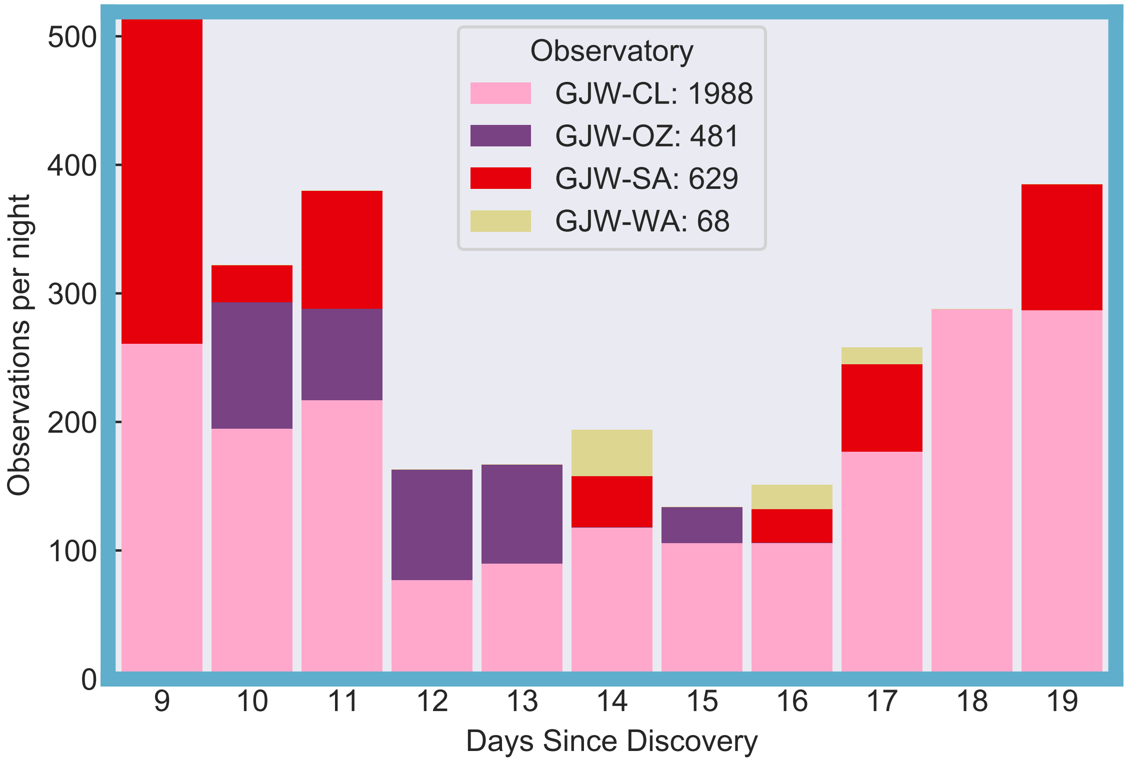

Figure 9.1: Stacked histogram showing the number of spectra taken each day per observatory during the early days of observation. CL — Chile, OZ — east Australia, SA — South Africa, WA — Western Australia.

Table 9.1: List of exposure times for observations of V906 Car, in decreasing order of number of observations, as of 2020-02-04. The exposure times were chosen to suit the instantaneous brightness of the target, for example some exposure times are short to avoid saturating the H-α line, while other exposure times allowed this line to saturate while getting good signal-to-noise on the other emission lines.

Exposure time (s)	Number of spectra
100	1370
300	1194
200	217
150	140
10	110
60	64
250	52
60	64
3	29
3000	6
1000	3

classical novae change luminosity by several orders of magnitude. A detailed breakdown of exposure times is provided in Table 9.1.

9.2.2 Epochs of observations

Figure 9.1 is a histogram of the number of spectroscopic observations captured, binned daily for the first eleven days of observations. The different colours represent the different observatories. What is notable about this data set is that it contains an unbroken run of over 100 spectra per night during the fastest evolving part of the light curve (made available by the American Association of Variable Star Observers, or AAVSO[3] and shown in Figure 9.3). Furthermore, the first three days have over 300 spectra each. There is a cumulative exposure time of 22.5 h for the first 24 h of the observation set. This is made possible because the telescopes are separated in longitude around the Earth. In total 3256 spectra were taken up to 2018 Sep 30, extending over a time span of six months since the detonation. Investigations of other spectral features and their evolution are presented in Chapter 11.

---

[3] aavso.org

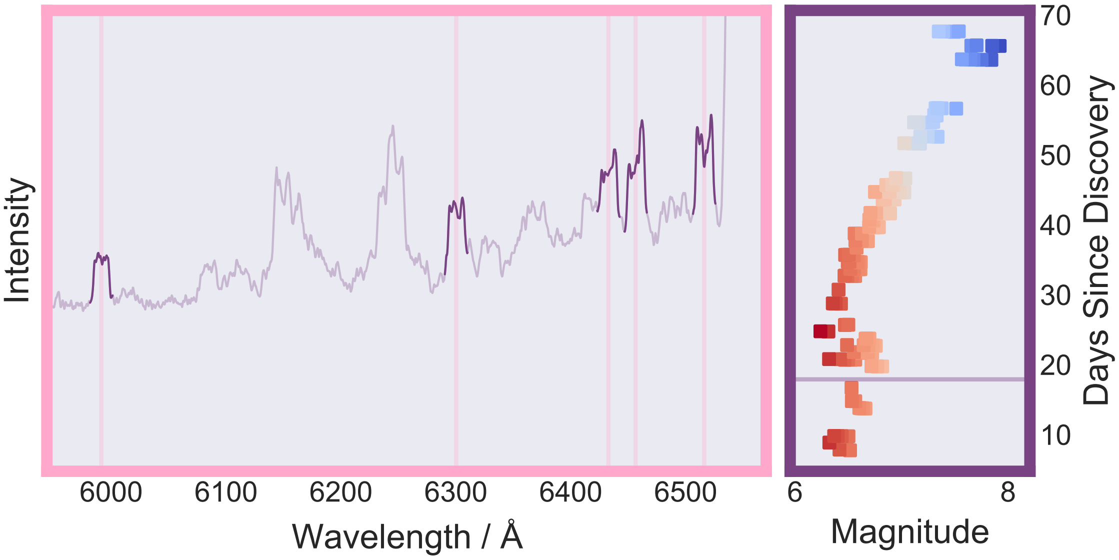

Figure 9.2: A single spectrum of exposure time 150 s observed at the GJW-CL observatory at +17.905 d, showing the characteristic signatures of the disc feature, at a number of different wavelengths (identified as four different Fe II transitions at 5991 Å, 6432 Å, 6456 Å, 6516 Å, and the forbidden [O I] 6300 Å transition) indicated in Fig 9.3.

9.2.3 Spectral features

Figure 9.3 presents a spectral feature, somewhat in the shape of the silhouette of a capital letter 'M' with sloping outer edges, noticed in several of the spectra while investigating other aspects of the evolving nova. The feature is most prominent from +13.5 d to +20 d, shown by the central column in Figure 9.3. Its origins can be traced as early as +13.1 d (top left panel), while hints of this structure remain until +50d (bottom right panel), but it does not constitute a detection at these extremes of times, as the signal-to-noise is too low. Each of five different colours represents an overlaid spectral feature corresponding to a different rest wavelength of an Fe II or [O I] line. There is no evidence for any spectral shape resembling this feature in either Hα or He. It is remarkable that the sloping outer boundaries of this feature overlay one another with such good agreement for each of the five lines. The modal velocities[4] correspond to the velocities of the outermost particles in a disc, per Elitzur et al. [2012]. This signal is consistent with the nova having very low ($< 50\,\mathrm{km\,s^{-1}}$) systemic velocity. This is corroborated by the pre-maximum spectra, which show

---

[4] The modal velocity is the Doppler-shifted velocity at which the spectrum has the most counts. This is used in preference over e.g. 'peak' velocity, because that would instead imply the highest speed.

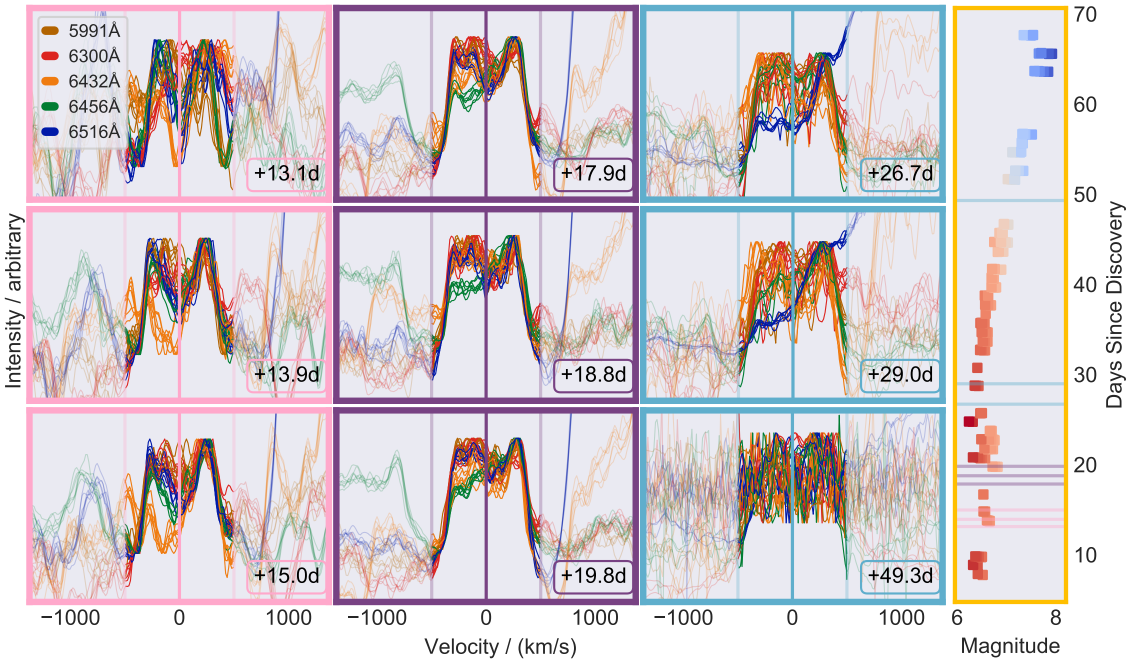

Figure 9.3: The central column shows a clear sloping M-shaped spectral feature in four lines of Fe II and one line of [O I]. Spectra taken at nine different epochs, shown in velocity space taking the rest wavelength of each line as zero radial velocity. The rightmost panel shows the AAVSO light curve, with time increasing up the page, and magnitude, hence faintness, increasing to the right. The observation times are shown superimposed on the light curve as horizontal green, peach and lilac lines. There is at best a marginal detection at +13.1 d — the Fe II and [O I] signatures emerge from +13.9 d, and last until +19.8 d. At each epoch, seven spectra taken within minutes of each other are superimposed to demonstrate that these are a persistent signal with common outer boundaries in velocity space, not just spectral noise.

narrow absorption lines with Doppler shifts roughly $\Delta\lambda < 1\,\text{Å}$ relative to the rest wavelengths, within the spectral resolution.

There is no significant change in the velocity width of the spectral features shown in Figure 9.3. I discuss the implications of this observation in Section 9.3.1 below.

9.2.4 Iron and Oxygen disc detection

Transient disc-like emission profiles appear at five different central wavelengths across our spectra, highlighted in bold in Figures 9.2 and 9.4. It is clear from Figure 9.3 that aligning the profiles on rest wavelengths of known Fe II emission lines (5991 Å, 6432 Å, 6456 Å, 6516 Å) and the 6300 Å [O I] line gives remarkable concordance in the velocity profiles at +17.8 d. The modal Doppler shift velocity, most naturally interpreted as the rotational velocity at the outer edge of a disc, is measured at $-266\,\mathrm{km\,s^{-1}}$ for the blue wing, and 254 km s^{-1} for the red wing. The M-shaped profile is unambiguously present from +13.9 d to +19.8 d, with marginal hints outside this time range.

9.2.5 Profile variations

For each epoch shown in Figure 9.3, six spectra taken within minutes of each other are superimposed, to demonstrate that this is a persistent signal, not simply spectral noise. While all these lines show the overall M-shaped profile, there is some variation between the profiles of different spectral lines within each epoch. I investigated whether telluric absorption lines had affected some of the Fe II and [O I] profiles, and established that while there is some overlap with such lines, they were too narrow and weak to significantly alter the profiles in our spectra.

There is a notable deviation from the clear M-shape seen in Figure 9.3 at early times in the 6432 Å line (yellow, left column), and at later times the 6456 Å line (green, middle column). For each of these two cases, the reduced emission relative to the other lines is limited to the blue-shifted peak. The discrepancy itself evolves with time, disappearing for some of the spectra and returning later in others, which suggests that a real physical evolution is unfolding.

9.3 Interpretation

9.3.1 Spectral shape

The double-peaked spectral features (showing such structure for many Fe II lines and for the forbidden [O I] line (very different transitions) with such similar boundaries in velocity) are naturally interpreted as rotating gas, as evinced by Fig 9.5. I fit this model to the Global Jet Watch data in Section 9.3.2. Other possible interpretations are discussed in detail in Section 9.5, although rotating gas best fits the data.

9.3.2 Disc model and fitting

Following Elitzur et al. [2012], I model the nova as having a kinematically double-peaked Keplerian ring. This simple model considers a flat disc viewed edge-on, with all particles on Keplerian orbits. Constant density is assumed throughout the ring for simplicity. The double-peaks arise naturally from the geometry; considering contours of constant line-of-sight velocity, the longest such contours correspond to the modal velocities; the peaks of the spectral signature. The contour of zero radial velocity is caused by material moving perpendicular to the line-of-sight. This (straight-line)

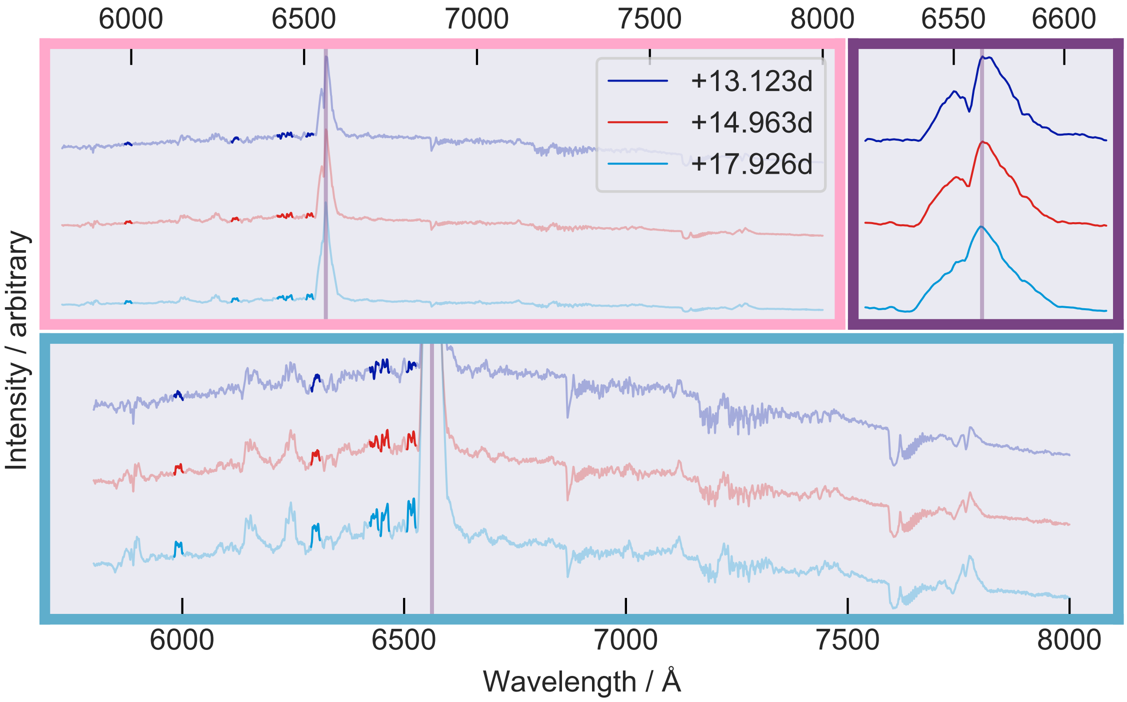

Figure 9.4: Three example spectra taken at +13.123 d (dark blue), +14.963 d (red), +17.962 d (light blue). The three panels clockwise from top left show the full spectra, a close-up of the Hα complex, and the full spectra cropped to show the spectral features other than Hα. The bold segments of these spectra are the same sections as those shown in bold in Figure 9.3.

Table 9.2: Fitting constraints for the uniform disc model using integration over disc area. Note that the mass and outer radius parameters are combined into one free parameter for the fit, since M/r is degenerate in this Keplerian model.

	Lower bound	Upper bound
Kepler Mass	$0.5\,M_\odot$	$3\ \ M_\odot$
Outer Radius.	$0.1\,$mAU	$10^4\,$AU
Radius ratio	0	1
$\mathrm{v}_{\mathrm{rot}}/\mathrm{v}_{\mathrm{thermal}}$	0	20

Table 9.3: Fitted parameters for the uniform disc model using integration over disc area. Since total enclosed mass and radius are degenerate in this model, the values are presented assuming the mass is $M_{\mathrm{tot}} = 1\,M_\odot$

Parameter	Fitted value
Inner Radius	$[0.017 \pm 0.001]\,$mAU
Outer Radius	$[14.7 \pm 0.1]\,$mAU
$\mathrm{v}_{\mathrm{rot}}/\mathrm{v}_{\mathrm{thermal}}$	5.1 ± 0.1
Outer Velocity	$[244 \pm 5]\,\mathrm{km\,s^{-1}}$

contour is shorter than those for higher absolute Doppler shifts (curved), explaining the dip between the peaks.

I fitted this model to the Global Jet Watch data as shown in Figure 9.5, using three free parameters: (i) total Keplerian central mass / radius of emissive material ($M_{\mathrm{tot}}/r_{\mathrm{disc}}$) (ii) radius ratio of disc ($r_{\mathrm{in}}/r_{\mathrm{out}}$) (iii) characteristic velocity ratio ($v_{\mathrm{rot}}/v_{\mathrm{thermal}}$). Table 9.2 details the upper and lower bounds used. The mass bounds were designed to incorporate the full range of what could be expected for this system. Considering that the likely mass range for the white dwarf in a classical nova system is at least restricted within $0.6\,M_\odot$ to $M_{\mathrm{CH}} = 1.4\,M_\odot$ (the Chandrasekhar limit)[Shara et al., 2018], the combined mass of the two stars is most likely contained within the interval $0.5\,M_\odot < M_{\mathrm{tot}} < 3\,M_\odot$. The ratio of inner/outer radius of the orbiting material was allowed to vary from zero (a disc), to unity (a vanishingly thin ring). The relative velocity scale was allowed to range from 0 to 20, as per Elitzur et al. [2012].

The results of the fitting are in Table 9.3. I utilised an MCMC fitting procedure based on the `emcee` python package [Foreman-Mackey et al., 2013]. Since the mass and outer radius are degenerate, the system parameters derived from the combined mass/outer radius free parameter are listed for the case of the inner binary having $M_{\mathrm{tot}} = M_\odot$.

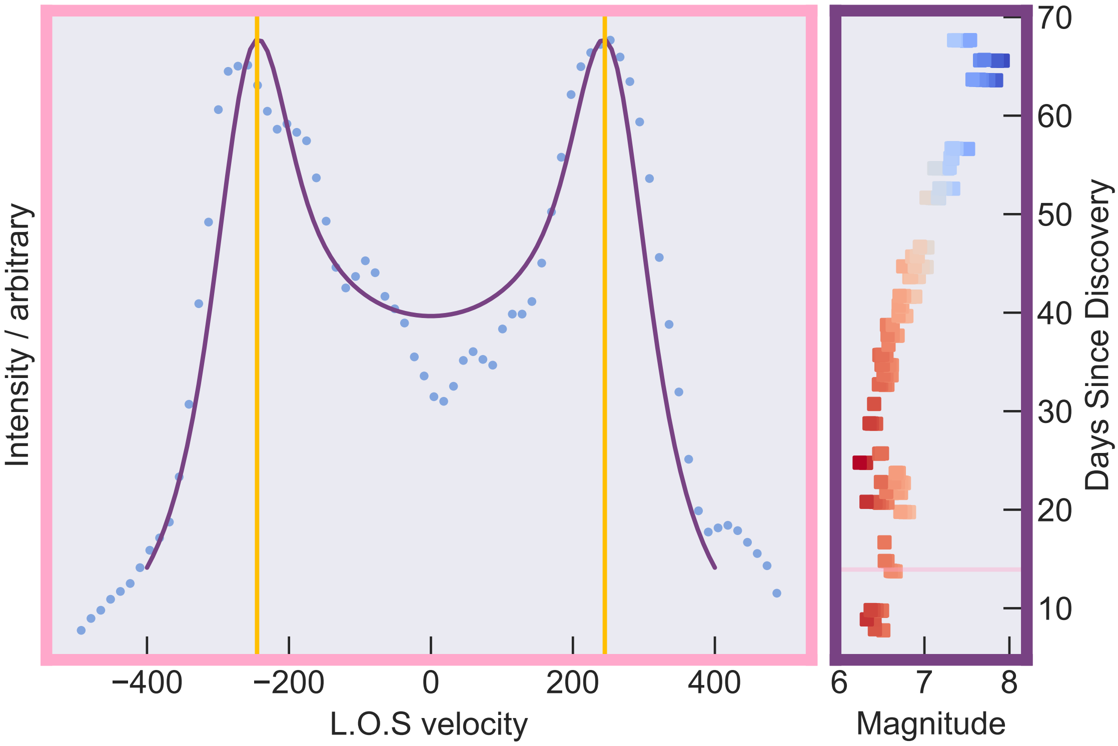

Figure 9.5: The bins are 0.3 Å wide, or 14 km s^{-1} at this wavelength. The left panel shows the best fit profile for a rotating Keplerian disc with uniform surface density, viewed edge-on, as a function of line-of-sight velocity. Data are taken from the 6516 Å Fe II line signal, taken at +13.928 d at the GJW-CL observatory in Chile, indicated on the light curve with a horizontal line. The vertical lines indicate the velocity of the fit at the outer edge of the disc.

9.4 Circumbinary discs

9.4.1 Theoretical predictions

Covered in greater detail in Section 2.4, the n-body, test particle, 3D simulations of Doolin and Blundell [2011] showed that there are three families of remarkably stable circumbinary orbits around binary star systems throughout eccentricity/mass ratio parameter space.

The distance from the binary centre of mass to the innermost stable radius for test particle orbits has been studied numerically and analytically [Szebehely, 1980, Holman and Wiegert, 1999]. It is given as $r_c \approx f(q)a_B = f(q)r(1+q)$ where a_B is the binary separation, q is the mass ratio of the two stars m/M, and r is the orbital radius of the less massive component. $f(0.1) = 2$ and $f(0.3) \approx 2.3$.

9.4.2 Significance of timing

As shown in Figure 9.3, the M-shaped signal becomes clear by +13.9 d. The signal then fades into noise after a week, with very little trace of the detailed shape remaining after +19.8 d. Taking the speed of sound in the intervening medium between the explosion and a hypothetical circumbinary disc to be $1\,\mathrm{km\,s^{-1}}$, and the circumbinary radius to be two times bigger than the semimajor axis and approximately 10 mAU, the time for a shockwave to reach a circumbinary disc would be around 17 d. This is the same ballpark time that it takes for the signal in these observations to appear.

What might determine the switch-off time? The signal appears to melt away ~10 days after appearing, which could be due to cooling of the shock-heated circumbinary disc. Under this interpretation, a shockwave travels radially outwards from the initial event for 14 days covering the distance to the inner edge of the circumbinary disc (estimated to be at ~ 10 mAU), then arrives at the disc and shock-heats it, rapidly causing emission in a profile that informs us of the underlying pre-existing circumbinary disc geometry at +13.9 d. This recently heated disc is visible clearly for a week, but loses energy through radiation, and cools, which we observe as the disc signal growing fainter and fainter over the following week (from +19.8 d to +27 d). Crucially, in any alternative model involving a sustained wind or ejection, it isn't trivial to conjure up a reason for the ejection to turn off in a week-long decay.

There is also an observed change in dynamics over the week of strong signal. Close inspection of the spectrum at +17.8 d in Figure 9.3 reveals that the left peak is in fact dynamically split at this epoch. The 6516 Å and 6456 Å lines at +13.9 d show no sign of such splitting, but this develops in the next four days. Under the shock-heated circumbinary rotating disc interpretation, it is reasonable that the shock disrupts the equilibrium and introduces dynamical changes as it spreads outwards through the gas.

9.4.3 Significance of chemistry

The observations of the disc signature are identified as being Fe II and [O I] ions, and this structure is not observed in H or He . Whatever the mechanism for establishing a circumbinary body of gas, if the gas were formed of recent ejecta then we would expect to see the M-shaped profile in hydrogen and helium spectral lines (Hα 6563 Å, He 7065 Å, He 5876 Å). I therefore speculate that the signal is dominated by relatively heavy species because lighter particles do not remain in these orbits for long times. It is possible that, much like the Jeans escape mechanism for the atmospheres of planets, the lighter elements pick up enough kinetic energy given the temperature (via the equipartition theorem) to cease to be bound in the circumbinary disc. On the other hand, for the heavier elements, particles have less speed for the same temperature, and do not reach the required escape velocity. The notable absence of M-shaped profiles in our spectra from hydrogen or helium lines corroborates the conjecture that this circumbinary system pre-existed the recent explosion.

9.4.4 Circumbinary discs in classical novae

A longstanding problem in this field is the so-called “white dwarf mass problem", detailed in Zorotovic and Schreiber [2019]. The fundamental issue is a discrepancy whereby the theoretical model predictions of white dwarf masses consistently underestimate the mean observed value in well-determined spectroscopic binaries of $M_{WD} = 0.83\, M_{\odot}$. Consequential angular momentum loss, or angular momentum loss additional to gravitational radiation and magnetic braking, is necessary to resolve the discrepancy [Schreiber et al., 2016, Liu and Li, 2018]. One such mode of angular momentum transfer from the inner binary could be provided by a circumbinary disc.

Spruit and Taam [2001] sets out the theory behind circumbinary discs surrounding cataclysmic variables. They argue that cataclysmic variables have a high variance in the mass-transfer rate $\dot{M}$ given the small variation in white dwarf masses. A circumbinary disc applies a tidal torque to the secondary star undergoing Roche lobe overflow, increasing both angular momentum transfer and $\dot{M}$. They show that if only a small fraction from 10^{-4} to 10^{-3} of the total $\dot{M}$ lost from the secondary is trapped in a circumbinary disc, this disc will become massive enough on short enough timescales to influence the mass transfer. Furthermore, as this is a positive feedback mechanism, it can continue to accelerate mass transfer, which could act as the “trigger” for a cataclysmic event such as a classical nova. It is interesting in this regard that a circumbinary disc whose existence is posited on the basis of Fig 9.5 could explain the mass problem and could act as a trigger for the nova eruption.

9.4.5 Kepler argument for disc being circumbinary

This section aims to compare the detected rotation speed of the disc signal with an estimate of the rotational speed of the secondary star. If the disc signal is orbiting at a slower speed than the secondary then we can conclude that the disc is orbiting at a larger radius and is therefore external to the inner binary, given the total central mass for these (assumed Keplerian) orbits is the same in both cases.

Consider the two-body problem in a circularised Keplerian system with a white dwarf, $M_{\rm WD}$, and a lower mass companion, $M_2 < M_{\rm WD}$. Defining the mass-ratio by $q = M_2/M_{\rm WD}$, the speed of the secondary is bounded by

$$v_{2,\max} = \frac{2\pi}{(1+q_{\min})} \cdot \left(\frac{G}{4\pi^2} \frac{M_{\max}}{P_{\min}} \right)^{\frac{1}{3}} \tag{9.1}$$

$$v_{2,\min} = \frac{2\pi}{(1+q_{\max})} \cdot \left(\frac{G}{4\pi^2} \frac{M_{\min}}{P_{\max}} \right)^{\frac{1}{3}} . \tag{9.2}$$

The total binary mass then satifies $M_\odot < M_{\rm tot} < 2M_{\rm Ch}$, since the companion cannot in this paradigm be more massive than the WD, and the WD itself cannot be more massive than the Chandrasekhar limit $M_{\rm Ch}$. I take the orbital period $81\,\mathrm{min} < P < 6\,\mathrm{h}$, where the limits come from population studies of novae [Kolb and Baraffe, 1999, Fuentes-Morales et al., 2021]. Setting the mass ratio lower limit at $q \geq 0.1$ gives $v_{2,\min} = 343\,\mathrm{km\,s^{-1}}$, which is comfortably larger than the observed $260\,\mathrm{km\,s^{-1}}$ rotation speed of the disc in V906 Car, for inclination angles $\theta < 60°$. The upper bound is given by $v_{2,\max} = 783\,\mathrm{km\,s^{-1}}$, which is significantly larger than the detected signal. Assuming moderately low inclination angles, this is evidence that the detected disc is at a larger orbital radius from the inner binary centre of mass than the secondary is — in other words, it is a circumbinary disc.

The profile of the first strong detection of this feature in the 6516 Å Fe II line gave the outer radius as $r_{\rm out} = 14.7\,\mathrm{mAU}$. Assuming a Keplerian inner binary, with canonical total mass $M = 1.5\,M_\odot$, and a typical orbital period from Townsley and Bildsten [2005] for a classical nova of $P = 4\,\mathrm{hr}$, gives a semi-major axis $a = 5.9\,\mathrm{mAU}$

for the inner binary, in the absence of any certain evidence of eccentricity. This puts the outer radius of $r_{\rm out} = 14.7\,{\rm mAU}$ outside the binary. Crucially, $r_{\rm out} > 2 \times a$, which puts it in the regime of stable circumbinary orbits [Doolin and Blundell, 2011].

9.4.6 Low-density hot gas

The same disc spectral signature detected in the Fe II lines is detected in the 6300 Å [O I] line. This is evidence that its origin is in a low-density region. It is also hot (with a temperature in excess of 10^5 K), as evinced by the appearance of ionised Fe II which has an ionisation potential of 16.2 eV. This is naturally explained in a model in which it is a diffuse and expansive circumbinary disc which has just been shock-heated by a blastwave emanating from the recent explosion, causing it to emit.

9.5 Possible alternative models

There are *a priori* four mechanisms for rotation in a binary system: (i) the underlying binary orbit, (ii) spin of one of the stellar components, (iii) rotation of gas within the accretion disc and (iv) rotation of gas in a family of circumbinary orbits far outside of the binary. The observed spectral system appears too late after the brightening to be readily explained by any of (i) to (iii). If the spectral signature were evidence of an accretion disc, it should have been seen within the first day after the first flaring episode, and the same argument applies to (ii) and (iii). The speeds are also lower than what would be expected if they derived from inner binary activity — as detailed in Section 9.4.5.

9.5.1 Bipolar outflow model

If this were a bipolar outflow, it should be apparent in Hα. It would also be expected to be on the order of $\sim 1000\,{\rm km\,s^{-1}}$ (not a few hundred) as per Harvey et al. [2018] and McLoughlin et al. [2021a]. Also, if it were a bipolar outflow into a decreasing density profile, you would expect it to widen in velocity space. This is an oddly transient feature, appearing at +13.9 d and lasting only until +19.8 d — although long-lasting bipolar phenomena such as those described in Chapters 10 and 11 are common in classical novae, transient bipolar outflows have not been observed previously and there is no reason to invoke them here.

9.5.2 Pinwheel model

I now consider the merits of a pinwheel model compared to the circumbinary disc model. In this model, instead of a reservoir of orbiting circumbinary material, some of the secondary star's wind escapes in a powerful stream of gas launched radially outward from the primary, funneled through the unstable L2 Lagrange point. As the inner binary rotates, so does this outward beam, resulting in a spiral/pinwheel shape [Tuthill et al., 2007]. I present a toy model in which the system launches blobs of matter in a pinwheel, and the intensity contribution of each expanding blob decreases. Simulated spectra generated via this model show the characteristic double peaks that we observe in V906 Car.

Let us first consider the timing data, per Section 9.4.2. The M-shaped signal in these observations commences abruptly at +13.9 d, lasting with clear signal for approximately one week before fading into the low signal-to-noise regime. This is naturally explained by the pre-existing circumbinary disc model, as explained in Section 9.4.2. In the pinwheel model however, there is no natural geometric reason why the gas blown off the secondary would only light up two weeks after the initial peak of luminosity. Moreover, it is not clear why, having been established, the pinwheel would turn off after seven days.

The second piece of evidence is the chemistry, detailed in Section 9.4.3. These observations show M-shaped profiles clearly for Fe II and [O I], but not for either Hα or He. For the circumbinary disc model, since we are seeing a stable orbiting structure, formed prior to the most recent outburst event, there has been sufficient time for the lighter elements to have escaped. This is similar to the Jeans escape mechanism whereby the kinetic energy of lighter gases causes them to be unbound, and they leave a planet's atmosphere. The pinwheel model allows no time (or reason) to lose the inevitable H and He, in disagreement with our observations.

The third piece of evidence is the spectral shape itself. In Figure 9.5, the drop-off edge at speeds higher than the peak intensity is slanting, rather than being a vertical immediate drop-off. The disc model predicts this naturally, since there is matter closer to the centre which orbits faster than the modal velocity. Under the pinwheel model, fluid parcels are launched with a sinusoidally varying line-of-sight velocity, such that any bolus launched parallel to our line of sight has the maximum perceived speed. To fit the observations at higher speeds than the peak speed, one would need to invoke that the pinwheel engine is launching matter at a variety of speeds. This assumption is not required by the circumbinary disc model.

There is spectral profile asymmetry in some of the spectral shapes shown in Figure 9.3 which is not explained by a purely axisymmetric circumbinary disc. However, the simulations of Doolin and Blundell [2011] contend that the stability of many circumbinary orbits comes at the expense of those orbits precessing (this is especially the case for inner binaries having mass ratios that differ from unity or that have significant eccentricity). Precession would act to tilt parts of the circumbinary disc in and out of the direct line-of-sight, readily explaining the asymmetric spectral profiles.

Occam's razor favours the circumbinary disc. Circumbinary orbits have been shown to be very stable (under certain conditions, satisfied by V906 Car) by Doolin and Blundell [2011], and what is stable and possible must be expected to occur in nature. There are many examples of circumbinary orbits, including Be star HR 2142 [Peters et al., 2016], SS433 [Blundell et al., 2008, Perez and Blundell, 2010], and 99 Herculis [Kennedy et al., 2012]. While there exist elegant manifestations of pinwheel features in nature, for example in a few Wolf-Rayet stars [Tuthill et al., 2007], a pinwheel model is neither warranted by our data, nor consistent with it.

9.5.3 Outflowing disc model

Could the M-shaped profiles be caused by an outflowing disc as opposed to a rotating disc? I argue that the fundamental spectral shape, as discussed for the pinwheel model in Section 9.5.2, cannot match the observations. Demonstrated in Section 9.5.4, the edge ascent is vertical for an outflowing disc model, while only the rotating disc model can account for the steady decrease in flux at higher radial speeds than the modal speed. Despite this, it is possible that the rotating disc is also expanding radially outwards with a similar speed to the rotation. The shock arriving at the circumbinary disc heats it up, and it is possible that it also accelerates it. Any expanding disc is nonetheless likely to be circumbinary pre-existing material because of the timing reasons detailed in Section 9.4.2.

9.5.4 Spectral comparison of models

I have presented four models to explain this spectral feature. The bipolar outflow model is ruled out because it predicts much greater speeds than observed, together with its manifestation in H and He. Figure 9.6 shows a comparison between the spectral shape determined through the remaining three models — rotating disc, pinwheel and expanding disc. Due to geometric considerations, only the rotating disc model predicts a sloping drop-off beyond the peaks. Conversely, the pinwheel and expanding

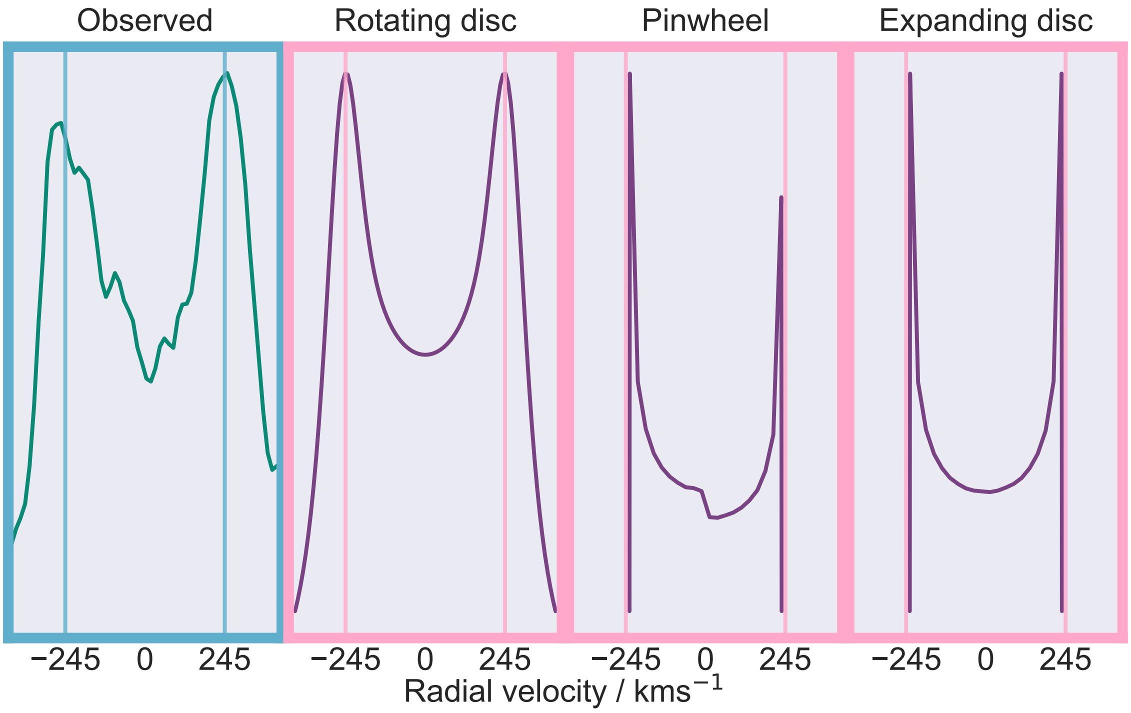

Figure 9.6: Data, left, compared to three models - rotating disc, pinwheel, and expanding disc. Spectral feature from the 6516 Å Fe II line signal, taken at +13.928d at the GJW-CL observatory.

disc models are characterised by a sharp cut-off at higher velocities. Convolving even with 45 km s^{-1} is insufficient to obtain qualitative agreement. Even if only lower-resolution observations were available, the pinwheel and expanding disc models could not satisfactorily explain the observed switch-on and switch-off timescales. In addition, the rotating disc best reproduces the width of the peaks and is therefore the favoured model.

9.6 Conclusions

I have presented observations of Fe II and [O I] emission features in time-resolved spectroscopy of V906 Car. Its consistency with being a circumbinary disc was demonstrated through fitting a model to the spectral shape. The possibility that the signature is in fact related to two jets was rejected due to the non-detection of this feature in Hα and He . The rotating circumbinary disc model fits the data better than a pinwheel scenario and two other models. I have presented three separate strands of evidence — timing data, chemistry, and spectral shape. I provided a simple orbital mechanics argument which gives a strong indication that this disc is in fact circumbinary. The theoretical reasons why circumbinary discs may have a large role to play in explaining the white dwarf mass problem in close binary evolution for the classical nova scenario were noted. Hence I conclude that the M-shaped profile signal rotation most likely arises from circumbinary gas.

To establish the prevalence of circumbinary discs in novae is challenging, because the signal is even shorter-lived than the nova event itself. Both nova evolution studies in general, and investigations of their putative circumbinary discs in particular, necessitate sustained highly time-resolved spectroscopic observations.

Chapter 10

The precessing jets of classical nova YZ Reticuli

The work presented in this Chapter has been published as McLoughlin et al. [2021a].

Abstract

The classical nova YZ Reticuli was discovered in July 2020. Shortly after this the Global Jet Watch observatories commenced a sustained, highly time-sampled coverage of its subsequent rapid evolution with time-resolved spectroscopy. Its H-alpha complex exhibited qualitatively different spectral signatures in the following weeks and months. These H-alpha complexes are found to be well described by the same five Gaussian emission components throughout the six months following eruption. These five components appear to constitute two pairs of lines, from jet outflows and an accretion disc, together with an additional central component. The correlated, symmetric patterns that these jet/accretion disc pairs exhibit suggest precession, probably in response to the large perturbation caused by the nova eruption. The jet and accretion disc signatures persist from the first ten days after brightening – evidence that the accretion disc survived the disruption. I compare another classical nova (V6568 Sgr) that erupted in July 2020 whose H-alpha complex can be described analogously, but with faster line-of-sight jet speeds exceeding 4000 km/s, and suggest that classical novae with higher mass white dwarfs bridge the gap between recurrent novae and classical novae such as YZ Reticuli.

10.1 Introduction

When classical novae erupt, their optical luminosities increase dramatically and rapid spectral changes are observed which are exemplified in their evolving Balmer Hα complexes. Changes occur on rapid and varied timescales ranging from hours to weeks. One such classical nova, YZ Reticuli, underwent an outburst in July 2020 and the Global Jet Watch observatories implemented a sustained, time-resolved, spectroscopic follow-up programme, monitoring it extensively throughout the first few months after it was discovered. This revealed properties of key dynamical components that drive its observed evolution. YZ Reticuli constitutes a particularly convenient example of a very energetic classical nova, having broad emission lines (FWHM $\sim 3000\,\mathrm{km\,s^{-1}}$). The high speeds in this classical nova make it possible to discern distinctly resolved components: I reduce the dimensionality of the data by fitting composite models of multiple Gaussians to the Hα profile. This reveals that much of the dramatic spectral behaviour can be explained by simple changes to these underlying components giving insights into the underlying dynamical story.

Section 10.2 briefly recaps the instrumentation and describes the observations this chapter is based upon. The methodology is presented in Section 10.3, and the results of the fits in Section 10.4. Section 10.5 suggests a physical interpretation for YZ Reticuli comprising a prominent and precessing accretion disc, and jets, in the aftermath of the nova event. Section 10.6 then confronts this model for YZ Reticuli with data from another classical nova, V6568 Sgr, which also erupted in July 2020. The conclusions are presented in Section 10.7.

10.1.1 Ejecta of classical novae

It is widely agreed that many classical nova eruptions eject roughly spherical shells, which move at slow speeds of up to ($\sim 1000\,\mathrm{km\,s^{-1}}$). The evidence for this includes remnant shells of past novae [Krautter et al., 2002, Santamaría et al., 2020], and the near-universally observed absorption lines in the early spectra around maximum and immediately following it. These absorption lines often show slow and fast components, known as the principal and diffuse systems respectively [McLaughlin, 1942, Williams, 1992], and are linked to a slow, equatorial outflow and a faster polar outflow respectively. In some novae, these flows are thought to collide [Williams and Mason, 2010], causing shocks. Recent multi-wavelength studies have shown correlations between gamma-ray detections interpreted as shocks between the various ejecta, and changes in the light curve [Aydi et al., 2020a]. Slavin et al. [1995] present evidence that in addition to ellipsoidal shells, long tails of emission can appear. A popular idea is that in addition to an initially approximately spherical slowly-ejected shell, there is a significantly faster outflow along the polar axes which can disrupt the shell; this picture is clearly supported in the case of DQ Her in figure 3 of Toalá et al. [2020].

10.1.2 Jets and accretion discs in novae

The launch of jets in astrophysical systems undergoing accretion is widely accepted. However, typical cataclysmic variable systems have been thought to be accreting at too low a rate to launch jets according to established disc theory. Classical novae could therefore present an important window on the emergence and survival of the jet phenomenon, because they have recently undergone disruptive eruption and associated enhanced accretion, which under favourable circumstances may put these systems over the threshold for the jet-launching mechanism to be triggered [Lasota and Soker, 2005].

There have been theoretical predictions and suggested observations of jets in classical novae: Retter [2004] discusses the theoretical viability of jets, using the V1494 Aql eruption in 1999 [Iijima and Esenoglu, 2003] as an example. Csák et al. [2005] find that there is evidence in favour of a link between the transition phase in novae and the establishing of an accretion disc. Kawabata et al. [2006] show jet-like winds in V475 Scuti, a moderately fast but otherwise unremarkable classical nova that erupted in 2003. They used Okayama Astrophysical Observatory spectropolarimetry to show that the nova had blue- and red-shifted winds with a constant position angle of linear polarisation in their polarised flux spectrum.

In V5668 Sgr, Harvey et al. [2018] presents a jet-like biconical outflow as an explanation for the profiles of [O III], although their model-dependent fits do not capture the asymmetry fully (see their figures 7 & 8). Ribeiro et al. [2013] present modelling of a fast bipolar outflow based on their fits to emission line structures. Chapter 11 contains results from the following paper in this series (McLoughlin et al, submitted), which presents V5668 Sgr in the context of the model set out in this chapter.

There is also evidence of jets in recurrent novae. Sokoloski et al. [2008] present direct evidence of highly collimated outflows and lobes in radio images of recurrent nova RS Oph during its 2006 eruption, captured by the VLBA with milli-arcsecond resolution. These lobes were found to be moving across the sky with proper motions implying underlying jet launch speeds of a few thousand $\mathrm{km\,s^{-1}}$. Figure 5 of Taylor et al. [1989] shows a radio map of the object, with clearly aspherical lobes apparent within the first few months of the 1985 outburst. Rupen et al. [2008] observed synchrotron jets with the VLBA, noting that this implies emission is occurring far from the central source. This has important consequences, as it implies that fast jets may extend beyond the inner system long before the photosphere fully recedes. Skopal [2015] presents spectra of the Hα complex of recurrent nova RS Oph, and their figure 3 shows these are qualitatively similar to what is seen here in the classical nova YZ Reticuli, although they do not present any fitting analysis or decomposition of the complex.

Darnley et al. [2017] discuss the jets in the recurrent nova M31N 2008-12a, which contains a white dwarf near the Chandrasekhar limit, and thus accretes at a high rate, hence enabling the very rapidly (yearly or more frequently) recurring eruptions. They found broad shoulders around Hα (offset from line centre by about $4800\,\mathrm{km\,s^{-1}}$ and about $5900\,\mathrm{km\,s^{-1}}$) and attributed these to fast bipolar outflow — intriguingly faster than the jets that might be expected from a lower mass white dwarf such as those in classical nova systems.

V1721 Aquilae, a very fast classical nova (according the the Payne-Gaposchkin scheme set out in Section 2.3), within three days of discovery was found to host an accretion disc that was either re-established fast, or was indeed never fully disrupted [Hounsell et al., 2011]. It seems probable then, that there exist generalised jet-like outflows in erupting nova systems, with faster and more collimated jets occurring in recurrent systems containing white dwarfs that are more massive, through very fast novae such as V1721 Aquilae with moderate jets, to more typical classical novae, such as the one for which evidence is presented in this chapter.

10.1.3 Classical nova YZ Reticuli

YZ Reticuli was discovered by Robert H. McNaught (Coonabarabran, NSW, Australia) at magnitude 5.3 on 2020 July 15.590 UT (CBET 4811). It is also identified as Nova Reticuli 2020, and MGAB-V207, with Right Ascension of 03 58 29.55 and Declination −54 46 41.2 (J2000). There is a light curve from the All Sky Automated Survey for SuperNovae (hereafter ASAS-SN) for this object which first shows evidence of brightening at 2020-07-08.17081, or JD 2459038.67081, and this is the epoch taken as the zero point. For the rest of this chapter, I use the convention that +10 d should be interpreted as 10 days after this zero point. There is a four-day gap in the ASAS-SN observations (magnitude 15 at 2020-07-04.18, magnitude 6 at 2020-07-08.17), so it cannot be precluded that it could have started brightening as early as 2020-07-04.18 [Aydi, 2020a].

Figure 10.1 shows the AAVSO[1] light curve of YZ Reticuli for the 125 days following eruption. The decline time of this classical nova is $t_2 = 15$ days. YZ Reticuli exhibits a plateau-type light curve [Strope et al., 2010] between approximately +30 d and +60 d, which may relate to the plateau in the light curve of classical nova V407 Cyg, interpreted by Hachisu and Kato [2012] as a surviving accretion disc emerging out of the receding photosphere.

YZ Reticuli was detected in gamma-rays between 2020-07-10 and 2020-07-15 by Fermi-LAT [Kwan-Lok, 2020]. [Sokolovsky, 2020] detected it in X-rays using NuSTAR, and found evidence for non-solar abundances of N , O and/or Fe , in 67 ks of total exposure commencing 2020-07-17.98. A light curve is available from ASAS-SN[2] for the 2000 days preceding eruption. The quiescent brightness was around 16 magnitudes. I did not detect periodicities in the nova's quiescent luminosity in that data set despite using the CLEAN algorithm together with deconvolution following Roberts et al. [1987].

Songpeng [2020a] notes that the X-ray luminosity of YZ Reticuli is significantly lower than expected, which hints at the existence of a spatially extended accretion disc which may be obscuring a significant portion of the hot surface of the white dwarf. Coronal line emission is reported by Songpeng [2020b], who also comment on the absence of dust and on the appearance of double-peaked profiles with deep troughs at the centre that are explored in Section 10.5.3. This X-ray phase, deemed by Sokolovsky et al. [2020a] to be the super-soft phase, commenced at +58 d, indicating the epoch by which the photosphere has receded enough to see hydrogen burning on the surface of the white dwarf.

This classical nova is in the GAIA DR2 catalogue [Prusti et al., 2016, Brown et al., 2018] (source identification number: #4731746232846281344), with a recorded parallax of [0.3781 ± 0.038] mas, giving a distance of $2.7\,^{+0.4}_{-0.3}$ kpc [Bailer-Jones et al., 2018]. As it was not previously classified as a variable star, it does not have a GAIA light curve.

10.1.4 Classical nova V6568 Sgr

Classical nova V6568 Sgr, also known as Nova Sagittarii 2020 #3, or PNV J17580848-300537, was discovered at magnitude 9.9 by Shigehisa Fujikawa on 2020 July 16.51939 UT (JD 2459047.01939). V6568 Sgr has Right Ascension 17 58 08.48 and Declination

---

[1] aavso.org

[2] https://asas-sn.osu.edu/light_curves/ee130388-139d-4f3a-ab8a-7427a7545e21

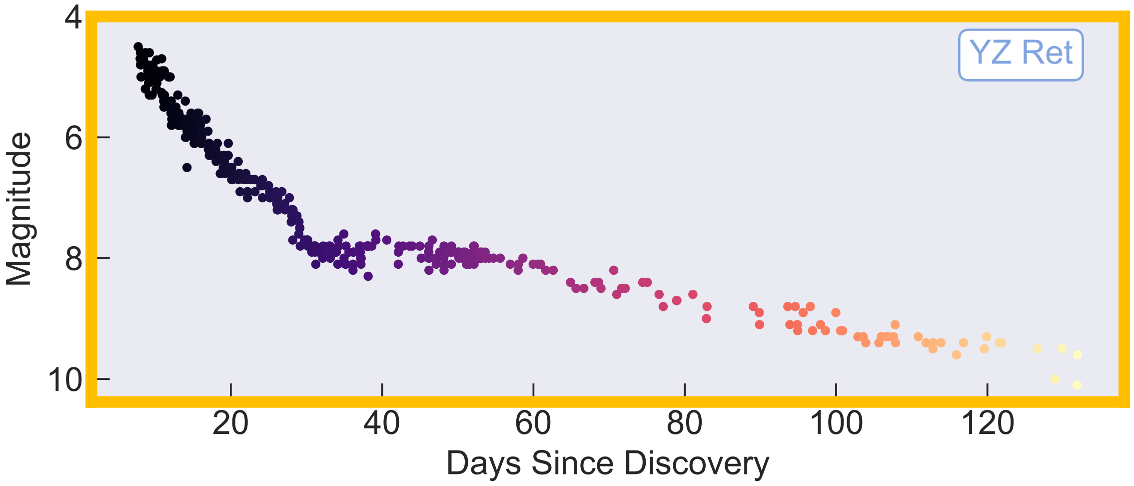

Figure 10.1: The AAVSO light curve of YZ Reticuli ranging from magnitude 4.5 to 10.1 in the visible band over the first 125 days since discovery.

−30 05 37.6 (J2000). It was classified as a classical nova by Aydi [2020b] using high-resolution optical spectroscopy from the Southern African Large Telescope (SALT).

10.2 Observations

This chapter focuses on a sequence of observations of YZ Reticuli from +8.29 d to +160 d predominantly from the Aquila spectrographs in the Global Jet Watch observatories but supplemented by some LHires spectra at the Rainbow Observatory (IAU #430) (reduced using the methodology presented in Chapter 8). A discussion of the comparison between the data from the different instruments is provided in Section 10.3.3. The cumulative exposure time on this target at the time of publication was 3.1 weeks, across a range of different exposure times which lengthened as the classical nova declined in brightness. Some spectra were intentionally saturated in Hα so as to enhance the signal-to-noise in the less intense parts of the spectrum, and these spectra are used in the analysis of weaker lines.

10.2.1 Global Jet Watch observatories

The Global Jet Watch[3] observatories are a network of five telescopes distributed in longitude (South Africa [GJW-SA], Chile [GJW- CL], east Australia [GJW-OZ], Western Australia [GJW-WA], India [GJW-IN]), described in Blundell et al (in prep),

---

[3] www.GlobalJetWatch.net

and in Chapter 3. This configuration allows for sustained observations of an astrophysical target while the Earth rotates. Each observatory contains a Ritchey-Cretien carbon fibre telescope having a 0.5m-diameter primary mirror, and a fibre-fed spectrograph known as Aquila (Lee et al, in prep).

10.2.2 Aquila spectroscopy

The predominant data product used in this chapter is a set of 5129 optical spectra captured by the Aquila spectrographs, introduced in Chapter 3. The Aquila observations start soon after discovery and, although this chapter focuses on the spectra captured prior to the submission of the paper, continue until the present day.

10.2.3 LHires spectroscopy

As mentioned in Chapter 8, this instrument was commissioned to support the Aquila spectrographs, complementing them with high-resolution ($R \sim 20,000$) auxiliary spectra. These spectra confer a high level of confidence in these fits that deconstruct the highly time-sampled Aquila spectra presented in Sections 10.3 and 10.4. It has been a fruitful strategy to obtain a high temporal density sustained data set using Aquila, with occasional supplementation from LHires.

10.2.4 Epochs of observations

Figure 10.2 presents the epochs at which these observations were captured, marked against the AAVSO light curve on the back wall. The colour of the light curve at the back is mapped to Julian Day using the same colour map as that of the Aquila observations themselves, shown as spectra in the foreground; this light curve is shown in Figure 1. The LHires spectra are also shown, interspersed in gold. This gives a general impression of the data set, as a continuous follow-up programme with highly regular observations and indicates how such an extended yet thorough campaign is made possible by the network of Global Jet Watch observatories.

10.2.5 Data reduction

The spectra were reduced with the bespoke data reduction pipeline ENDEAVOUR, introduced in Chapter 3, a full end-to-end image-processing solution which takes data captured at the telescopes, performs the reduction, and delivers one-dimensional spectra as a product to the fitting applications. It is a bespoke pipeline suited to the

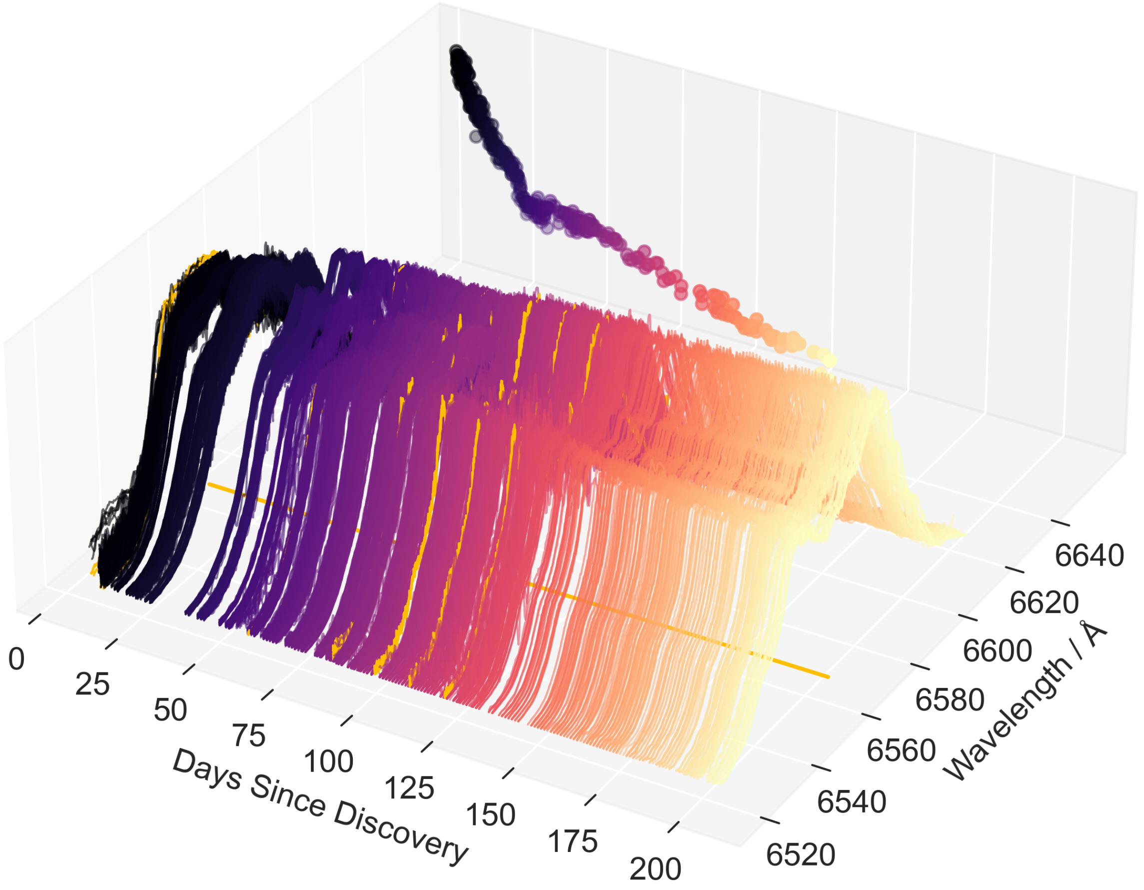

Figure 10.2: This plot indicates the epochs for which there are Global Jet Watch spectra of YZ Reticuli. The LHires spectra are depicted in gold, while the light curve and the Aquila epochs are coloured according to their Julian Date. The rest wavelength of Hα is projected to the floor (in gold) as a fiducial marker. The AAVSO light curve of YZ Reticuli is reproduced for context against the back wall, coloured as in Figure 10.1. The emission complexes shown are normalised continuum-subtracted H-alpha profiles.

specific instrumentation at Global Jet Watch observatories, but is built on widely-used open-source scientific libraries available for the Python programming language. The packages used include NumPy [Harris et al., 2020], SciPy [Virtanen et al., 2020], Pandas [Mckinney, 2010], scikit-learn [Pedregosa et al., 2011] and Astropy [Robitaille et al., 2013, Price-Whelan et al., 2018]. The pipeline ensures every spectrum is dark corrected, and calibrates wavelengths against Thorium-Krypton bulb arcs. The one-dimensional spectra are heliocentrically corrected to the barycentre of the solar system.

10.2.6 Post-processing application

Post-processing analysis is performed by `poirot`, which takes as an input the reduced 1D spectra from ENDEAVOUR. This chapter focuses on the Hα complex, and so data are typically normalised to the peak of Hα. In the fits presented here, the background has been fitted with a simple constant offset, valid since the background does not vary strongly over the Hα region. This is performed at the point of fitting the physical model, and is allowed to freely vary. The background is consistently responsible for $< 1\,\%$ of the signal in our region of interest, and so is subtracted.

During the data-cleansing step of `poirot`, it removes any spectra with either saturation on the feature of interest or adverse weather effects (determined by manual inspection). In the later stages of the eruption when changes are happening on longer timescales, the spectra taken at a given observatory on a given night are summed. When necessary, as the nova became fainter, to further improve signal-to-noise, any spectra which were not the longest exposure taken on a particular night are rejected as saturation of the Hα complex ceased to be a problem.

10.3 Fitting methods

Deducing the underlying physical model from spectral complexes is notoriously difficult in astronomy, in part because of the intrinsic richness of the phenomena, and in part because of the trade-off between spectral resolution and signal-to-noise. The Global Jet Watch Hα spectra of classical novae are no exception, with several dynamical components evolving rapidly. There are different strategies for deconstructing the underlying components that comprise emission lines profiles e.g. Schmidt et al. [2019], but the following results use a simple set of Gaussians. The merits of this approach are demonstrated elsewhere e.g. Blundell et al. [2008], Grant et al. [2020].

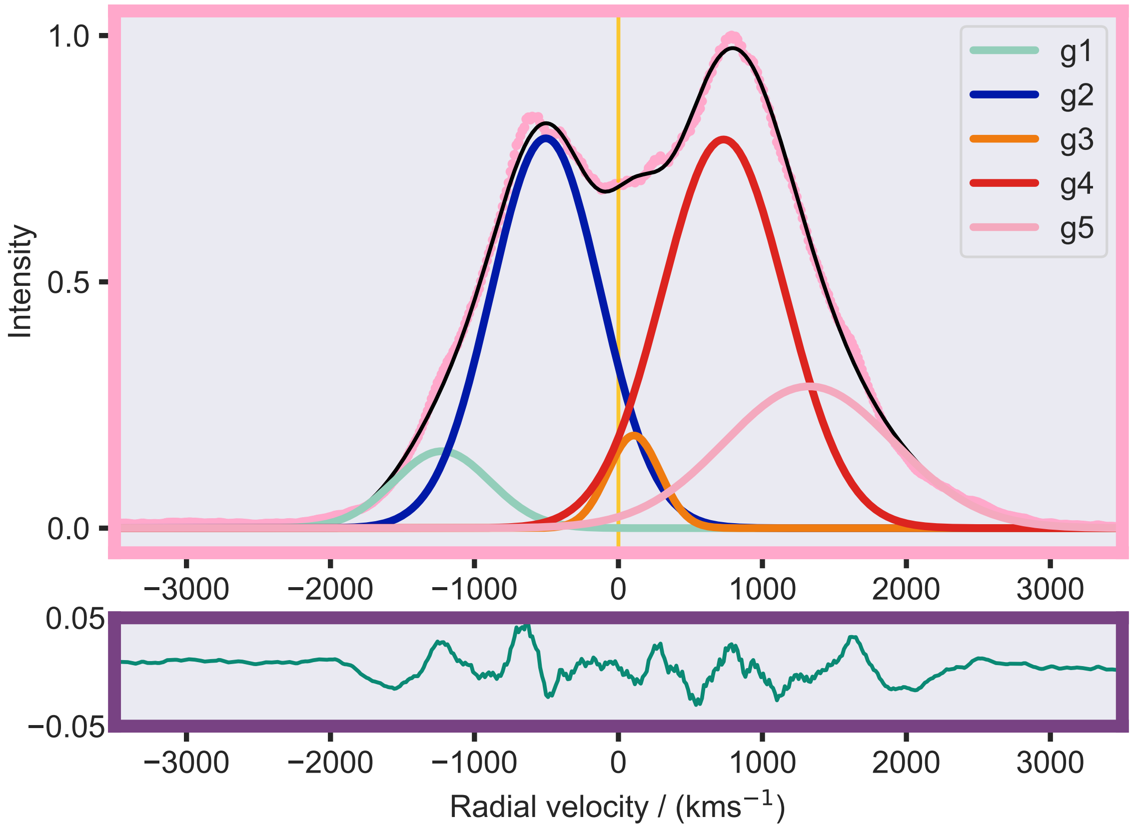

Figure 10.3: Fitted model showing two pairs of emission lines, and a central component. The spectrum (shown as pink points) was taken at the GJW-OZ observatory, at +129.56 d since discovery, with an exposure time of 1000 s.

10.3.1 Composite fits

Multiple one-dimensional Gaussian models (and a subsequently subtracted constant offset background, discussed in Section 10.2.6) were utilised in a composite fit to these data, using the Levenberg-Marquardt algorithm with a least squares statistic [Price-Whelan et al., 2018] to minimise the difference between observation and model. Residuals were analysed, and inspected to ensure that there only remained variation on velocity scales smaller than or comparable to the limit from spectral resolution. Section 10.3.3 discusses the robustness of the fits. Figure 10.3 demonstrates typical results of this fitting template, with residuals plotted in the panel below. Figure 10.4 illustrates how the qualitatively changing profile of YZ Reticuli is nonetheless well-modelled by these five Gaussian components.

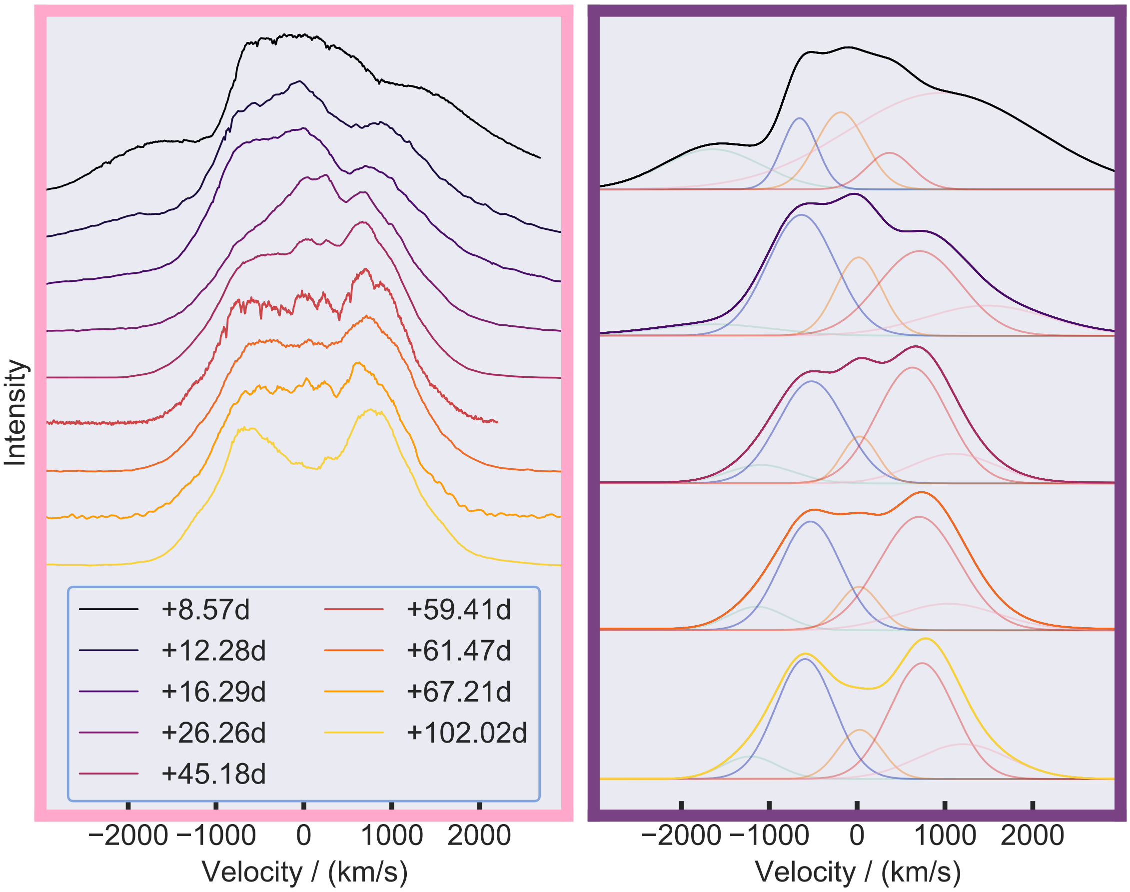

Figure 10.4: The left panel shows a sequence of H-alpha profiles. An example LHires spectrum (fourth from the bottom) covers a shorter wavelength range than Aquila. These profiles can be deconstructed into a similar set of five emission Gaussian components, shown in the right panel. Section 10.4 discusses these components in more detail.

10.3.2 Importance of coverage

The extended coverage spanning 180 days together with the consistency of the succession of fits were key to our confidence in these fits. While there are multiple possible fits which might reasonably approximate any one individual spectrum, most *a priori* possible fits may be rejected by comparing across time ranges within this data set, shown in Figure 10.2 alongside the light curve.

Section 10.4.2 discusses various changes which occur in the evolution of the fits to YZ Reticuli Hα profiles. The data set contains features with variation timescales that are significantly sub-night, which would have been missed with sparser sampling.

10.3.3 Robustness

This section considers different signals that might give a minor influence to our determinations of the underlying Hα emission components. Figure 10.5 shows Earth's own atmospheric lines (bottom panel), which consist of many narrow absorption lines. These lines do not have a significant impact on the fits, as the second panel shows clearly. Many of these lines are in doublets, although these are rarely fully resolved, even by the LHires instrument.

Another feature usually observed in the early spectra of classical novae is blue-shifted absorption (see Section 2.3). Denoted the low- and high-velocity components respectively, these hundreds of $\mathrm{km\,s^{-1}}$ and thousands of $\mathrm{km\,s^{-1}}$ absorption systems are thought to be due to the complex ejecta [Arai et al., 2016]. While these are often significant and usually need to be fitted for early spectra, in the particular case of these Global Jet Watch observations, by the time of the first spectrum of this nova and thereafter only emission components dominate; any such blue-shifted absorption components had already disappeared.

Let us also carefully consider the possible presence of forbidden [N II] lines at 6548 Å and 6584 Å. These lines could in principle be very problematic for fitting the broad Hα complexes in classical novae such as YZ Reticuli or V6568 Sgr as a full width at half maximum of $\sim 3000\,\mathrm{km\,s^{-1}}$ at late times is easily enough to encompass these wavelengths. However, the nitrogen lines would correspond to Hα speeds of $-685\,\mathrm{km\,s^{-1}}$ and $959\,\mathrm{km\,s^{-1}}$, and should exhibit similar dynamics to Hα. Indeed, Figure 10.5 shows that the profiles of two forbidden [O III] lines (4959 Å, 5007 Å) in the spectrum are strikingly similar to Hα (the [O III] emission at 5007 Å is a factor of three less strong than Hα at late times for our spectra). If the nitrogen 6548 Å and 6584 Å lines were strong, evidence of a similarly spaced pattern of emission might be expected — but there is no such evidence, and nor is there any emission component at the rest wavelengths either in YZ Reticuli, although this is not the case for V5668 Sgr, as detailed in Chapter 11.

10.4 Results

When experimenting with fits to the YZ Reticuli Hα spectra with multiple Gaussians, there was one solution set which persisted throughout the entire time series. While individual parameters changed over time as the nova progressed, two pairs of lines behaved in a correlated fashion, leading to the recognition that these fits do indeed reflect an underlying physical reality rather than a mathematical degeneracy. The qualitative shape of the overall spectra seems to change dramatically over the course of the observations (see Figure 10.4 for examples), but this section shows that only small changes to the Gaussian components reasonably explain these variations in an intuitive way. While of course it is possible to model classical nova emission line profiles with elaborate spatial models of the ejecta, the results emerging here show that a much simpler physical picture can explain a lot of the spectral behaviour without presuming any geometry in advance.

10.4.1 Best-fitting model

The template which best fits the Global Jet Watch Hα spectra consists of five Gaussian components in emission. In Figure 10.6, the left plot is a convenient representation of the evolving characteristics of each fit because it represents the Doppler-shifting movements in the centroids of each of the five emission components, and reveals any time-varying patterns where these are present. The colour code for the five components is persistent throughout this chapter (and also Chapter 11): red and blue denoting the red- and blue-shifted double peaks of an inner pair of emission lines, and pink and turquoise correspondingly for an outer (higher-speed) pair. Orange is consistently the colour of the component roughly stationary with respect to the system, which was called the central or the additional component in the paper. There seems to be a natural pairing of the two outermost emission components, since the movement of their centroids is associated throughout time. The same is true for the inner pair (coloured red and blue). This leaves the single central component, close to the Hα rest wavelength, which seems to vary only slightly in wavelength, though independently of the other four emission components. This is a persistent characteristic of the spectra throughout our spectroscopic monitoring.

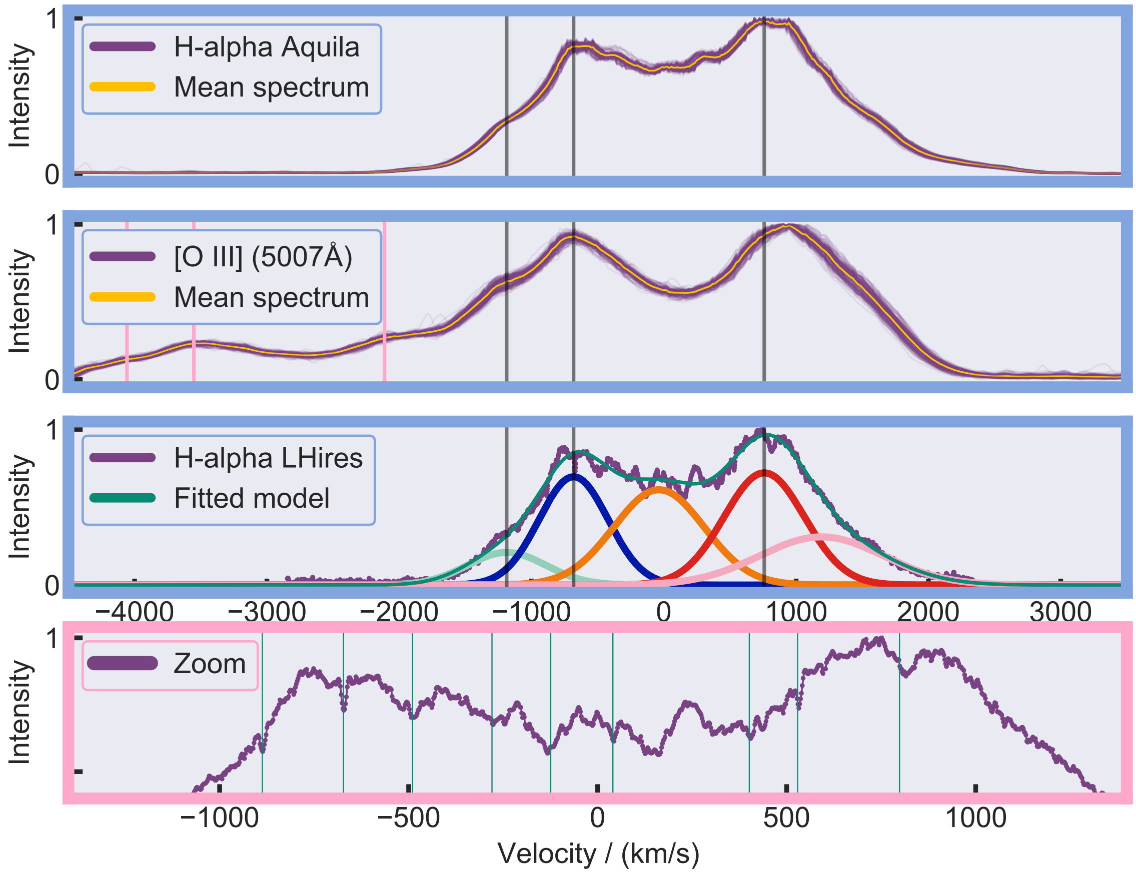

Figure 10.5: The first three panels have vertical lines marking peaks at radial velocity $-1200\,\mathrm{km\,s^{-1}}$, $-680\,\mathrm{km\,s^{-1}}$ and $760\,\mathrm{km\,s^{-1}}$ for the relevant component. Top panel: Hα complex taken on an Aquila instrument over the course of several nights (+120 d to +160 d) plotted in purple, with the mean of these overlaid in gold. Second panel: [O III] (4959 Å and 5007 Å) complex in the same spectra, with corresponding sets of markers at the same radial velocities for each of these transitions. Third panel: an example Hα spectrum from LHires at roughly the same time but significantly higher resolution shown in purple, with the sum of five fitted Gaussian emission components overlaid in teal. The vertical markers match the centroids of three of these peaks. Fourth panel: a close-up of the top of the previous panel's Hα complex, with known telluric absorption features denoted with vertical lines given by Curcio et al. [1964]. Several of these have only slightly weaker doublet partners whose wavelengths are not known as accurately as the ones depicted, hence some of these lines appearing wider than others. The close matching of the fit despite the telluric features confirms the appropriateness of the template model presented in this chapter to explore the dynamics of this classical nova. The wavelength alignment required no additional correction beyond the calibration as part of the ENDEAVOUR pipeline. Each spectrum is normalised to the peak of the complex being plotted, but the oxygen 5007 Å peak is a factor of three weaker than that of Hα.

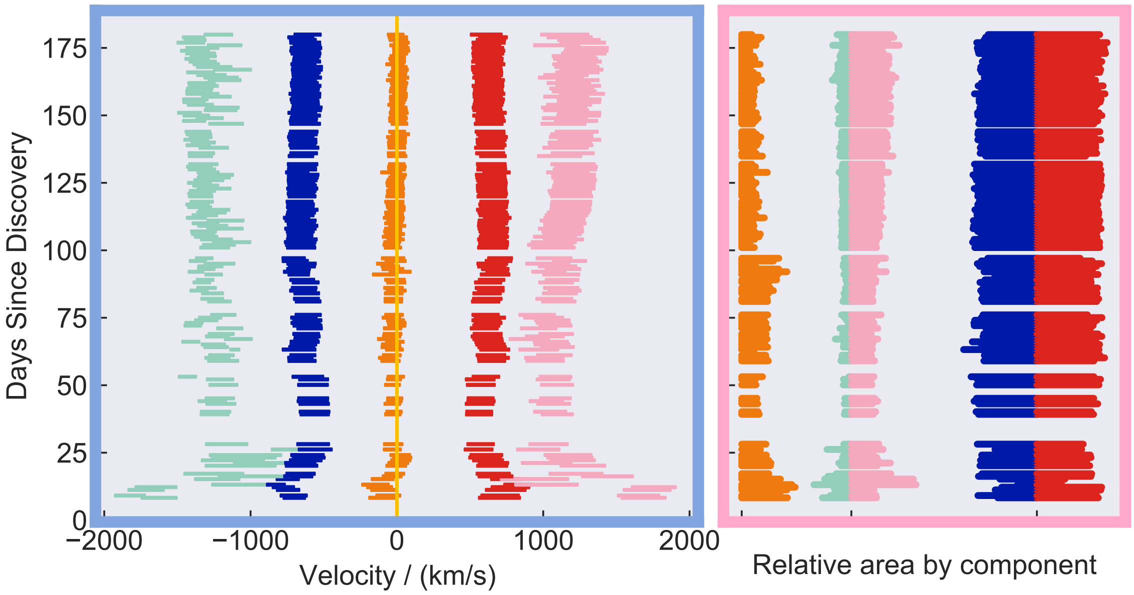

Figure 10.6: The left plot shows the centroid velocity of each of the five emission components relative to the centroid of the inner pair over time, ensuring that no subtle wavelength calibration trends are at play. The widths of the lines are proportional to the full width at half maximum of that component. The right plot shows how the area of each component develops with time, and the identified pairs of lines are plotted back-to-back to more easily see correlations. The width of each bar in this figure is proportional to the FWHM of the line, with an overall scaling that avoids horizontal overlap. The area scaling is performed after normalisation, and so these bars represent the fraction of flux contained in each component.

What is truly remarkable about this five emission component model is that it fits the data well over a large range of dates, and for a variety of seemingly qualitatively different spectral complexes. The overall shape of the Hα complex at early times displays broad emission, then from about +15 d to +95 d a generally flat peak is observed, with a surplus of red emission at about 1000 km s^{-1}. After +95 d, there is a sudden transition to a doubly-peaked profile with emission at -1000 km s^{-1} and 1000 km s^{-1}. This set of five peaks fits all of these rich data across the whole time series, with only minimal changes to the model parameters, thus preserving the relationships between the pairs.

10.4.2 Observed trends in derived parameters

In the context of Figure 10.6 (left panel), let us now consider not just how each emission component is behaving, but the relationships between them. Figure 10.7 presents the calculated relative metrics from the fitted parameters for each spectrum,

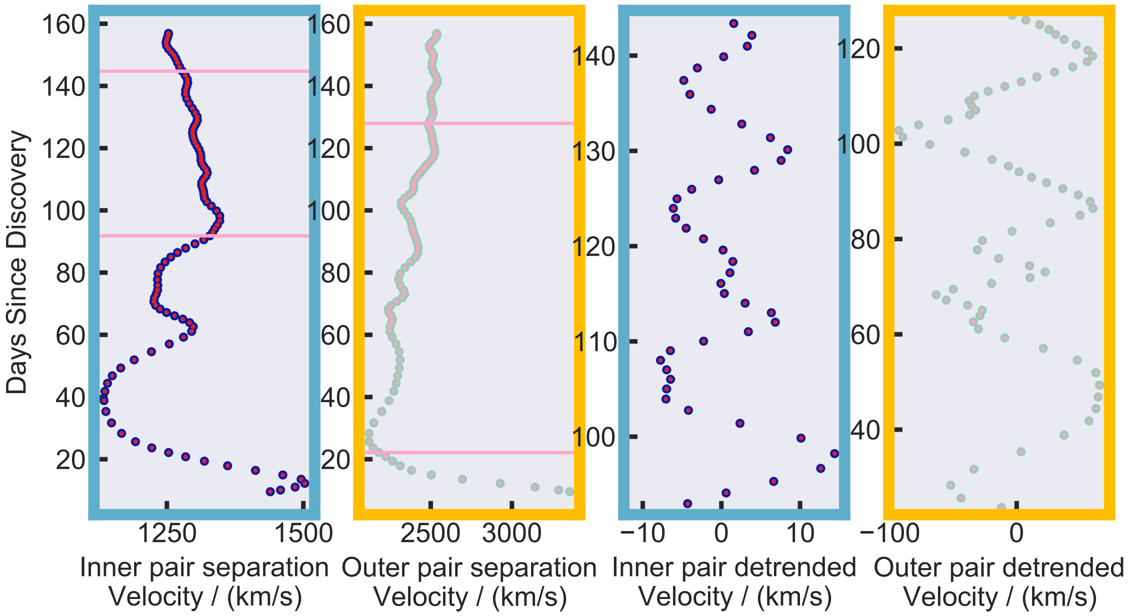

Figure 10.7: The two left panels show how the separations in velocity between each of the inner components and each of the outer components respectively vary over the full time range. Each observation is the median of the highest-quality spectra we have for that observatory on each night, and these are smoothed using a mean filter with a moving window size of three nights. The right two plots represent the same underlying data as the left two plots, but over limited date ranges, and with the linear trend removed from each to show the oscillations more clearly. The date ranges used for the right-hand plots are represented by fiducial horizontal lines on their respective left-hand analogues.

plotted over time. This technique is essential for reducing the dimensionality of the highly complex data set, and gives insights as to the underlying physics. The pertinent derived parameters include the differences between velocity centroids of the inner (blue, red) and outer (turquoise, pink) pairs of components, and the integrated area under each Gaussian component. The right panel of Figure 10.6 shows how the area of each Gaussian component varies with time. The areas of the inner (blue, red) pair of components are remarkably steady relative to one another at every epoch. The central component, depicted in orange, shows a step-change decrease in area just before +100 d. The red-shifted outer component, depicted in pink, initially fluctuates to larger and then steadily smaller areas by +50 d, after which it gradually increases in area.

Figure 10.7 depicts the time-variation of the difference in velocity between the red- and blue-shifted component centroids within each pair. This quantity (henceforth the separation) is smoothed for each pair over time using a mean filter, revealing pronounced quasi-periodic oscillations in both the inner and outer pairs. Note that this is a separation in velocities rather than pure distances. These pairs behave in a similar fashion to each other, with large velocity separations at early times ($< +50\,\mathrm{d}$) which reach a minimum separation within the first few weeks. They then both recover slightly, but continue to oscillate. The oscillations happen at different times for each pair, and the timescale of recurrence decreases over time, which is discussed further in Section 10.4.3.

At around +100 d there is an abrupt change in behaviour, whereby the oscillation of the inner pair decreases in amplitude (moving to lower radial speeds) but the outer pair maintains its steady progression to higher-magnitude radial speeds. After this date, the predominant general trend is of a much more steady increase in velocity by $\sim 2.5\,\mathrm{km\,s^{-1}}$ per day, but superimposed on this is the distinct but much more subtle oscillation of $\sim 10\,\mathrm{km\,s^{-1}}$ with a characteristic timescale of ~ 10 days.

On removing the overall trends in the two leftmost panels of Figure 10.7 by subtracting a straight-line fit from each, the underlying oscillations are recovered more clearly and these are plotted in the two rightmost panels for a restricted time range (marked in the left plots with horizontal lines). This representation shows the oscillations in the clearest manner. Upon folding by the strongest peaks in the Fourier transform, it became apparent that the period of both these pairs was in fact changing by several days each cycle, so this is in fact a quasi-periodic oscillation in the velocity separation. There were no statistically significant persistent periodicities, most likely because the duration of each cycle was quite different from that of the preceding cycle.

However, the pattern of irregular oscillations with similar velocity amplitude is clear, and we take this to mean that these pairings of lines reflect pairings in the underlying physical reality. In the two rightmost panels of Figure 10.7, there also appear to be hints of smaller-scale nutations superimposed on the large oscillations, but even finer time resolution would be needed to quantify these accurately and fully separate the various perturbations.

The inner pair velocity separation decreases by $120\,\mathrm{km\,s^{-1}}$ between +100 d and +150 d, while the outer pair separation steadily increases from $2200\,\mathrm{km\,s^{-1}}$ to $2600\,\mathrm{km\,s^{-1}}$ between +60 d and +150 d, which is discussed in more detail in Section 10.5.1.

10.4.3 Dynamical changes

What is fascinating about the oscillatory behaviour of the inner and outer pair seen in both Figure 10.6 and Figure 10.7 is that it is not fixed in either amplitude or period, but instead these quantities both decline over the course of the first six months. The implication here is that this reflects some property of the underlying system which is returning to an equilibrium. Furthermore, these timescales are substantially longer than both the canonical classical nova binary orbit timescale ($\sim 4\,\mathrm{hr}$) and the canonical white dwarf spin period in classical novae ($\sim 76\,\mathrm{s}$), meaning this effect takes tens to hundreds of orbital periods to perform a cycle.

10.5 Physical interpretation

The data clearly fit well with a template of 5 Gaussian emission components: a central component, an inner pair of components with radial velocity relative to the centre of $\sim 600\,\mathrm{km\,s^{-1}}$, and an outer pair of $\sim 1200\,\mathrm{km\,s^{-1}}$. There are several possible models which may initially appear to explain this clear signal. The key components in the system are thought to be the accretion disc, the hot white dwarf, the secondary star, and various forms of ejecta. We disregard the Hα flux as arising from the secondary itself, since the magnitude of the whole system was $>= 15$ during quiescence (ASAS-SN data in Section 10.1.3).

10.5.1 Jets and accretion disc

The natural interpretation of the outermost widely-separated and broad emission lines is as oppositely-directed collimated outflows, which may reasonably be described as jets [Iijima and Esenoglu, 2003]. Given the prevalence of accretion discs within jet-ejecting systems, it is possible that the inner pair of lines could arise from the accreting material either because of its rotation or via winds off the accretion disc.

I propose that the five Gaussian emission components within the Hα complex are best explained by a model involving jets, an accretion disc and an additional component henceforth the 'jets' model. This model therefore posits that all five observed Gaussians have some connection with an accretion disc. The outer pair of lines (turquoise, pink) is caused by bipolar anti-parallel jets launched as a consequence of an accretion process, while the inner pair of lines (blue, red in Figure 10.6) may be a manifestation of the standard double-peaked rotating disc signature or the wind off this rotating matter. This is a simple suggestion, and is exactly what is seen arising from the accretion disc in the microquasar SS433 [Blundell et al., 2008]. As noted by Long and Knigge [2002], the evidence for winds off the discs in cataclysmic variable systems is compelling. The fifth, central component has a significantly lower line-of-sight Doppler shift and may be a shell arising from local matter that is not dynamically important.

This model is supported by the fact that there seem to be distinctive epochs where characteristics change in a related way in each of the components, sometimes after a short delay; these changes occur most notably around +100 d. This is explained naturally under the jets model, because the lines all have a common origin in the accretion disc. Furthermore, the evolution of the pairs of lines and their oscillatory behaviour makes most sense when considering them as being derived from a precessing (perhaps unbalanced) disc.

After about +100 d, the inner pair is observed to move to steadily lower speeds, while the outer pair increases its velocity separation. Classical nova events differ from normal cataclysmic variable systems by their enhanced mass-transfer rates. The increased accretion rate may lead to their exceeding the threshold necessary for launching jets [Lasota and Soker, 2005], a situation that can persist for at least a few months following eruption. Since not all the ejecta receive enough energy to be fully ejected, some gas can be trapped and forms an envelope surrounding the inner system. As this gas steadily rejoins the accretion disc, it is expected that there exists a region of temporary overdensity, which could then feed the jet. As proposed in the paper, it is reasonable to suggest that the trends in pair speeds are caused by such an effect. The inner pair decreases in speed, as the accretion disc is depleted at its inner radius (where it is fastest). The outer pair increases in speed, as the accretion disc processes this extra material and launches increasingly powerful jets. This assumes the outer pair of lines are produced near the base of the jet, which is expected since at high rates of adiabatic expansion into relative vacuum, the plasma becomes optically thin.

Simulations conducted by Figueira et al. [2018] under particular conditions result in the accretion disc being fully disrupted by the shell ejecta within the first forty minutes after outburst. This occurs when the mass of the ejecta is assumed to be high compared with the mass of the accretion disc. A typical timescale for reforming a disrupted accretion disc has been estimated at around hundreds of years. However, as Figueira et al. [2018] mention, accretion discs have been observed in novae within months of eruption, implying that they either are re-established more rapidly than theoretical predictions or indeed are never fully disrupted. Moreover, it is possible that accretion can be efficacious even when the disc has not yet attained a state of perfect equilibrium. Hernanz and Sala [2002] show the resumption of accretion after only a brief period of disruption. Also, the best-fitting model to the progenitor spectral energy distribution in V906 Car includes an accretion disc, shown comprehensively in Wee et al. [2020] and used to explain changes in the colour evolution of that classical nova around 50 days after outburst. The spectral signature of the accretion disc, or potentially the wind off its surface, is visible beyond the photosphere. The very early time-series data on YZ Reticuli shows a rapidly-evolving regime which may be caused by an accelerated resumption of accretion processes.

A succession of over-pressured jet ejecta could plausibly expand into each other, causing shocks and hence particle acceleration giving rise to the X-ray shocks observed by NuSTAR reported in Sokolovsky [2020]. This is exactly analogous to the behaviour exhibited in the microquasar SS433 as its Hα-emitting jet bolides expand then collide into one another, causing non-thermal, polarised synchrotron emission after shock acceleration [Blundell et al., 2018].

10.5.2 Precession of the accretion disc and its jets

The leftmost teal panel in Figure 10.7, which plots the velocity separation of the inner pair with time, shows oscillations on two different characteristic timescales, one being about 40 days prior to about +100 d, and thereafter about 10 days (zoom-in in the third panel). These oscillations show a general tendency to decrease in both amplitude and period for the entire duration. The first phase of oscillations might be a response of the accretion disc to the rapid influx of material in the immediate aftermath of the eruption. The existence of the inner lines only a few days after the initial classical nova event indicates a pre-existing (albeit perturbed) accretion disc in the YZ Reticuli system. After about +100 d, the amplitude of the oscillations is significantly smaller, suggesting damping and equilibration have happened to some extent.

The parameter space occupied by cataclysmic variables, and a discussion of the precessing accretion discs in these systems, is outlined in Patterson [2001]. Charles et al. [2008] give a summary of the observational properties of accretion discs in X-ray binaries, many of whose time-varying/periodic characteristics have been interpreted as precession on a wide variety of different timescales.

It is plausible that the oscillations seen in the leftmost teal panel in Figure 10.7 may arise from reactive precession in a (likely asymmetrically) perturbed accretion disc. It is therefore interesting to consider that the oscillations seen in the gold panels (2 and 4) of Figure 10.7 arise from a change in axis of the jet ejecta given the ongoing re-orientation of its launch platform i.e. accretion disc. In the microquasar SS433, its accretion disc has been shown to exhibit a persistent periodic pattern of precession and nutation tightly coupled to that exhibited by its famous precessing jets (figure 2 of Blundell et al. [2008]). In contrast to the quasi-persistent periodic precession of SS433, YZ Reticuli's accretion disc appears to precess on shortening timescales as a response to the disruption following the classical nova outburst, perhaps due to mass-loss through the jets themselves, so a tight link between this and consequent jet precession might not be expected. Regardless of the duration of the precession period, the accretion disc and jets are expected to exhibit a quarter-phase lag because of their perpendicularity from one another. It is therefore interesting to note in the leftmost two panels of Figure 10.7 that there are distinct hints of such a phase difference between jet oscillations and accretion disc oscillations that are consistent with such a picture.

10.5.3 The behaviour of the central component

One phenomenon observed in the succession of YZ Reticuli spectra is that several emission species appear to transition from a profile with a central peak to one characterised by wider double-peaks. The first such transition event occurs suddenly at +63.21 d in Hα. It then recurs multiple times, most dramatically at +93.93 d (lasting until +95.92 d). At +96 d, the transition repeats, steadily deepening across the next few hours and drastically deepening again at +101.9 d. Eventually, the Hα complex remains in a split state. This unsteady approach to eventual stability could be due to the temperature of the emitting body of gas dropping below a relatively sharp cutoff, with some fluctuations in temperature causing it to temporarily return above the threshold before settling below.

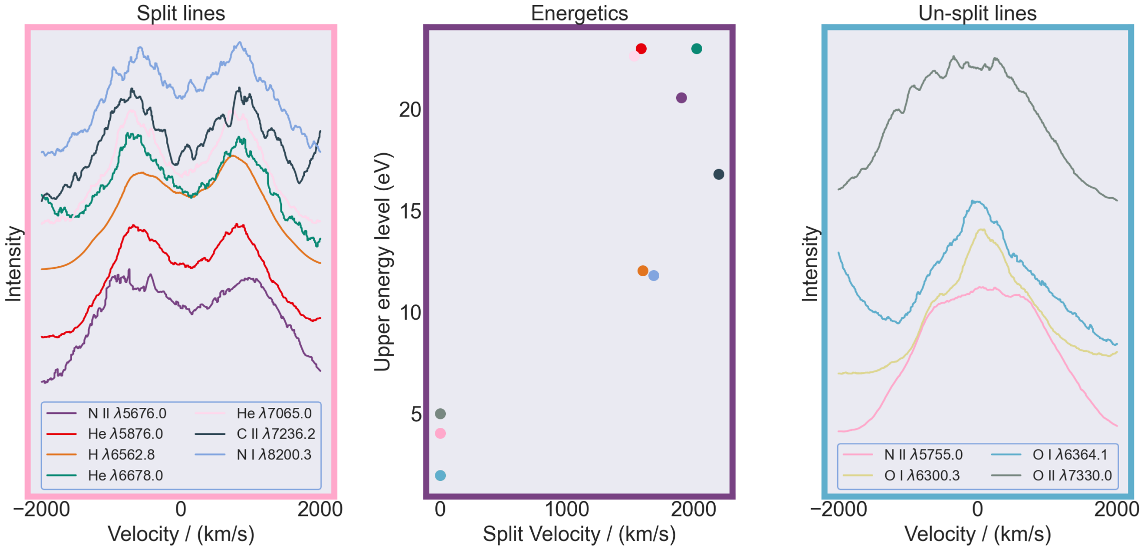

Figure 10.8: Average spectra of different emission species taken at +94 d of YZ Reticuli. The left panel shows certain species which display the splitting phenomenon, the right panel shows the lines which do not exhibit this splitting. The central panel shows the excitation energy of the upper level of the identified transitions responsible for each line, plotted against the velocity difference of the two peaks in the case of the split lines. Note that this is not to be confused with the quantum effect of line splitting, it is purely dynamical.

Figure 10.8 shows the velocity profiles of the most prominent lines in our spectra. Some of the lines (shown in the left panel) display this split feature, and others do not (shown in the right panel). The central panel shows the split velocity (difference between the two peaks, zero if only one peak) against the upper energy level of the atomic transition responsible, giving a strong positive correlation. The transitions that do not appear to split show no qualitative sign of significant changes before and after +95 d.

It is apparent, upon close examination of the fits to the Hα complex in a time series around this epoch, that in fact this splitting is caused by the central dynamical component growing weaker relative to the inner and outer pair. While the other lines are fainter and as such have lower signal-to-noise than Hα they correspond very closely in velocity space, so it is likely that they are also appearing to split because a corresponding central component is growing dimmer.

Since the splitting seems to only occur in the higher upper energy level transitions, this is evidence that the origin of the central component has cooled to below 10 eV by +100 d.

10.5.4 Why are radio jets not observed in classical novae?

While radio jets are indeed observed in *recurrent* novae such as RS Oph [Sokoloski et al., 2008], as detailed fully in Section 10.1.2, they have not yet been found in *classical novae*. However, the existence of a jet whose radiant gas gives rise to Hα emission does not necessarily imply the existence of a radio-emitting synchrotron jet. In the case of the jets from the microquasar SS433, the former leads to the latter only because of the expansion of the jet bolides into one another causing shocks and consequent particle acceleration [Blundell et al., 2018] evinced by particular signatures in the synchrotron polarisation coinciding with where the expansion and collision happens. In the case of a binary whose Hα emitting jet ejecta either (i) did not expand sufficiently so as to collide into one another causing a shock-accelerated synchrotron emitting population or (ii) were not propagating through a magnetised region of space, no synchrotron radio jet would be observed.

It is interesting to note that Toalá et al. [2020] require a non-thermal component to fit the spectral shape of elongated X-ray emission of DQ Her shown in their figure 3; in McLoughlin et al. [2021a], we conjecture that this X-ray emission arises from inverse Compton emission of spent synchrotron electrons with the radiation field from the hot post-eruption white dwarf surface.

10.5.5 Why is the inner pair not just a spherical shell?

Given the resolved images of approximately spherical shells around many classical novae it might be tempting to associate the inner pair of emission lines with a spherically ejected shell as discussed in Section 10.1.1. However, a ballistically-ejected outflow would not continue to modulate in an apparently symmetric way nor change speed several months after the eruption (Figure 10.7, panel 3). Moreover, this could not explain the sometimes correlated changes in the characteristics of the inner pair and the outer pair of emission components. As described in Section 10.4.3, the magnitude of oscillations in the relative velocity of the inner pair components consistently decreases over time. While this is expected behaviour for a precessing accretion disc re-establishing equilibrium, there is no clear explanation for why this would happen to opposite sides of a spatially-extended ejected shell.

In addition, a spherical shell would be too hot within the first few weeks of the eruption ($\sim 10^8\,\mathrm{K}$) to emit Hα, which requires a more modest temperature regime (a few $10^4\,\mathrm{K}$). Indeed, the reason why the jets manifest themselves in Hα spectra is because they are far cooler, which is consistent with their having been launched further out in radius by the accretion disc rather than being the products of thermonuclear runaway on the surface of the white dwarf.

10.5.6 Alternative models

Several models were carefully considered before settling on the jets model. The primary such alternative model is that the inner pair is a slow-moving outflow, considered in Section 10.5.5.

Another is that one of the pairs could be formed in a circumbinary disc. While there is preliminary evidence pointing towards circumbinary discs in other novae such as V906 Car (as presented in Chapter 9, following McLoughlin et al. [2020]), the speeds involved here are significantly faster, implying either rotation within the binary system or high-speed outflow.

Another alternative model is that one pair of lines is formed in an accretion disc while the other pair are the fast bipolar ballistic ejecta. The difference between this and our primary model is that in this case, the ejecta are not launched by the accretion, but are direct ejecta from the surface of the white dwarf. This model acknowledges that there has been a recent explosion and as such matter should be flying outwards, but it fails to explain why the inner pair and the outer pair should continue (i) at all and (ii) to oscillate, sometimes in related ways, to late times. Under this scenario, it would be difficult to explain why the bipolar outflow would be visible in Hα, as it would be far too hot.

10.6 Similar characteristics in V6568 Sgr

Given the theoretical basis for jets in classical novae set out in Kato and Hachisu [2003], it is tempting to consider whether this model applies to other systems in a prevalent manner, or if YZ Reticuli was an unusual event. In the same month, July 2020, there was another classical nova eruption, V6568 Sgr, of the very fast variety. Confronting this with the same model revealed that it was largely similar, albeit with a non-detection of the central component. The speeds of the outermost pair were $-4100\,\mathrm{km\,s^{-1}}$ and $4500\,\mathrm{km\,s^{-1}}$, while the speeds of the inner pair were $-1000\,\mathrm{km\,s^{-1}}$ and $1900\,\mathrm{km\,s^{-1}}$. Although V6568 Sgr was substantially faster to fade than YZ Reticuli, it nonetheless displays strikingly similar pairings. Figure 10.9 shows that the model presented in this chapter for YZ Reticuli fits to this other classical nova well, despite the vast difference in speed class. Note that the spectrum appears to be noisier than those of YZ Reticuli — this is expected, because it is so much

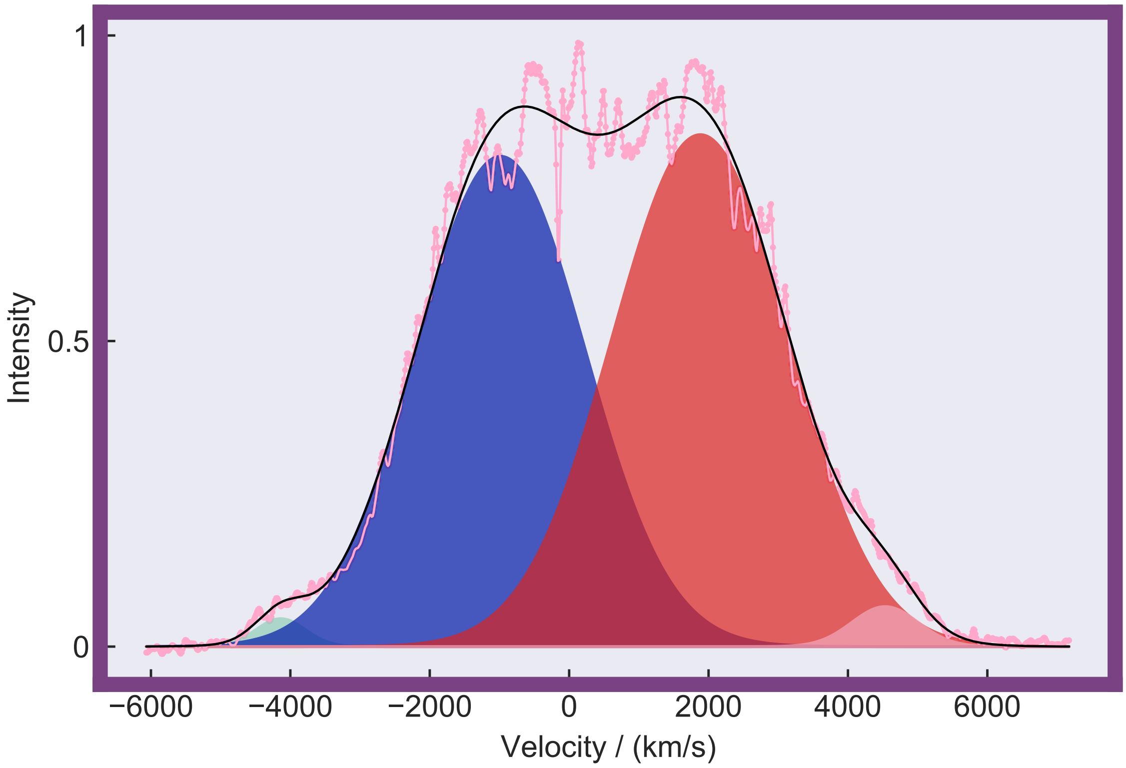

Figure 10.9: V6568 Sgr spectrum taken at +1.03d after discovery. This fits the same model as YZ Reticuli, apart from the notable absence of a central component.

fainter, but also because the jet speeds were higher, and this was taken with the LHires instrument which has a much smaller angstrom-to-pixel ratio. As such, there are fewer photons per wavelength bin, and the telluric absorption lines play a bigger role. Notwithstanding this, the same model fits, with the exception of a greatly diminished central component, which may be interpreted as any optically thick shell having disappeared. The two pairs of lines are very symmetric, and it was not possible to achieve a sensible fit without an outer pair; despite their relatively lower amplitude in comparison to the inner pair, they are nonetheless crucial. I therefore conjecture that V6568 Sgr bridges the gap in parameter space, specifically that of white dwarf mass, between YZ Reticuli and the recurrent novae.

10.7 Conclusions

In this section, I have investigated the rapidly evolving spectral characteristics of the classical nova YZ Reticuli, especially its Hα complex, since shortly after its discovery in July 2020 with time-resolved spectroscopy from the Global Jet Watch. Remarkably, a model for the Hα complex comprising the same family of Gaussians throughout the first six months fits the data well across a variety of qualitatively different spectral shapes. These five components appear to be naturally categorised as two pairs of lines, representing an accretion disc and jet outflows, together with a small central component.

Correlated and symmetric changes in these two pairs of lines suggest that the accretion disc and jets appear to be precessing, probably in response to the recent large perturbation caused by the nova eruption.

The jet and accretion disc signatures are followed from the first spectra at around +10 d until the present day, which we take as evidence that the accretion disc survived the blast, albeit in a non-equilibrium state for the first ~ 100 days.

Confirmation of this interpretation for YZ Reticuli could be possible from future VLBI observations giving milli-arcsecond-scale resolution radio imaging akin to the images of the radio jets in recurrent nova RS Oph observed by Sokoloski et al. [2008] but only if the Hα jets give rise to particle acceleration and if there is a sufficiently strong magnetic field present. However, it would undoubtedly be fruitful to consider time-resolved spectroscopy from other classical novae to establish the generality of Hα jets, and I present the results of such an analysis in the following chapter.

Chapter 11

The jets phase of classical novae

The work presented in this Chapter has been published as McLoughlin et al. [2021b].

Abstract

This chapter presents two further classical novae, V906 Car and V5668 Sgr, that show jets and accretion disc spectral signatures in their H-alpha complexes throughout the first 1000 days following their eruptions. From extensive densely time-sampled spectroscopy, the appearance of the first high velocity absorption component is measured in V906 Car, along with the duration of the commencement of the main H-alpha emission. The time taken for V5668 Sgr to transition to the nebular phase is constrained using [N II] 6584Å. These timings are found to be consistent with the jets and accretion disc model for explaining optical spectral line profile changes in classical novae, and the implications of this model for understanding of the enrichment of the interstellar medium are discussed.

11.1 Introduction

Classical novae are explosive events which occur on the surface of white dwarfs found in close binaries. Hydrogen accretes onto the degenerate surface presumed to be via a disc mechanism, and heats up until it reaches temperatures of around 10^7 K, when it ignites and thermonuclear runaway occurs. This ejects a shell of the accreted hydrogen, whose mass is dependent on that of the underlying white dwarf. The exact nature of the outflows from such an eruption are not well known, and there are competing theories about interactions between the potentially complex ejecta and the immediate surroundings.

One particular topic about which not much is known is the survival of the accretion disc in the immediate aftermath of the eruption. Theoretically, it is thought that such a disc should be destroyed by the eruption, yet there is a growing body of observational evidence that discs may indeed survive in some form. Chapter 10 (published as McLoughlin et al. [2021a]) presented, from an analysis time-resolved spectroscopy, evidence for the persistence of both jets and an accretion disc in the weeks and months that followed the eruption of classical nova YZ Reticuli. This chapter add two new examples, V5668 Sgr and V906 Car, which also seem to display the spectral signature of accretion discs and jets from early times after the eruption, timing the onset of their becoming spectrally significant in Hα observations.

Objects classified as classical novae display a rich variety of behaviours, yet there are many phases which most novae appear to share. Further to the discovery of jets, this chapter also presents densely time-sampled observations of V906 Car and V5668 Sgr to tightly constrain the durations of certain transitions between phases for these particular objects. Section 11.2 explains the data collected for each target. Section 11.3 quantifies the duration of certain phase changes for V906 Car. Section 11.4 contains a similar set of analyses for V5668 Sgr, which reveal the spectral signature of jets and an accretion disc. Section 11.5 discusses the jets and disc of V906 Car and the overall picture of classical nova evolution in light of this new model, and considers the impact nova jets may have on enrichment of the interstellar medium. Section 11.6 concludes by recommending new observations to target direct observations of such jets.

11.1.1 Classification schemes for classical novae

Over the years, there have been several efforts to categorise different aspects of the hugely varied range of behaviours exhibited by classical novae in spectroscopic, photo-metric and polarimetric data. Also discussed in Section 2.3 of this thesis, a thorough account of our modern understanding of the phases typically exhibited by classical novae is given in section 4 of Della Valle and Izzo [2020], and so here just a brief recap of a relevant subset of this material is provided.

Williams et al. [1991], Williams [1992, 1994] performed a deep analysis of optical spectra for a set of 13 novae, and found that they separate into two spectral classes - Fe II and He / N . They also found that some classical novae exhibit a "hybrid" evolution, changing between the two types.

Light curves of classical novae can be organised into several broad types, such as "dust dip", "jitter" or "plateau", based on features they exhibit [Strope et al., 2010]. They also have speed classes, with typical designations based on the time to decline by two optical magnitudes (t_2) from maximum, although this can be difficult to assign in cases when the maximum is not well-sampled or the decline is not smooth. It is thought that the faster a classical nova is, the more massive the white dwarf is, and the more rapidly it will recur [Kato, 2002, Shara et al., 2018]. Speed class therefore may relate to the propensity of a system to launch jets along the axis of spin of the white dwarf.

The study of classical novae in the X-ray has revealed a clear distinction between hard and soft X-ray radiation, with the former being associated with shocks in outflows and the latter with the surface burning. Schwarz et al. [2011] performed a population study of the X-ray behaviour of classical novae, and determined that systems with a faster speed class show early hard X-rays, which are not present for slow novae.

11.1.2 Jets and shell ejecta in novae

Chapter 10, which followed the previous paper in this series [McLoughlin et al., 2021a], studied the Hα complex of classical novae YZ Reticuli, and found it to be well-fit by the same five Gaussian emission components throughout many months. Those components naturally divided into two pairs, with each consisting of a red-shifted and blue-shifted component which exhibited anti-correlated quasi-oscillatory velocity evolution, as well as an additional component which was not dynamically significant. The pairs were found to be precessing, with the higher-speed pair of lines lagging the lower-speed pair, implying that the former pair of lines should be interpreted as having been formed in jets and in accreting material related to the disc.

McLoughlin et al. [2021a] further found that a similar model accounted for the emission in V6568 Sgr, a classical nova which rapidly faded and so did not warrant

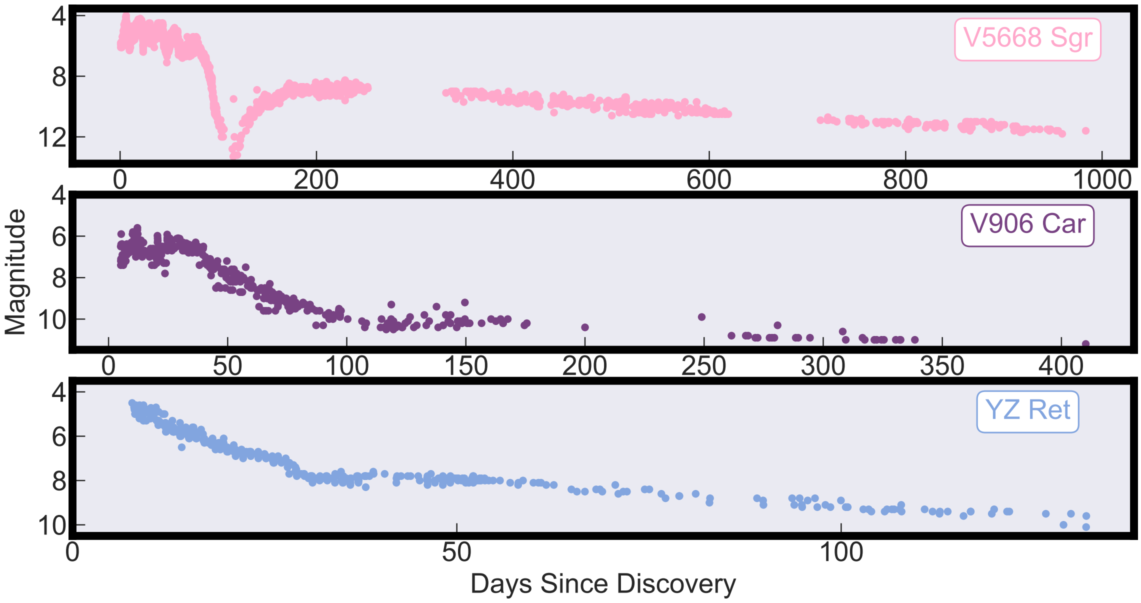

Figure 11.1: AAVSO light curves for the three classical novae studied in this chapter in chronological order. V5668 Sgr shows a deep dust-dip phase between +80 d and +200 d, with some oscillations immediately following the maximum lasting until at least the start of the dust-dip. V906 Car has initial flares, shown to be correlated with X-ray emission [Aydi et al., 2020b]. YZ Reticuli exhibits a plateau between +31 d and +62 d. The gaps in the V5668 Sgr light curve (commencing at roughly +250 d and +620 d) correspond to when it was a daytime object, as this classical nova lies in the zodiacal plane.

a major observing campaign. Given how clearly this signal persisted for YZ Reticuli, this chapter extends earlier work by examining historical data captured in two major campaigns by the Global Jet Watch on two previous classical novae, V906 Car and V5668 Sgr, in light of the jets and disc model found to describe the behaviour of YZ Reticuli.

11.1.3 Classical nova V5668 Sgr

Classical nova V5668 Sgr was discovered by John Seach on 2015 March 15.634 UT (= JD 2457097.134), which is taken as +0 d for this paper. It is also known as Nova Sgr 2015b, and PNV J18365700−2855420. It has (J2000) Right Ascension 18 37 39.9 and Declination −29 04 03, and Galactic coordinates of 005.3211, −10.0678. It was classified as being of type Fe II by Banerjee et al. [2015] using near-infrared observations. This classical nova showed a prominent dust-dip in the AAVSO[1] light

---

[1] aavso.org

curve (shown in the first panel of Figure 11.1) from approximately +80 d to +200 d [Banerjee et al., 2015]. It was found to have a white dwarf mass of 0.85 $M_\odot$ by Hachisu and Kato [2019], using their scaled light curve method.

Gehrz et al. [2018] describes in detail the X-ray story according to *Swift* observations of V5668 Sgr. Their first detection in X-ray was at JD 2457191.5, which corresponds to +94 d. They derive supersoft source switch-on and switch-off times as t_{on} =+168 d and t_{off} =+240 d to +340 d. They find the mass of the underlying CO white dwarf to be 1.1 $M_\odot$.

Molaro et al. [2016] found V5668 Sgr to be a significant producer of ^{7}Be II, and used this to suggest that classical novae may be the source of all the ^{7}Li in the Milky Way. Diaz et al. [2018] discovered structures in ALMA imaging of the ejecta of V5668 Sgr which were the smallest scale of those found in the remnants of any stellar explosion to date, although the veracity of the details depends on UV-plane coverage of the interferometric ALMA observations.

Jack et al. [2017] performed a detailed identification of four optical spectra of V5668 Sgr and compared their spectroscopy with the light curve to show that certain blue-shifted absorption features moved to slower radial velocities during the multiple rises to maximum (present until ∼ +80 d in the light curve of V5668 Sgr in the upper panel of Figure 11.1), and to faster radial velocities during the decline. Section 11.4.4 briefly expands on this work.

11.1.4 Classical nova V906 Car

The earliest detection of classical nova V906 Car was at JD 2458193.819444 by the Evryscope-South observatory, according to Corbett et al. [2018], which is taken as the discovery date for this system. It was independently discovered after this by Stanek et al. [2018] using the All Sky Automated Survey for SuperNovae (hereafter ASAS-SN). This chapter references times as days since commencement of brightening (JD 2458193.819444). V906 Car has (J2000) Right Ascension 10:36:15.426 and Declination −59:35:53.731, and its Galactic coordinates are 286.580, −01.088.

The mass of the white dwarf in V906 Car was estimated at $0.71^{+0.23}_{-0.19}\,M_\odot$ through an analysis of the progenitor system by Wee et al. [2020] using an accretion disc model to understand the light curve. Their decline times of $t_{2,v} = 26.2$ d and $t_{3,v} = 33.0$ d point to V906 Car's being a moderately fast classical nova, although Aydi et al. [2020a] suggest a value of $t_2 = 44$ d. Sokolovsky et al. [2020b] determined that this classical nova occured on a CO white dwarf, and further that the mass must have been low as there was no distinct supersoft phase detected in the X-ray. Their thorough analysis of the NuSTAR, Fermi/LAT, XMM–Newton and Swift/XRT X-ray data revealed that the nova shell was expelled by +24 d. The white dwarf was found to be of CO-type by Sokolovsky et al. [2020b].

The optical AAVSO light curve shows some initial flares up to around +30 d, and exhibited weak cusp-like behaviour between +100 d and +150 d. Aydi et al. [2020b] used correlated flares in the optical and gamma-ray light curves of V906 Car to show that the emission in both bands is probably powered by the same underlying mechanism, which they took to be fast outflow colliding with slower outflows, and shock-heating it. Observations by Pavana et al. [2020] show V906 Car to be of the hybrid class of classical novae. Wee et al. [2020] associate V906 Car with the Carina Nebula, and thus assume a distance to the classical nova of $d = 2.3 \pm 0.5$ kpc.

11.1.5 Classical nova YZ Reticuli

YZ Reticuli was discovered by Robert H. McNaught (Coonabarabran, NSW, Australia) at magnitude 5.3 on 2020 July 15.590 UT (CBET 4811). It is also identified as Nova Reticuli 2020, and MGAB-V207, with (J2000) Right Ascension of 03 58 29.55 and Declination −54 46 41.2 (J2000), its Galactic coordinates 265.397, −46.395. Following Chapter 10 and McLoughlin et al. [2021a], this chapter takes the date of discovery as 2020-07-08.1708, or JD 2459038.6708.

The lower panel of Figure 11.1 shows the AAVSO light curve of YZ Reticuli for the 125 days following eruption. YZ Reticuli exhibits a plateau-type light curve [Strope et al., 2010] between ∼ +30 d and ∼ +60 d, which may relate to the plateau in the light curve of classical nova V407 Cyg, interpreted by Hachisu and Kato [2012] as a surviving accretion disc emerging out of the receding photosphere. The decline time of this classical nova is $t_2 = 15$ days during the initial decline (although this changes during the plateau event) [Sokolovsky et al., 2020a]. A summary of the behaviour in X-rays and in gamma rays is given in McLoughlin et al. [2021a]. As discussed in that paper, the GAIA distance is $2.7^{+0.4}_{-0.3}$ kpc [Bailer-Jones et al., 2018]. It was classified as a He /N classical nova [A. Carr et al., 2020], while [L. Izzo et al., 2020] suggest that this is an ONe white dwarf.

11.1.6 A heterogenous sample

The set of novae examined includes one of each kind of the Tololo scheme for classical nova spectra; one Fe II, one He /N , and one hybrid as detailed in Table 11.1. The light curves (shown in Figure 11.1) are thoroughly distinct, through their classifications,

Table 11.1: A comparison of properties for the three classical novae.

[a] Banerjee et al. [2015]; [b] Pavana et al. [2020]; [c] A. Carr et al. [2020]; [d] Aydi et al. [2020a]; [e] McLoughlin et al. [2021a]; [f] Sokolovsky et al. [2020a]; [g] Gehrz et al. [2018]; [h] Sokolovsky et al. [2020b]; [i] L. Izzo et al. [2020]; [j] Wee et al. [2020];

Target	V5668 Sgr	V906 Car	YZ Ret
Tololo class	Fe II[a]	Hybrid[b]	He / N []c
Light curve	Dust-dip[a]	Jitter/Cusp[d]	Plateau[e]
Decline time t_2		+44 d[d]	+15 d[f]
SSS on	+168 d[g]		
SSS off	+240 d to +340 d[g]		
White dwarf type	CO[g]	CO[h]	ONe [i]
White dwarf mass	$1.1\,M_\odot$[g]	$0.71\,M_\odot$[j]	

peculiar idiosyncrasies and timescales. Table 11.1, summarises the key details of three classical novae.

11.2 Observations

Since this chapter applies the results of Chapter 10 to data captured before the commissioning of the Global Jet Watch LHires instrument, it is predominantly based off Aquila data, detailed extenstively in previous chapters. The observational set-up and reduction pipeline for the data used in this paper are largely identical to the procedures described in Chapter 10 and McLoughlin et al. [2021a], so only a brief summary is provided here.

11.2.1 Observatories

The data in this chapter is based on three major observational campaigns carried out by the Global Jet Watch telescopes, a system of five terrestrial observatories separated around the world by longitude to enable continuous monitoring of rapidly evolving astrophysical targets. The observatories house 0.5 m telescopes, with most of the observations taken using the Aquila spectrographs which cover the wavelength range of approximately 5800 Å to 8400 Å with a resolution of $R \sim 4000$.

11.2.2 Major campaigns

The exact time-sampling for the three major campaigns discussed in this chapter differ qualitatively, and as such the insights gained from each varies accordingly. The

Table 11.2: Data set summary for the three major observational campaigns.

	V5668 Sgr	V906 Car	YZ Ret
First spectrum	2015-03-20	2018-03-25	2020-07-15
Last spectrum	2020-10-01	2020-12-28	2021-01-28
Discovery JD	2457097.134	2458193.819	2459038.670
Aquila spectra	2792	3166	5226
LHires spectra	1	5	36
Total exposures (h)	304	172	593

campaigns are summarised in Table 11.2.

V5668 Sgr erupted 6 years ago, and the Global Jet Watch has data spanning most of the intervening time. This makes it possible to trace different spectral components over broad changes, and allows quantitative comparisons between early spectra and later analogues. It was particularly fortunate that this classical nova erupted just after the system emerged from behind the sun, allowing for several months of consistent follow-up observations.

For V906 Car, the observational campaign was focused on dense time-sampling for the first three weeks, with somewhat sparser observations beyond this. The strength of this campaign is that it enables precise timing (to within hundreds of seconds) of certain early changes to the line profiles.

The coverage of YZ Reticuli was a hybrid of these two extremes, with a virtually uninterrupted sequence of nightly spectra captured across the four months between +50 d and +175 d. This effectively 'cracked the code', and revealed that a particular set of five Gaussian components (two pairs and an additional component) could satisfactorily explain the entire progression of spectral changes over several months [McLoughlin et al., 2021a]. This model was then used as a template to aid understanding of V5668 Sgr and V906 Car, and this chapter discusses the similarities and differences between these three examples of the varied field of objects that is classical novae.

11.2.3 Data reduction

The data reduction pipeline used is called ENDEAVOUR, and it performs the necessary dark subtraction, flat-fielding, wavelength calibration and heliocentric correction to produce one-dimensional output spectra from the input two-dimensional CCD data, as discussed at length in earlier chapters. The bespoke spectroscopic analysis tool called POIROT was then utilised to carry out fitting, following the same protocols as those detailed in section 2.6 of McLoughlin et al. [2021a]. The data are not photometrically calibrated, and so for visual consistency all the plots in this paper are normalised to unity.

11.3 Timing phase transitions for V906 Car

The existence of distinct phases in classical novae has been known for some time, as mentioned in Section 2.3. This section uses the intensive follow-up campaign for the first three weeks after V906 Car erupted to quantify the duration of transitions between these phases, shedding light on the underlying physical mechanisms driving the changes to the optical spectroscopy.

11.3.1 Timing the high-velocity absorption in V906 Car

Figure 11.2 shows the early evolution of the [O I] 7774 Å complex alongside the telluric-A band (included to show stability of the wavelength calibration). Time runs from bottom to top, and the whole series last only half a day. During this time, the principal absorption component at 7766 Å (300 $\mathrm{km\,s^{-1}}$) stays relatively steady, whilst a new and higher-velocity absorption complex appears at 7743 Å (1200 $\mathrm{km\,s^{-1}}$), shown with a fiducial vertical line. The last non-detection of the secondary absorption is at +12.21 d, and by +12.4 d there is a clear detection. The duration of this rapid phase transition can therefore be constrained to < 5 hours; Aydi et al. [2020a] had noticed such a transition for Hα and Hβ to have occurred in their spectra taken a day apart.

11.3.2 Timing the jets phase onset in V906 Car

While it has been well-understood for many years that classical novae can quickly transition from single to multiple absorption, these observations constitute the finest time-resolution yet for timing the transition.

Figure 11.3 shows the sudden onset of significantly increased emission in V906 Car. The left panel shows the spectra stacked with time increasing upwards, normalised from zero to one, while the right panel shows (with the corresponding colour for each spectrum) the depth of absorption and height of emission relative to that normalisation. Each vertical line in the right panel represents a single observation. The lower end of each vertical line represents the depth of the low-velocity absorption component, while the upper end is set to the height of the emission peak. This then shows that over the course of only half a day between +10.6 d and +11.3 d, there is

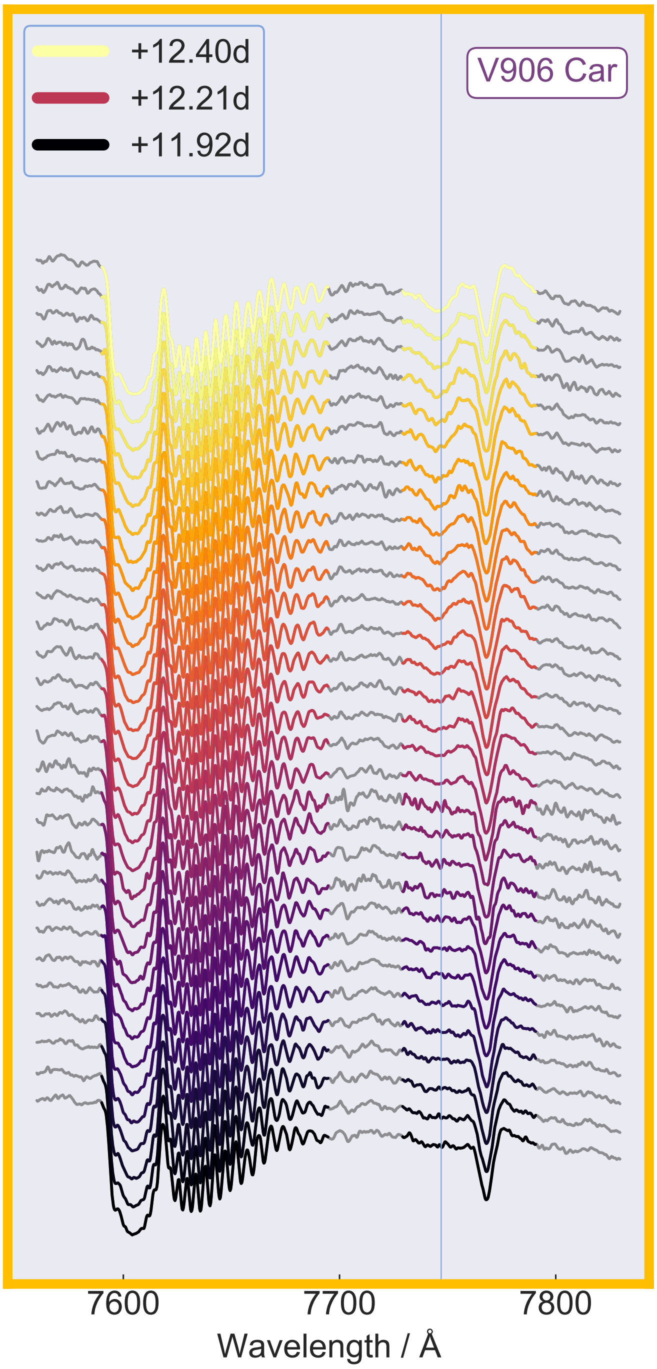

Figure 11.2: This series of spectra, with time increasing up the page, shows the development of a new high-velocity O7774 Å absorption line in V906 Car, with highly time-resolved sampling. The telluric-A band shown to the blue is included as an indicator of the robust wavelength calibration. The significant dates marked in the legend are the date of the last spectrum in the series, the date of the last spectrum with a non-detection of the high-velocity absorption component, and the date of the first spectrum in the series respectively.

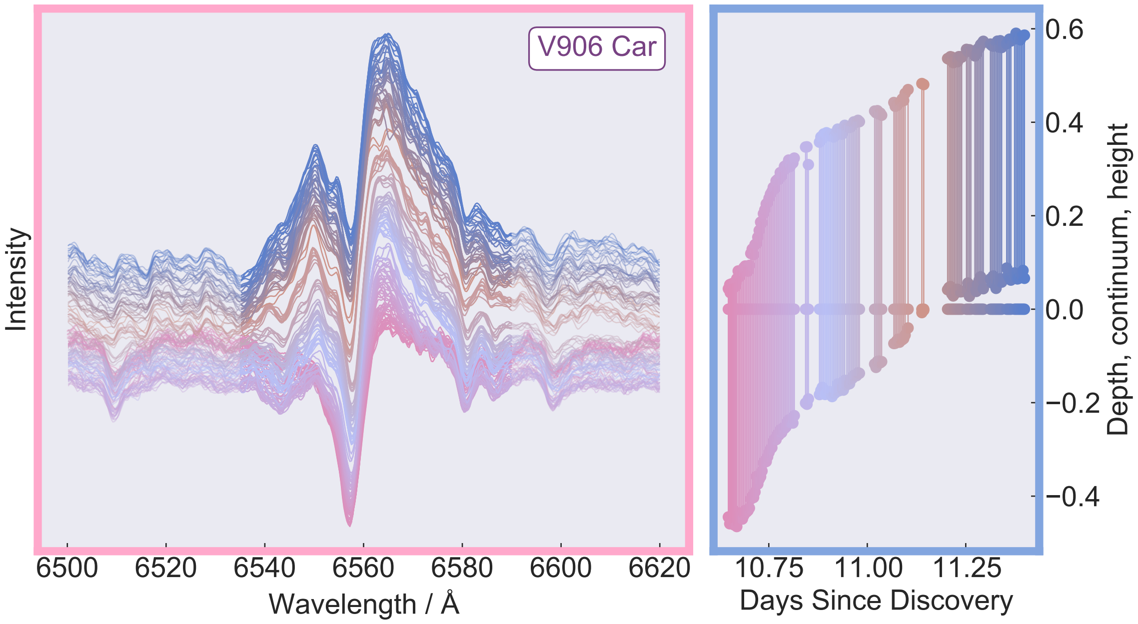

Figure 11.3: The left panel shows a stacked series of the Hα complex of V906 Car, ranging from +10.6 d to +11.3 d, with the earlier spectra shown at the bottom in pink and the later observations at the top in blue. Each vertical line in the right panel represents one of these observations; the lower end of each vertical line represents the depth of the low-velocity absorption component, while the upper end is set to the height of the emission peak, with the continuum subtracted and accordingly set to zero. The latest observation has the absorption dot above the background level, because by this time the emission is so strong and broad that even the deep low-velocity absorption trough does not descend below the continuum.

a transition from predominantly absorption to predominantly emission. At the end of this sequence, the low-velocity absorption component is still present, but the minimum is now above the continuum level, due to the immense emission at that same wavelength.

The emission is hard to fit well while P-Cygni absorption is present. With the diminishing of the absorption system by +30 d, an uncontaminated view of the emission profile is obtained, which can be fit by a clear jets and accretion disc pattern, as presented in Section 11.5.1.

There are two plausible explanations for this sudden onset of emission in the context of jets. One is that it takes a certain amount of time for the recently ejected material in the immediate vicinity of the eruption to be processed down the accretion disc and then launched as jets. In this scenario, one would expect to see the accretion disc emerging prior to the jets being formed; in the right-hand panel of Figure 11.3 there is a distinct change in the gradient of growth of emission occurring at +10.8 d

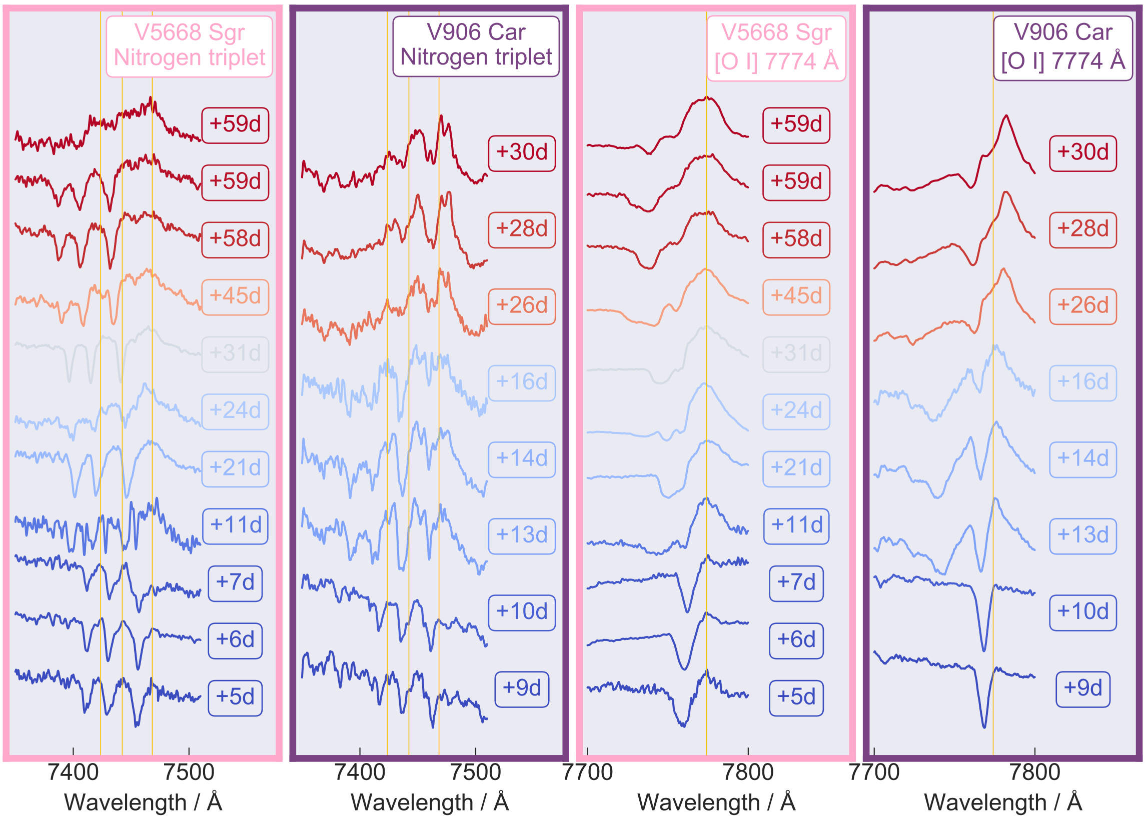

Figure 11.4: This figure compares the profiles of the nitrogen triplet (7423.641 Å, 7442.298 Å, and 7468.312 Å) and O7774 Å between the two classical novae, V5668 Sgr, and V906 Car. The left two panels show the nitrogen complex in the two novae, while the right two panels show the oxygen progression. For V5668 Sgr, two separate spectra taken hours apart at +59 d are plotted, which between them capture the disappearance of the nitrogen triplet — an event with no analogue in oxygen. The emission complex shown at the lower end of the wavelength range of V906 Car is not part of the oxygen complex, but a circumbinary disc signature, which was the subject of a previous paper [McLoughlin et al., 2020], and chapter 9.

which may mark the onset of jets.

The second possible conjecture is that the disc, fuelled by the recently ejected material, immediately forms jets, and that the signal presented here is the first time that these jets penetrate the shell-like ejecta and become spectrally significant.

11.3.3 Evolution of nitrogen and oxygen in V906 Car

The second panel in Figure 11.4 displays the evolution of the nitrogen triplet (7423.641 Å, 7442.298 Å, and 7468.312 Å) as it moves from absorption through to emission. From +13 d and beyond, hints of splitting in the peaks of some of these lines become apparent. Although this probes the limit of sensitivity given the signal-to-noise in this faint complex, separations of $\sim 200\,\mathrm{km\,s^{-1}}$ are suggested. These spectra also serve as a reference for the discussion of sudden changes in V5668 Sgr in Section 11.4.4.

The right panel of Figure 11.4 shows the progression of the O7774 Å complex. Perhaps more clearly than in the equivalent nitrogen panel, the oxygen low-velocity absorption line moves to increasingly blue-shifted wavelengths over time. The commencement of emission occurs between +10 d and +13 d, which mimics the behaviour detailed for the Hα complex in Section 11.3.2.

11.4 Timing phase transitions for V5668 Sgr

Figure 11.5 shows the progression of the Hα complex of V5668 Sgr through the several phases discussed in this chapter. The three distinct phases presently focused on each approximately span a year; in the following sections, each is considered in turn. This section comments on the rapidly evolving nitrogen triplet, contrasting it with that observed in V906 Car in Section 11.3.3.

11.4.1 V5668 Sgr: Initial jets and accretion disc

The fits in the first column of Figure 11.5 show a remarkably similar set of Gaussian components to those observed in YZ Reticuli in McLoughlin et al. [2021a]. These spectra of V5668 Sgr therefore also represent bipolar anti-parallel jets (shown in turquoise and pink), and an accretion disc (blue and red), with an additional component (orange) perhaps formed in a shell-like outflow. Initially, the additional component does not account for a significant proportion of the flux relative to the far stronger jets and disc signatures. However, by +113 d, the disc has begun to weaken in emission, and the jets decline slightly too.

By +154 d, two additional Gaussian components (other than the jets and disc) are required to fit the profiles well. Figure 11.5 shows +154 d twice: once with and once without an extra component to show its inclusion is necessary to obtain a close fit. Following the assumption that this extra flux most likely originates in the shell-like ejecta, this change is concordant with the time taken for it to become visible from behind the explosion, as the ejected material thins out and becomes transparent at this wavelength.

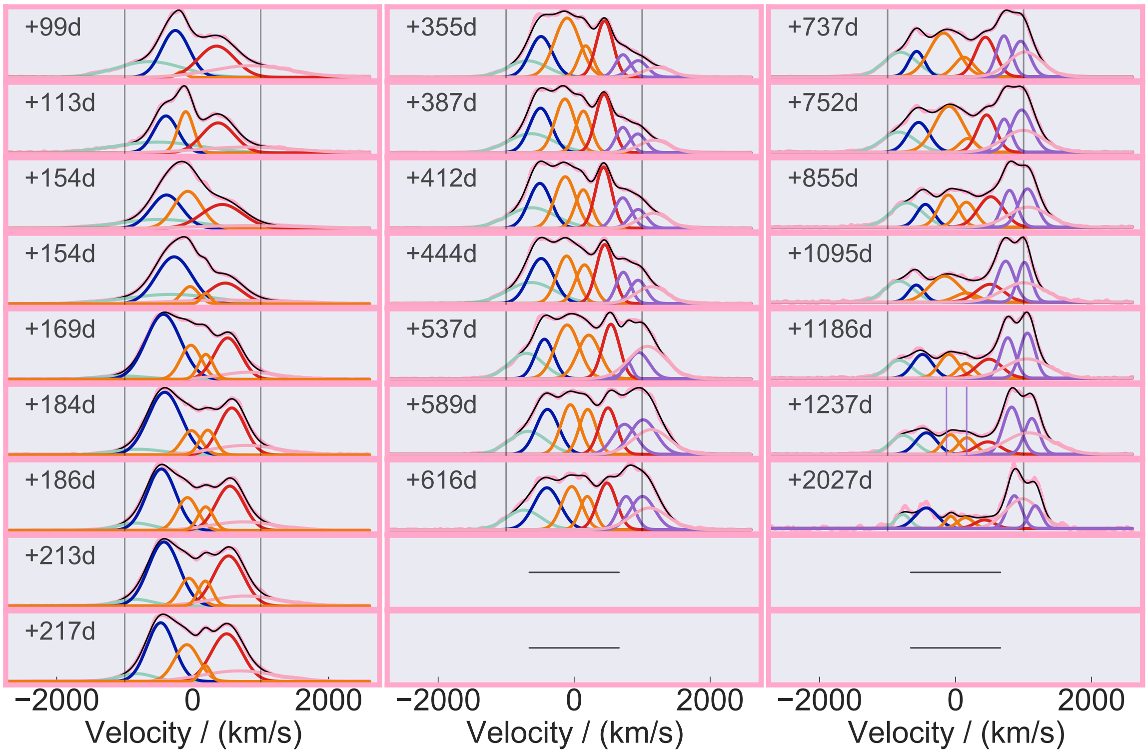

Figure 11.5: This figure shows an extensive sequence of Gaussian fits to the H-alpha complex of V5668 Sgr, from +99 d to +2027 d, by which time the spectrum is dominated by nebular nitrogen emission. The first column shows roughly the first year, when a simple jets and accretion disc model suffices to explain the profiles — the blue-shifted and red-shifted jet components are shown in turquoise and pink respectively, while the double-peaked signature of the accretion disc is represented by red and blue Gaussians. Following McLoughlin et al. [2021a], there is an additional orange component which may be associated with an ejected shell. Two separate models are presented for +154 d, since at this epoch it becomes necessary to include a second additional component to explain the data. The second column shows the classical nova during its transition to the nebular phase. Here, the same jets and accretion disc model remains valid, but requires two additional components (shown in purple) which are consistent with being due to [N II] 6584 Å. The third column shows a continuation of this pattern during the final years, with [N II] 6584 Å becoming the dominant emitter in this region. For the panel at +1237 d, two vertical lines (purple) to the plot to represent the radial velocities of the two nitrogen Gaussian components shown in purple, thus showing that these nebular lines are consistent with being from a dynamically similar parcel of gas as the shell (orange components) of H-alpha.

11.4.2 V5668 Sgr: Transition to nebular phase

The middle column of Figure 11.5 approximately spans the second year after eruption. During this time, additional lines appear to the red. The more recent spectra are used to get an accurate initial guess for the central wavelength of the extra components, revealing that they are consistent with being a double-peaked [N II] 6584 Å line. These spectra are well fit by the same set of six Gaussian components as those fitting the spectrum at +154 d, but with the addition of two nitrogen lines ([N II] 6584 Å). With the goal of using the Hα complex alone to determine the phase of a classical nova, this would provide a straightforward marker for the onset of the nebular phase. At this point, there is still comparable flux coming from Hα, and so the nebular phase does not have a sharp start time when the nebular lines dominate.

11.4.3 V5668 Sgr: Nebular phase dominant

The third column of Figure 11.5 shows the nebular phase beginning to dominate. By +750 d the nitrogen profile gives the strongest emission in this wavelength range, and by +1100 d nitrogen is totally dominant. Gehrz et al. [2018] identify the presence of the high-excitation forbidden coronal line of [Fe VII] 6085 Å. Figure 11.6 demonstrates the arrival and departure times of the [Fe VII] 6085 Å line; the [O I] 6300 Å appears throughout the existence and demise of the [Fe VII] 6085 Å line. The [Fe VII] 6085 Å emission is contemporaneous with the nitrogen lines shown in Figure 11.5 (in purple), which begin to emit strongly after approximately one year since detonation. Unlike these nitrogen lines, which establish themselves as the strongest line in the Global Jet Watch optical spectra after +800 d, the [Fe VII] 6085 Å line dwindles in comparison to Hα at late times.

11.4.4 V5668 Sgr: evolution of nitrogen and oxygen

The left panel of Figure 11.4 shows the progression of the nitrogen triplet (7423.641 Å, 7442.298 Å, and 7468.312 Å) in V5668 Sgr. What is fascinating about this is the suddenness of the disappearance of absorption lines between two spectra both taken on +59 d. Using time-resolved spectroscopy, an upper bound on the duration of this phase transition is set at $t <= 17$ h. The absorption was not present in any of the subsequent spectra.

There is a trend in the radial velocity of nitrogen absorption lines, which moves to higher speeds as time progresses from +7 d ($-480\,\mathrm{km\,s^{-1}}$) to +59 d ($-1440\,\mathrm{km\,s^{-1}}$). In the very early spectra taken before +7 d though, it decreases in radial speed, which

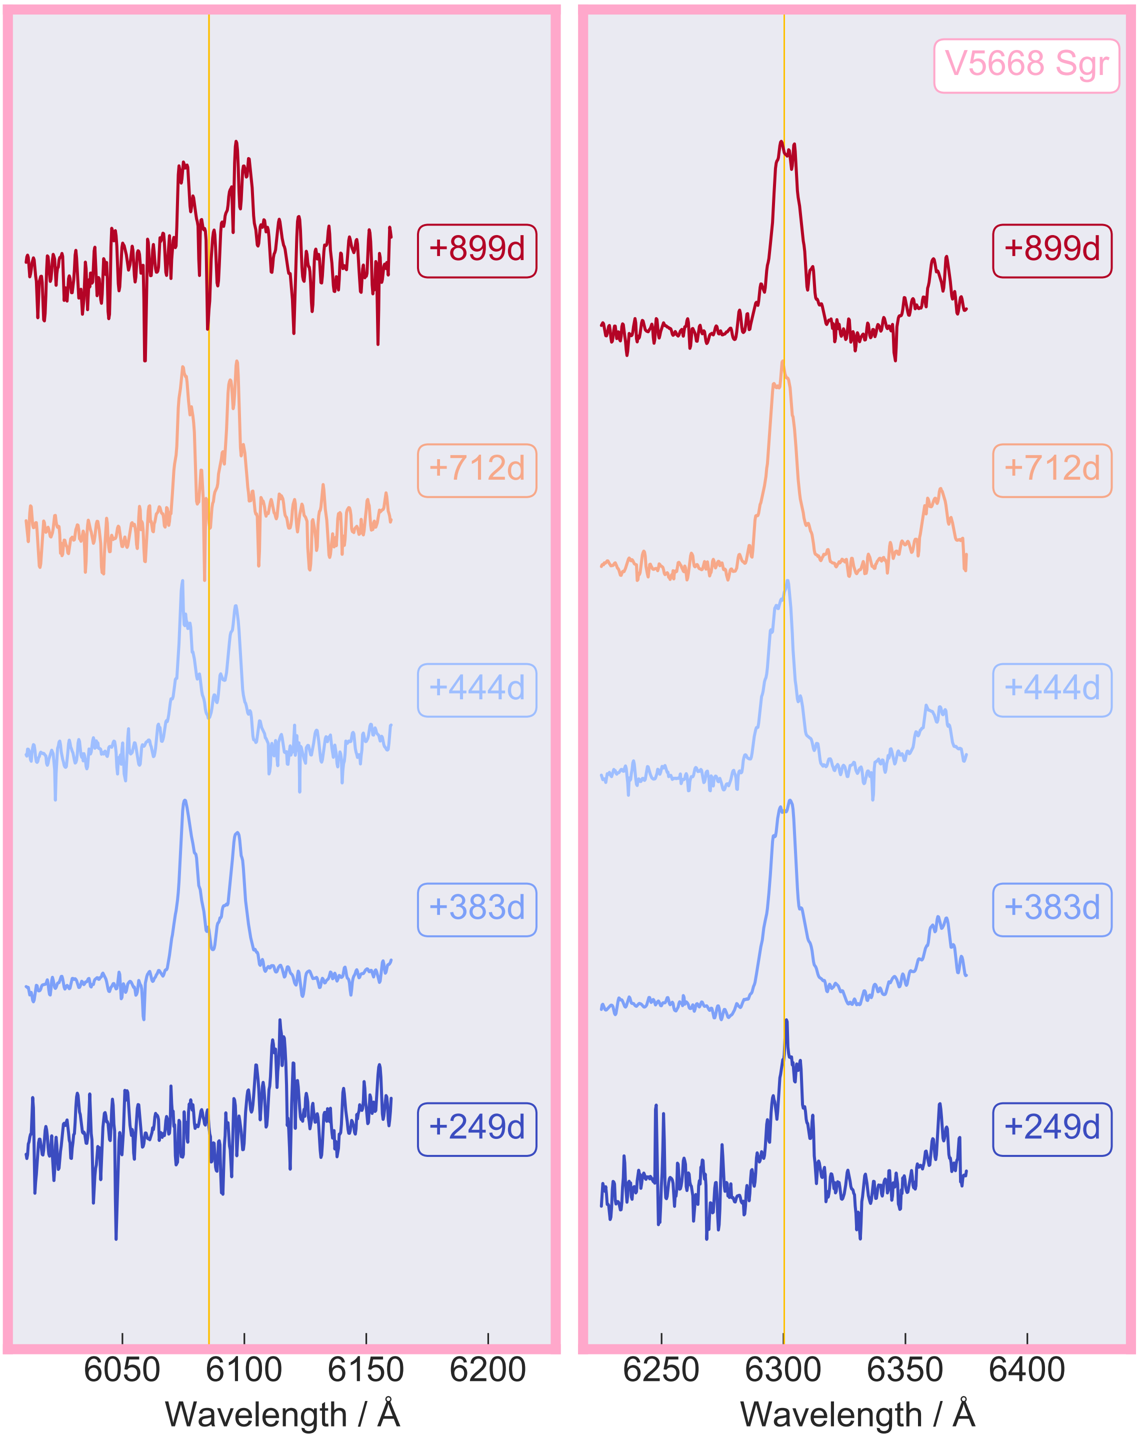

Figure 11.6: The left panel shows a double-peaked profile of the high-excitation coronal line of [Fe VII] 6085 Å in V5668 Sgr, from when it first appears in Global Jet Watch spectra around one year after eruption, to the last detection at +899 d. The right plot shows the [O I] 6300 Å line, which is present throughout this spectroscopic monitoring campaign of V5668 Sgr.

is discussed in more detail below. The spectrum taken at +11 d shows at least two distinct features, a narrow low-velocity component at 540 km s^{-1} and a broader high-velocity complex between 802 km s^{-1} and 1000 km s^{-1}. It is therefore likely that the dominant line in later spectra is related to the high-velocity component, rather than a drastically shifted low-velocity component.

The third panel of Figure 11.4 follows the O7774 Å line through a series of changes. It starts off at +5 d and +6 d with a single low-velocity P-Cygni absorption feature with a speed of 580 km s^{-1} along the line-of-sight. This then slows, reaching a minimum speed of 420 km s^{-1} on +7 d, before progressively increasing in blueshift up to a maximum of 730 km s^{-1} on +45 d. The pattern of slowing before continuing to advance to the blue, which is also present in the nitrogen complex, has previously been recorded in V906 Car by Aydi et al. [2020b], and is only apparent when spectral data are captured close to or before optical maximum. V5668 Sgr has many large-scale jitters in the early light curve (Figure 11.1), which is probably why this early 'retrograde' pattern is apparent.

From +11 d, the low-velocity component is joined by a high-velocity system, with a diffuse range of absorption velocities. The diffuse high-velocity absorption system of V5668 Sgr was the subject of a thorough study with high-resolution data [Tajitsu et al., 2016], which studied a snapshot of the Balmer lines of hydrogen and other metal lines, and measured the precise wavelengths of various narrow components within the complex.

11.5 Overview: A dynamical model for nova phases

Figure 11.7 shows the striking similarity between three seemingly very different eruptions. For each of these systems, the putative jet components are coloured in turquoise and pink, with the accretion disc in blue and red. The jet speeds differ substantially, although these observations cannot separate the effects of inclination angle from the actual rest frame jet-launch speed.

11.5.1 Onset of jets and accretion disc in V906 Car

The observational campaign for V906 Car focused keenly on the first three weeks, which made it possible to capture several key early transitions in the spectroscopic development of this classical nova. However, as shown in panels two and four of Figure 11.4, the emerging emission spectrum is interspersed with several highly variable absorption systems during the first month following eruption. Figure 11.8 therefore

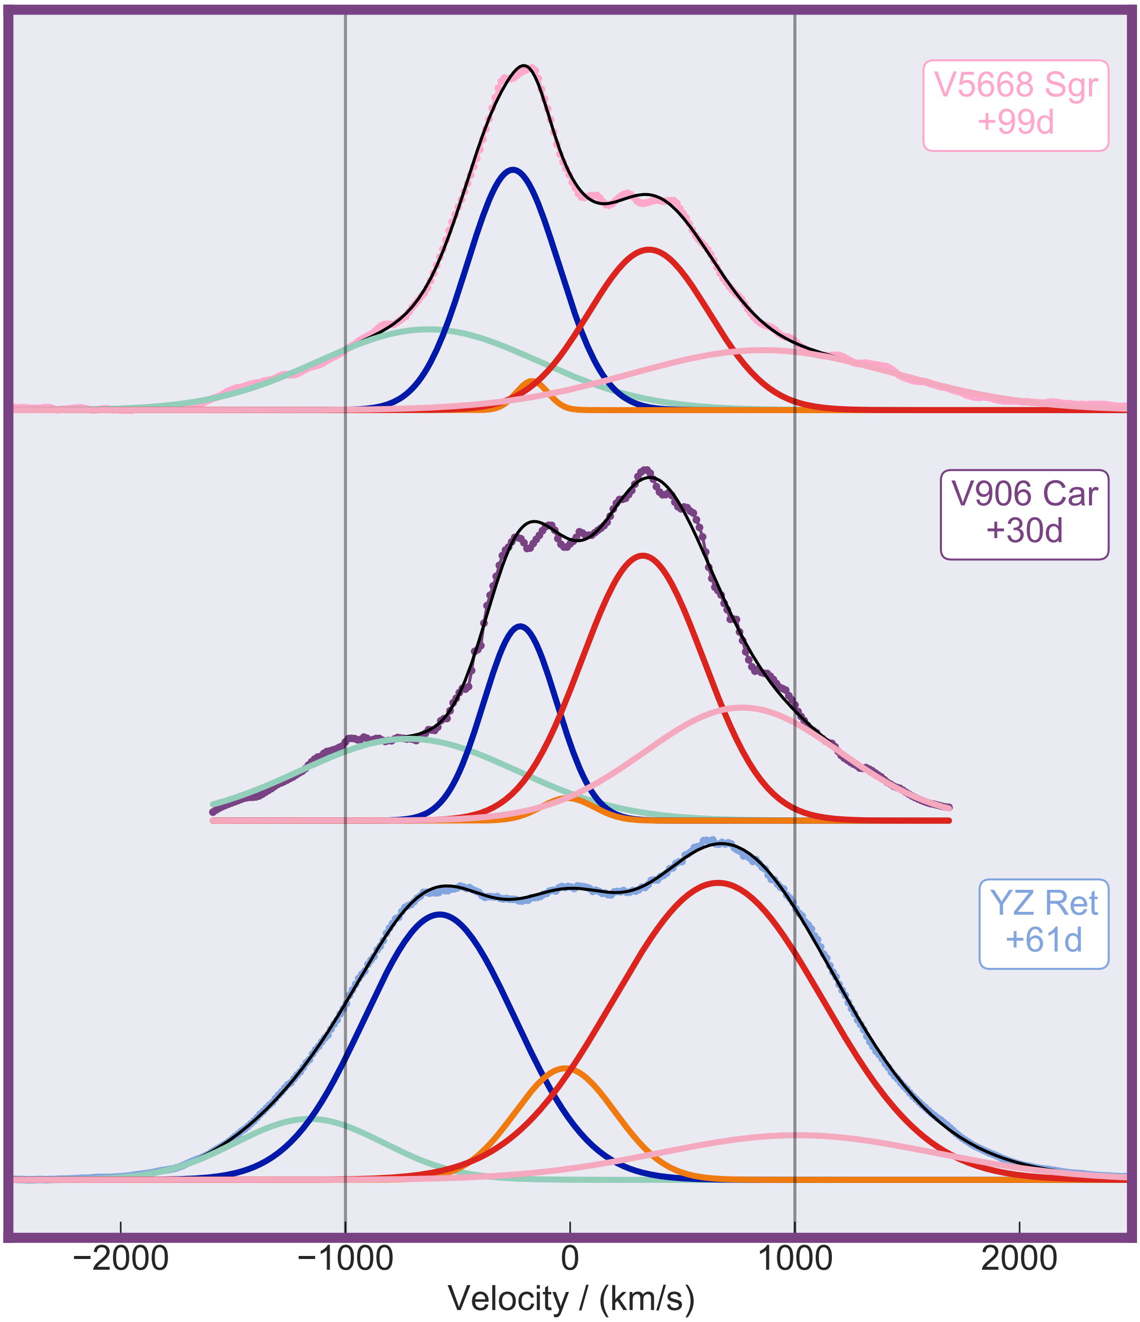

Figure 11.7: The five-component Gaussian model fits to the H-alpha complex of three separate and quite different classical novae. The top panel shows a spectrum of YZ Reticuli at +61 d, the middle is of V906 Car at +30 d, and the lower is of V5668 Sgr at +99 d. These classical novae had very different light curves (shown in Figure 11.1), yet all are well fit by two pairs of Gaussians and an additional component. The colour scheme and interpretation of lines follows Figure 11.5. There are two vertical fiducial lines, at ±1000 km s^{-1}.

begins fitting the emission complex for V906 Car (middle column) from +30 d, by which time the complex is well-approximated by the standard five emission Gaussians model previously identified as jets and accretion disc for YZ Reticuli and V6568 Sgr in McLoughlin et al. [2021a]. This fit remains accurate at least until +198 d, beyond which there is a two-year gap in observations and it becomes impossible to contiguously trace the evolution. Following Section 11.3.2 which discusses the onset of substantial emission in V906 Car, it is reasonable to assume that the emission complex that emerges suddenly around +10 d would also, in the absence of complex absorption systems, be well-fit by the same model as the one which remains in place from +30 d to +200 d. V906 Car therefore joins the growing sample of classical novae which display jets and an accretion disc in their Hα profiles, and it is most likely that such jets and an accretion disc are the source of the emission responsible for the initial rise in Hα flux from +10.6 d to +11.3 d.

11.5.2 An updated prescription for classical nova optical evolution

There is a small but growing sample of classical novae which show jets and an accretion disc in their Hα spectra, which are the two examples presented in this chapter (V906 Car, V5668 Sgr), together with the two in Chapter 10 [McLoughlin et al., 2021a] (YZ Reticuli and V6568 Sgr, in their figures 3 & 8 respectively). It is sensible to propose an update to the traditional physical model underpinned by optical observations. At first, there is a P-Cygni spectrum, with a relatively low velocity absorption component on the order of hundreds of $\mathrm{km\,s^{-1}}$ (the LVC), most likely associated with a slowly-ejected shell. This is promptly followed by the arrival of secondary absorption, with far higher radial velocities on the order of thousands of $\mathrm{km\,s^{-1}}$ (the HVC), which often consists of multiple separate shells. Very shortly after this, emission features originating from material at the base of jets and from the accretion disc emerge, taking only a few hours to push the minimum of the LVC absorption trough above the continuum level. The jets and the accretion disc appear to be precessing in YZ Reticuli, and similarly systematic follow-up of other classical novae may well yield similar results. At late times, nebular lines of forbidden transitions become strong, a process through which [N II] 6584 Å took approximately a year to compete with Hα in V5668 Sgr, and approximately another year to become the dominant source of emission in our optical spectra.

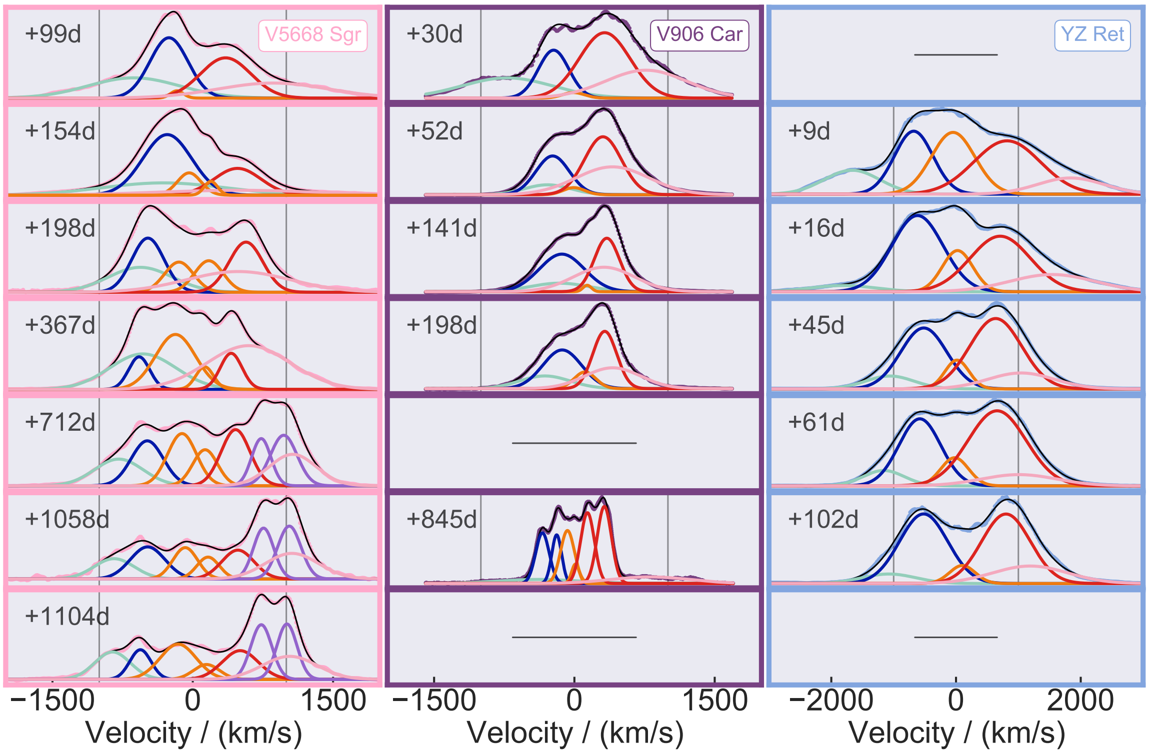

Figure 11.8: The jets and accretion disc model used to successfully describe the H-alpha evolution of classical nova YZ Reticuli [McLoughlin et al., 2021a] is extended to two additional examples — V5668 Sgr in the left column, and V906 Car in the central column. V5668 Sgr shows two additional Gaussian components (purple) which are due to [N II] 6584 Å. Since the campaign for V906 Car was primarily focused on the first three weeks after eruption during the multiple absorption component phase, the data set for this target has a substantial break before the final spectrum. Despite not being able to fully track the components through a time series, the +845 d spectrum of V906 Car shows hints of splitting occurring in the jet lines, coloured turquoise and pink. The right column shows fits for YZ Reticuli, as a comparison.

11.5.3 Jets, not clumps, responsible for H-alpha complexes

What is most remarkable about the fits shown in Figure 11.8 is how well this model accounts for variability in the line profiles. It is not warranted to claim that mere “clumpiness" can account for all this complexity. While the ejecta most likely does form clumps of some description, as evinced by imaging of classical nova remnants (such as the ALMA imaging of V5668 Sgr by Diaz et al. [2018], although this may arise in part because of sparse UV-plane sampling), this data set shows that it is not necessary to invoke clumps in order to explain the complex Hα profiles. Instead, a dynamical model consisting of jets, an accretion disc and a simple shell outflow is sufficient to explain the spectra from the early fading of absorption systems until the late nebular phase. Of course, the [N II] lines must be considered also, but they appear at late times, following the same dynamical prescription as the Hα shell, and as such are easily understood.

11.5.4 Implications for interstellar enrichment

Another interesting consequence of the developments to the jets model for nova outflow is the implication for the interstellar enrichment capabilities of classical novae. Considering the ejection to take place from near the surface of the white dwarf, where the escape velocity is approximately given by $V_{\rm esc}^2 \sim GM_\odot/R_{\rm WD}$, yields a value of around 4500 km s^{-1}, which is in excess of the ejection speeds observed in the majority of systems and thus implies that classical novae do not contribute significantly to the enrichment of the interstellar medium (ISM). However, if instead the launch radius is at the inner edge of an accretion disc, a quite different picture emerges. The standard prescription for a stellar accretion disc inner radius is that it terminates at the surface of the star. However, the jets of classical novae in this sample are launched while the white dwarf is still very hot and irradiance remains very high. It is conceivable that the disc inner edge is disrupted, and so the real radius of launch and therefore effective radius of the bottom of the gravitational potential well, may be significantly higher. It may then be quite feasible for jets of $\sim$2000 km s^{-1} such as those seen in YZ Reticuli or V6568 Sgr to propogate much further and hence play a role in enrichment of the ISM. In the case of V5668 Sgr, the jets appear to last in excess of 5 years; in this time, the initial ejecta will have reached a distance of at least 1000 AU, increasing the enthalpy and temperature of the ISM out to these radii as they do so. With $\sim$100 classical novae per year in the Milky Way, the cumulative effect of classical novae on the ISM is likely to be non-negligible.

11.5.5 Prediction for spatially resolved spectroscopy

The sprouting of the lines of [N II] 6584 Å (shown in purple in Figure 11.5) around at the anniversary of the explosion, and their resemblance to the central lines of Hα, is a reminder that detailed spectrally-resolved imaging of classical novae would likely reveal structures that could confirm or refute the jets model. In the case of the B[e] star MWC 137 studied in Mehner et al. [2016], an 8 m class telescope (VLT) together with a modern IFU (MUSE) revealed a slender jet in the [N II] 6584 Å line (their figure 8). With sensitive, high-angular resolution observations, it would be very instructive to see nitrogen-free Hα imaging, and to see which shells and which elongated outflows appear in lines other than hydrogen. There is however, a trade-off between angular size and brightness, which might be well balanced by a couple of months to years after a classical nova eruption.

11.6 Conclusions

The landscape of phases for the optical spectroscopic progression of classical novae has typically focused on relatively early-time absorption lines, as these elucidate what is happening specifically along the line-of-sight to the eruption, but a full understanding of the nature and evolution of the classical novae phenomena necessitates sustained time-resolved studies beyond this regime. Traditionally, the Hα emission complex has been deemed too contaminated to be straightforwardly fit and understood. The chapter extends the two jet-and-disc classical novae in McLoughlin et al. [2021a] to two further classical novae, showing that a jets and accretion disc model for understanding the Hα emission is perhaps more widely applicable to classical novae than was previously imagined. With the sample size growing, I propose this observation should be added to the standard sequence of events and phases in optical spectra and our way of understanding them. The variety of the classical novae now in the sample is encouraging, since it suggests that the jets and accretion disc model may be a feature of classical novae in general, rather than just restricted to one type.

Tight constraints on the timing of transitions between these classic phases were provided, lending support to the conjecture that jets and an accretion disc are key to understanding the optical emission of classical novae in the first thousand days.

I also suggested that thinking of the emission in terms of jets launched by an accretion disc rather than from the white dwarf surface could alter the relevant escape velocity, and may therefore allow for enhanced enrichment of the interstellar medium compared to what was previously thought possible from classical novae.

Beyond the spectroscopic evidence presented for these novae, independent support for this conjecture would come from a spatially-resolved spectroscopic image of these targets in a few years time, when the angular sizes would have grown sufficiently large to be well-matched to modern IFUs.

Chapter 12

Discussion

It is remarkable that the simple act of counting light reveals so much about nature. Of course, the light and resulting counts must be treated with the utmost care, but robotic instruments and automated computation together lend their assistance. This thesis respects the fact that the final fidelity and therefore value of scientific results rests upon the community's collective confidence in the processes involved in their creation, and takes pride in carefully detailing the meticulous measures that guarantee the integrity of Global Jet Watch spectra.

Careful inspection of the one-dimensional spectra resulted in the hitherto unobserved discovery of jets in the immediate aftermath of a classical nova eruption, and subsequent realisation that these are manifestly widely prevalent. Another key result was the first observation of a circumbinary disc, detected in emission after its shock-heating by the 2018 explosion of V906 Car. This final chapter sets out the findings of Chapters 9, 10 and 11 in the wider context of the field.

12.1 Circumbinary discs around classical novae

Chapter 9 presents the first direct observation of emission from pre-existing circumbinary material surrounding a classical nova (published as McLoughlin et al. [2020]). Its existence should come as no surprise, because of the well-established theoretical basis for the existence of families of stable circumbinary orbits. The circumbinary discs of classical novae have previously been postulated to exist but also to be inobservable due to their low luminosity [Spruit and Taam, 2001, Dubus et al., 2002]. However, these studies only considered systems in quiescence, not taking into account the energetic outbursts of classical novae. Careful fitting to the sustained time-resolved optical spectroscopy of classical nova V906 Car throughout the three weeks immediately following its 2018 eruption revealed that the disc had been shock-heated by the recent explosion, causing it to radiate line emission at detectable levels. This discovery of emission from a circumbinary disc lends strong support to the theory proposed by Williams et al. [2008] that such a disc may explain their observations of absorption systems in other classical novae.

This discovery ties together various other strands of other research. One is the theoretical possibility of circumbinary material, which has already been observed in a number of other systems, with several circumbinary exoplanets now known. Another is the white dwarf mass problem in classical novae; namely that in order to explain the observed mass distribution, systems must be losing angular momentum through a mechanism other than gravitational radiation, magnetic braking, mass loss and friction. The discovery of a circumbinary disc resolves this issue for V906 Car, perhaps suggesting that properly accounting for circumbinary discs might yield further advances in our understanding of the orbital mechanics of classical novae. Perhaps the most exciting idea is the possibility proposed by Dubus et al. [2002] and discussed by Williams [2011], that the presence of a circumbinary disc would drastically enhance the mass transfer rate, essentially acting as the trigger for classical nova eruptions.

It is not a coincidence that V906 Car revealed evidence of a circumbinary disc. As mentioned in Chapter 11, the observing campaigns of the different targets analysed during this thesis were substantially different (to some extent determined by Right Ascension and Declination, and the time of year of the eruption) despite all falling into the category of time-resolved spectroscopy. The Global Jet Watch spectra of V906 Car were front-loaded into the first three weeks after eruption, with the disc signature persisting for only seven days before vanishing. I conjecture that further similarly intensive and early observational campaigns for other classical novae would expose their circumbinary discs, which should now be strongly expected to exist. Indeed, the Global Jet Watch data set for V5668 Sgr, despite having a primary focus on capturing long-term trends, also contains some early spectra with hints of the same signatures.

A deep high-resolution observation of the circumbinary signal from various emission complexes in the spectrum of a classical nova would allow for more detailed modelling of the properties of the disc. In Chapter 9, there were hints of minor variations between different species emitting lines, and it would be interesting to analyse these in high-resolution spectra. A difficulty might lie in the rapid identification of these lines, as these may differ substantially based on the composition of the pre-existing circumbinary material, and perhaps therefore the nature of the white dwarf. Since the circumbinary signal is so transient, disappearing a week after first appearing, spectra would have to be observed, then processed to determine which lines are present, and then observed again with a high resolution instrument. The Global Jet Watch observatories, now equipped with LHires and Aquila spectrographs, is well-placed to carry out such additional high-resolution observations.

12.2 Classical nova jets

Chapter 10 presents spectroscopic evidence of jets and an accretion disc in the classical nova YZ Reticuli, and shows the same model also fits V6568 Sgr data. Astrophysical jets arising at a broad range of scales in the universe, although crucially only in the presence of ongoing accretion. While it had been posited that a recent classical nova eruption might totally disrupt the accretion disc, Retter [2006] found that this was only the case in magnetic systems or intermediate polars, which make up $< 15\%$ of the classical nova population. Indeed, Hachisu and Kato [2012] show that irradiation of the accretion disc by the hot white dwarf is a plausible explanation for the plateau phase in some novae. Given the probable survival of an accretion disc, and the recent ejection of a large quantity of hydrogen by the eruption, enhanced mass-transfer through the accretion disc is likely to occur, creating circumstances favourable for the launching of jets. Since this set of conditions commonly exists for classical novae, jets may be found in the spectra signature of many classical novae.

Chapter 11 takes over where Chapter 10 leaves off, bringing the sample of classical novae which show jets and an accretion disc to a total of four. The fact that all four of the classical novae observed thus far by the Global Jet Watch had an Hα complex that was well-fit by this model (after the disappearance of absorption systems) emphatically supports the consensus reached in 11, that this phenomenon is likely a very common feature of classical novae. Indeed, model-dependent fitting conducted using the SHAPE algorithm has hinted that classical novae might have bipolar morphologies (e.g. Ribeiro et al. [2013]).

Milli-arcsecond radio jets [Sokoloski et al., 2008], and collimated Hα jets [Skopal, 2015], have been observed in recurrent novae and this has been posited to be because of the higher white dwarf masses in these systems. However, it is also perhaps due to the relative ease with which they can be observed — repeated eruptions enable the preparation of detailed observational plans. It would be very interesting to confirm the relationship between white dwarf mass and jet speed, and to do this would require knowledge of one additional observable: the inclination angle. Polarimetric observations such as those performed by Harvey et al. [2018] may provide a useful way of unlocking this third variable, which would allow calculation of the jet speed in the rest frame of the eruption.

It would be fruitful indeed to capture spatially resolved spectroscopy of a classical nova showing jets over a wavelength range centred about Hα and sufficiently wide to capture the fast jets, as this would reveal the opening angle, presence or absence of knots along the jet axis (taking Mehner et al. [2016] as an example of this in a non-nova system), and potentially an imprint of precession. To optimally plan such an observation would be a challenge since classical nova systems fade as they expand to spatially resolvable scales, in ways that depend on the particular nature and orientation of the target. However, the results shown in Chapter 11 for V5668 Sgr indicate that a good time might be between roughly +100 d and +1000 d; late enough that it has had time to become spatially resolved, but early enough that Hα jet emission is still bright. A modern IFU would easily be capable of resolving some interesting dynamics. Since there were hints of this also in [O III] lines (4959 Å, 5007 Å) for YZ Reticuli, that spectral range may also yield interesting results.

12.3 Closing remarks

With an estimated ∼ 50 Galactic nova eruptions per year [Shafter, 2017], and ever-improving observational data sets, it is unsurprising that the field of classical novae continues to yield exciting new results. This thesis has presented novel insights derived from densely time-sampled optical spectroscopy, and I am confident that there are many more and varied qualities of classical novae waiting to be discovered through this and similar methods.

References

A. Carr, K. Said, and T. M. Davis. ATel #13874: WiFeS follow-up observations of the naked-eye nova associated to MGAB-V207. *The Astronomer's Telegram*, 13874 (1):1–1, 2020. URL http://www.astronomerstelegram.org/?read=13874.

M Ackermann, M Ajello, A Albert, L Baldini, J Ballet, G Barbiellini, D Bastieri, R Bellazzini, E Bissaldi, R. D. Blandford, E. D. Bloom, E Bottacini, T. J. Brandt, J Bregeon, P Bruel, R Buehler, S Buson, G. A. Caliandro, R. A. Cameron, M Caragiulo, P. A. Caraveo, E Cavazzuti, E Charles, A Chekhtman, C. C. Cheung, J Chiang, G Chiaro, S Ciprini, R Claus, J Cohen-Tanugi, J Conrad, S Corbel, F D'Ammando, A de Angelis, P. R. den Hartog, F de Palma, C. D. Dermer, R Desiante, S. W. Digel, L Di Venere, E do Couto e Silva, D Donato, P. S. Drell, A Drlica-Wagner, C Favuzzi, E. C. Ferrara, W. B. Focke, A Franckowiak, L Fuhrmann, Y Fukazawa, P Fusco, F Gargano, D Gasparrini, S Germani, N Giglietto, F Giordano, M Giroletti, T Glanzman, G Godfrey, I. A. Grenier, J. E. Grove, S Guiriec, D Hadasch, A. K. Harding, M Hayashida, E Hays, J. W. Hewitt, A. B. Hill, X Hou, P Jean, T Jogler, G Jóhannesson, A. S. Johnson, W. N. Johnson, M Kerr, J Knödlseder, M Kuss, S Larsson, L Latronico, M Lemoine-Goumard, F Longo, F Loparco, B Lott, M. N. Lovellette, P Lubrano, A Manfreda, P Martin, F Massaro, M Mayer, M. N. Mazziotta, J. E. McEnery, P. F. Michelson, W Mitthumsiri, T Mizuno, M. E. Monzani, A Morselli, I. V. Moskalenko, S Murgia, R Nemmen, E Nuss, T Ohsugi, N Omodei, M Orienti, E Orlando, J. F. Ormes, D Paneque, J. H. Panetta, J. S. Perkins, M Pesce-Rollins, F Piron, G Pivato, T. A. Porter, S Rainò, R Rando, M Razzano, S Razzaque, A Reimer, O Reimer, T Reposeur, P. M. Saz Parkinson, M Schaal, A Schulz, C Sgrò, E. J. Siskind, G Spandre, P Spinelli, Ł Stawarz, D. J. Suson, H Takahashi, T Tanaka, J. G. Thayer, J. B. Thayer, D. J. Thompson, L Tibaldo, M Tinivella, D. F. Torres, G Tosti, E Troja, Y Uchiyama, G Vianello, B. L. Winer, M. T. Wolff, D. L. Wood, K. S. Wood, M Wood, S Charbonnel, R. H. D. Corbet, I De Gennaro Aquino, J. P. Edlin, E Mason, G. J. Schwarz, S. N. Shore, S Starrfield, F Teyssier, and Fermi-LAT

Collaboration. Fermi establishes classical novae as a distinct class of gamma-ray sources. *Science*, 345(6196):554–558, 8 2014. doi: 10.1126/science.1253947.

Sachiko Amari, Xia Gao, Larry R. Nittler, Ernst Zinner, Jordi Jose, Margarita Hernanz, and Roy S. Lewis. Presolar Grains from Novae. *The Astrophysical Journal*, 551(2):1065–1072, 4 2001. ISSN 0004-637X. doi: 10.1086/320235. URL https://iopscience.iop.org/article/10.1086/320235.

A Arai, H Kawakita, Y Shinnaka, and A Tajitsu. Origins of Absorption Systems of Classical Nova V2659 Cyg (Nova Cyg 2014). *The Astrophysical Journal*, 830 (30):–, 2016. ISSN 15384357. doi: 10.3847/0004-637X/830/1/30. URL http://iopscience.iop.org/article/10.3847/0004-637X/830/1/30/pdf.

E. Aydi, L. Chomiuk, L. Izzo, E. J. Harvey, J. Leahy-McGregor, J. Strader, D. A. H. Buckley, K. V. Sokolovsky, A. Kawash, C. S. Kochanek, J. D. Linford, B. D. Metzger, K. Mukai, M. Orio, B. J. Shappee, L. Shishkovsky, E. Steinberg, S. J. Swihart, J. L. Sokoloski, F. M. Walter, and P. A. Woudt. Early spectral evolution of classical novae: consistent evidence for multiple distinct outflows. *The Astrophysical Journal*, 905:62–94, 10 2020a. ISSN 15384357. doi: 10.3847/1538-4357/abc3bb. URL http://arxiv.org/abs/2010.07481.

Elias Aydi. ATel #13867: SALT spectroscopic classification of MGAB-V207 as a classical nova, 2020a. URL http://www.astronomerstelegram.org/?read=13867.

Elias Aydi. ATel #13872: SALT spectroscopic classification of PNV J17580848-3005376 as a very fast classical nova, 2020b. URL http://www.astronomerstelegram.org/?read=13872.

Elias Aydi, L Chomiuk, J Strader, S J Swihart, A Bahramian, E J Harvey, C T Britt, D A H Buckley, Ping Chen, K Dage, M J Darnley, S Dong, F-J Hambsch, S Holoien, S W Jha, C S Kochanek, N P M Kuin, K L Li, L A G Monard, K Mukai, K L Page, J L Prieto, N D Richardson, B J Shappee, L Shishkovsky, K V Sokolovsky, K Z Stanek, T Thompson, and Elias Aydi. Flaring, dust formation and shocks in the very slow nova ASASSN-17pf (LMCN 2017-11a). Technical report, Michigan State University, 2019. URL https://asas-sn.osu.edu/.

Elias Aydi, Kirill V . Sokolovsky, Laura Chomiuk, Elad Steinberg, Kwan Lok Li, Indrek Vurm, Brian D. Metzger, Jay Strader, Koji Mukai, Ondřej Pejcha, Ken J. Shen, Gregg A. Wade, Rainer Kuschnig, Anthony F. J. Moffat, Herbert Pablo, Andrzej Pigulski, Adam Popowicz, Werner Weiss, Konstanze Zwintz, Luca Izzo, Karen R. Pollard, Gerald Handler, Stuart D. Ryder, Miroslav D. Filipović, Rami Z. E. Alsaberi, Perica Manojlović, Raimundo Lopes de Oliveira, Frederick M. Walter, Patrick J. Vallely, David A. H. Buckley, Michael J. I. Brown, Eamonn J. Harvey, Adam Kawash, Alexei Kniazev, Christopher S. Kochanek, Justin Linford, Joanna Mikolajewska, Paolo Molaro, Marina Orio, Kim L. Page, Benjamin J. Shappee, and Jennifer L. Sokoloski. Direct evidence for shock-powered optical emission in a nova. *Nature Astronomy*, 4:776–780, 4 2020b. ISSN 2397-3366. doi: 10.1038/s41550-020-1070-y. URL `http://www.nature.com/articles/s41550-020-1070-y`.

C A L Bailer-Jones, J Rybizki, M Fouesneau, G Mantelet, and R Andrae. Estimating Distance from Parallaxes. IV. Distances to 1.33 Billion Stars in Gaia Data Release 2. *The Astronomical Journal*, 156:58, 2018. doi: 10.3847/1538-3881/aacb21. URL `http://gaia.ari.uni-heidelberg.de/tap.html`.

D. P. K. Banerjee, Mudit K. Srivastava, N. M. Ashok, and V. Venkataraman. Near Infrared studies of the carbon-monoxide and dust forming nova V5668 Sgr. *Monthly Notices of the Royal Astronomical Society*, 455, 10 2015. ISSN 0035-8711. doi: 10.1093/mnrasl/slv163.

C. S. Beals. The Spectra of the P Cygni Stars. *Publications of the Dominion Astrophysical Observatory Victoria*, 9(1):1, 1 1953.

Katherine M Blundell, Michael G Bowler, and Linda Schmidtobreick. SS 433: Observation of the circumbinary disk and extraction of the system mass. *The Astrophysical Journal*, 678:47–50, 2008. doi: 10.1086/588027. URL `https://iopscience.iop.org/article/10.1086/588027`.

Katherine M. Blundell, Robert Laing, Steven Lee, and Anita Richards. SS433's jet trace from alma imaging and global jet watch spectroscopy: Evidence for post-launch particle acceleration. *The Astrophysical Journal Letters*, 867:25–29, 11 2018. ISSN 2041-8205. doi: 10.3847/2041-8213/aae890.

Michael F Bode. The Evolution of Nova Ejecta. In M Meixner, J Kastner, and N. Soker, editors, *Asymmetric Planetary Nebulae III*, volume XXX. ASP Conference Series, 2004.

Michael F. Bode and Aneurin Evans. Classical Novae. *Classical Novae, 2nd Edition. Edited by M.F. Bode and A. Evans. Cambridge Astrophysics Series, No. 43, Cambridge: Cambridge University Press, 2008.*, 43, 2008. URL http://adsabs.harvard.edu/abs/2008clno.book.....B.

Alan H Bridle and Richard A Perley. Extragalactic Radio Jets. *Annual Review of Astronomy and Astrophysics*, 22(1):319–358, 1984. doi: 10.1146/annurev.aa.22.090184.001535. URL https://doi.org/10.1146/annurev.aa.22.090184.001535.

A G A Brown, A Vallenari, T Prusti, J H J de Bruijne, C Babusiaux, C A L Bailer-Jones, M Biermann, D W Evans, L Eyer, F Jansen, C Jordi, S A Klioner, U Lammers, L Lindegren, X Luri, F Mignard, C Panem, D Pourbaix, S Randich, P Sartoretti, H I Siddiqui, C Soubiran, F van Leeuwen, N A Walton, F Arenou, U Bastian, M Cropper, R Drimmel, D Katz, M G Lattanzi, J Bakker, C Cacciari, J Castañeda, L Chaoul, N Cheek, F De Angeli, C Fabricius, R Guerra, B Holl, E Masana, R Messineo, N Mowlavi, K Nienartowicz, P Panuzzo, C Barata, D Barbato, F Barblan, P S Barklem, D Barrado, M Barros, M A Barstow, S Bartholomé Muñoz, J-l Bassilana, U Becciani, M Bellazzini, A Berihuete, S Bertone, L Bianchi, O Bienaymé, S Blanco-Cuaresma, T Boch, C Boeche, A Bombrun, R Borrachero, D Bossini, S Bouquillon, G Bourda, A Bragaglia, L Bramante, M A Breddels, A Bressan, N Brouillet, T Brüsemeister, E Brugaletta, B Bucciarelli, A Burlacu, D Busonero, A G Butkevich, R Buzzi, E Caffau, R Cancelliere, G Cannizzaro, T Cantat-Gaudin, R Carballo, T Carlucci, J M Carrasco, L Casamiquela, M Castellani, A Castro-Ginard, P Charlot, L Chemin, A Chiavassa, G Cocozza, G Costigan, S Cowell, F Crifo, M Crosta, C Crowley, J Cuypers, C Dafonte, Y Damerdji, A Dapergolas, P David, M David, P de Laverny, F De Luise, R De March, D de Martino, R de Souza, A de Torres, J Debosscher, E del Pozo, M Delbo, A Delgado, H E Delgado, P Di Matteo, S Diakite, C Diener, E Distefano, C Dolding, P Drazinos, J Durán, B Edvardsson, H Enke, K Eriksson, P Esquej, G Eynard Bontemps, C Fabre, M Fabrizio, S Faigler, A J Falcão, M Farràs Casas, L Federici, G Fedorets, P Fernique, F Figueras, F Filippi, K Findeisen, A Fonti, E Fraile, M Fraser, E Licata, H E P Lindstrøm, T A Lister, E Livanou, A Lobel, M López, S Managau, R G Mann, G Mantelet, O Marchal, J M Marchant, M Marconi, S Marinoni, G Marschalkó, D J Marshall, M Martino, G Marton, N Mary, D Massari, C Ordénovic, D Ordóñez-Blanco, P Osborne, C Pagani, I Pagano, F Pailler, H Palacin, L Palaversa, A Panahi, M Pawlak, A M Piersimoni, F-x Pineau, E Plachy, G Plum, E Poggio, E Poujoulet, A Prša, L Pulone, E Racero, S Ragaini, N Rambaux, M Ramos-Lerate, S Regibo, C Reylé, F Riclet, V Ripepi, A Riva, A Rivard, G Rixon, T Roegiers, M Roelens, M Romero-Gómez, N Rowell, F Royer, L Ruiz-Dern, G Sadowski, T Sagristà Sellés, J Sahlmann, J Salgado, E Salguero, N Sanna, T Santana-Ros, M Sarasso, H Savietto, M Schultheis, E Sciacca, M Segol, J C Segovia, D Ségransan, I-c Shih, L Siltala, A F Silva, R L Smart, K W Smith, E Solano, F Solitro, R Sordo, S Soria Nieto, J Souchay, A Spagna, F Spoto, U Stampa, I A Steele, H Steidelmüller, C A Stephenson, H Stoev, F F Suess, J Surdej, L Szabados, E Szegedi-Elek, D Tapiador, F Taris, G Tauran, M B Taylor, R Teixeira, D Terrett, P Teyssandier, W Thuillot, A Titarenko, F Torra Clotet, C Turon, A Ulla, E Utrilla, S Uzzi, M Vaillant, G Valentini, V Valette, A van Elteren, E Van Hemelryck, M van Leeuwen, M Vaschetto, A Vecchiato, J Veljanoski, Y Viala, D Vicente, S Vogt, C von Essen, and H Voss. Gaia Data Release 2. *Astronomy & Astrophysics*, 616(A1), 2018. doi: 10.1051/0004-6361/201833051. URL https://doi.org/10.1051/0004-6361/201833051.

M. R. Calabretta and E. W. Greisen. Representations of celestial coordinates in FITS. *Astronomy & Astrophysics*, 395(3):1077–1122, 12 2002. ISSN 0004-6361. doi: 10.1051/0004-6361:20021327. URL http://www.aanda.org/10.1051/0004-6361:20021327.

Phil Charles, Will Clarkson, Remon Cornelisse, and Chun Shih I. Observational properties of accretion discs: Precessing and warped? *New Astronomy Reviews*, 51 (10-12):768–774, 5 2008. ISSN 13876473. doi: 10.1016/j.newar.2008.03.025.

O. Chesneau, E. Lagadec, M. Otulakowska-Hypka, D. P.K. Banerjee, C. E. Woodward, E. Harvey, A. Spang, P. Kervella, F. Millour, N. Nardetto, N. M. Ashok, M. J. Barlow, M. Bode, A. Evans, D. K. Lynch, T. J. O'Brien, R. J. Rudy, and R. W. Russell. The expanding dusty bipolar nebula around the nova V1280 Scorpi. *Astronomy and Astrophysics*, 545, 2012. ISSN 00046361. doi: 10.1051/0004-6361/201219825.

H. Corbett, N. Law, E. Goeke, J. Ratzloff, W. Howard, O. Fors, D. del Ser, and R. M. Quimby. Pre-Discovery Detection of ASASSN-18fv by Evryscope. *The Astronomer's Telegram, No. 11467*, 11467, 2018. URL http://adsabs.harvard.edu/abs/2018ATel11467....1.

B. Csák, L. L. Kiss, A. Retter, A. Jacob, and S. Kaspi. Spectroscopic monitoring of the transition phase in nova V4745 Sgr. *Astronomy and Astrophysics*, 429(2): 599–605, 1 2005. ISSN 00046361. doi: 10.1051/0004-6361:20035751.

J A Curcio, L F Drummeter, and G L Knestrick. An Atlas of the Absorption Spectrum of the Lower Atmosphere from 5400A to 8520A. *Applied Optics*, 3(12):1401–1410, 12 1964.

M. J. Darnley, R. Hounsell, P. Godon, D. A. Perley, M. Henze, N. P. M. Kuin, B. F. Williams, S. C. Williams, M. F. Bode, D. J. Harman, K. Hornoch, M. Link, J.-U. Ness, V. A. R. M. Ribeiro, E. M. Sion, A. W. Shafter, and M. M. Shara. No Neon, but Jets in the Remarkable Recurrent Nova M31N 2008-12a? — Hubble Space Telescope Spectroscopy of the 2015 Eruption. *The Astrophysical Journal*, 847(1):35, 9 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/aa8867. URL https://iopscience.iop.org/article/10.3847/1538-4357/aa8867/pdf.

Massimo Della Valle and Luca Izzo. Observations of galactic and extragalactic novae. *Astronomy and Astrophysics Review*, 28(3):1–118, 12 2020. ISSN 09354956. doi: 10.1007/s00159-020-0124-6.

Marcos P Diaz, Zulema Abraham, Valério A R M Ribeiro, Pedro P B Beaklini, and Larissa Takeda. The structure of a recent nova shell as observed by ALMA. *Monthly Notices of the Royal Astronomical Society*, 480:54–57, 2018. doi: 10.1093/mnrasl/sly121. URL https://academic.oup.com/mnrasl/article-abstract/480/1/L54/5047315.

Samuel Doolin and Katherine M. Blundell. The precession of SS433's radio ruff on long timescales. *The Astrophysical Journal*, 698(1):L23–L26, 6 2009. ISSN 0004-637X. doi: 10.1088/0004-637X/698/1/L23. URL https://iopscience.iop.org/article/10.1088/0004-637X/698/1/L23.

Samuel Doolin and Katherine M Blundell. The dynamics and stability of circumbinary orbits. *Mon. Not. R. Astron. Soc*, 418:2656–2668, 2011. doi: 10.1111/j.1365-2966.2011.19657.x. URL https://academic.oup.com/mnras/article-abstract/418/4/2656/1029021.

Ronald A. Downes, Ronald F. Webbink, Michael M. Shara, Hans Ritter, Ulrich Kolb, and Hilmar W. Duerbeck. A Catalog and Atlas of Cataclysmic Variables: The Living Edition. *Publications of the Astronomical Society of the Pacific*, 113(784):764–768, 6 2001. ISSN 0004-6280. doi: 10.1086/320802. URL http://iopscience.iop.org/article/10.1086/320802.

Guillaume Dubus, Ronald E Taam, and H C Spruit. The structure and evolution of circumbinary disks in cataclysmic variable systems. *The Astrophysical Journal*, 569:395–404, 4 2002.

Moshe Elitzur, Andrés Asensio Ramos, and Cecilia Ceccarelli. Rotating discs and non-kinematic double peaks. *Mon. Not. R. Astron. Soc*, 422:1394–1402, 2012. doi: 10.1111/j.1365-2966.2012.20711.x. URL `https://academic.oup.com/mnras/article-abstract/422/2/1394/1036123`.

Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc. 2nd Int. Conf. Knowl. Discov. Data Min.*, pages 226–231. AAAI Press, 1996.

Joana Figueira, Jordi José, Enrique García-Berro, Simon W Campbell, Domingo García-Senz, and Shazrene Mohamed. Astronomy Three-dimensional simulations of the interaction between the nova ejecta, accretion disk, and companion star. *Astronomy & Astrophysics*, 613(8):8, 2018. doi: 10.1051/0004-6361/201731545. URL `https://doi.org/10.1051/0004-6361/201731545`.

Martin A Fischler and Robert C Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. In Martin A Fischler and Oscar Firschein, editors, *Readings in Computer Vision*, pages 726–740. Morgan Kaufmann, San Francisco (CA), 1987. ISBN 978-0-08-051581-6. doi: https://doi.org/10.1016/B978-0-08-051581-6.50070-2. URL `https://www.sciencedirect.com/science/article/pii/B9780080515816500702`.

Daniel (Center for Cosmology Foreman-Mackey, New Particle Physics, Department of Physics, York University), David W Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 2013. URL `http://dan.iel.fm/emcee.`

I Fuentes-Morales, C Tappert, M Zorotovic, N Vogt, E C Puebla, M R Schreiber, A Ederoclite, and L Schmidtobreick. Life after eruption VIII: The orbital periods of novae. *Monthly Notices of the Royal Astronomical Society*, 501(4):6083–6102, 1 2021. ISSN 0035-8711. doi: 10.1093/mnras/staa3482. URL `https://academic.oup.com/mnras/article/501/4/6083/5974297`.

Cecilia Helena Payne Gaposchkin. The Galactic Novae. *Amsterdam, North-Holland Pub. Co.; New York, Interscience Publishers, 1957.*, 1957. URL `http://adsabs.harvard.edu/abs/1957gano.book.....G`.

R D Gehrz, A Evans, C E Woodward, L A Helton, D P K Banerjee, M K Srivastava, N M Ashok, V Joshi, S P S Eyres, Joachim Krautter, N P M Kuin, K L Page, J P Osborne, G J Schwarz, D P Shenoy, S N Shore, S G Starrfield, and R M Wagner. The Temporal Development of Dust Formation and Destruction in Nova Sagittarii 2015 no. 2 (V5668 SGR): A Panchromatic Study. *The Astrophysical Journal*, 858 (78), 2018. doi: 10.3847/1538-4357/aaba81. URL https://doi.org/10.3847/1538-4357/aaba81.

Robert D Gehrz, James W Truran, Robert E Williams, and Sumner Starrfield. Nucleosynthesis in Classical Novae and Its Contribution to the Interstellar Medium. *Publications of the Astronomical Society of the Pacific*, 110:3–26, 1998. URL http://iopscience.iop.org/article/10.1086/316107/pdf.

David Grant, Katherine Blundell, and James Matthews. Uncovering the orbital dynamics of stars hidden inside their powerful winds: Application to η Carinae and RMC 140. *Monthly Notices of the Royal Astronomical Society*, 494(1):17–35, 5 2020. ISSN 13652966. doi: 10.1093/mnras/staa669.

I. Hachisu and M. Kato. An extremely massive white dwarf of the symbiotic classical nova V407 Cyg as suggested by the rs oph and u sco models. *Baltic Astronomy*, 21 (1-2):68–75, 2012. ISSN 13920049. doi: 10.1515/astro-2017-0360.

Izumi Hachisu and Mariko Kato. A Light-curve Analysis of 32 Recent Galactic Novae: Distances and White Dwarf Masses. *The Astrophysical Journal Supplement Series*, 242:18, 2019. doi: 10.3847/1538-4365/ab1b43. URL https://doi.org/10.3847/1538-4365/ab1b43.

Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, 9 2020. ISSN 14764687. doi: 10.1038/s41586-020-2649-2.

E J Harvey, M P Redman, M J Darnley, S C Williams, A Berdyugin, V E Piirola, K P Fitzgerald, and E G P O'Connor. Polarimetry and Spectroscopy of the

'Oxygen Flaring' DQ Herculis-like nova: V5668 Sagittarii (2015). *Astronomy & Astrophysics*, 611(A3), 2018. doi: 10.1051/0004-6361/201731741.

Margarita Hernanz and Gloria Sala. A classical nova, V2487 Oph 1998, seen in x-rays before and after its explosion. *Science*, 298(5592):393–395, 10 2002. ISSN 00368075. doi: 10.1126/science.298.5592.393. URL http://xxx.lanl.gov/astro-ph/0110265.

Matthew J Holman and Paul A Wiegert. Long-term stability of planets in binary systems. Technical report, Harvard-Smithsonian Center for Astrophysics, 1999.

R. Hounsell, M. J. Darnley, M. F. Bode, D. J. Harman, L. A. Helton, and G. J. Schwarz. A very luminous, highly extinguished, very fast nova — V1721 Aquilae. *Astronomy & Astrophysics*, 530:A81, 6 2011. ISSN 0004-6361. doi: 10.1051/0004-6361/201016085. URL http://www.aanda.org/10.1051/0004-6361/201016085.

Steve B. Howell, Lorne A. Nelson, and Saul Rappaport. An Exploration of the Paradigm for the 2–3 Hour Period Gap in Cataclysmic Variables. *The Astrophysical Journal*, 550(2):897–918, 4 2001. ISSN 0004-637X. doi: 10.1086/319776. URL https://iopscience.iop.org/article/10.1086/319776.

T. Iijima and H. H. Esenoglu. Spectral evolution of Nova (V1494) Aql and its high velocity jets. *Astronomy and Astrophysics*, 404(3):997–1009, 2003. ISSN 00046361. doi: 10.1051/0004-6361:20030528.

D. Jack, J. de J. Robles Pérez, I. De Gennaro Aquino, K. P. Schröder, U. Wolter, P. Eenens, J. H.M.M. Schmitt, M. Mittag, A. Hempelmann, J. N. González-Pérez, G. Rauw, and P. H. Hauschildt. Study of the variability of Nova V5668 Sgr, based on high-resolution spectroscopic monitoring. *Astronomische Nachrichten*, 338(1): 91–102, 1 2017. ISSN 15213994. doi: 10.1002/asna.201613217.

Mariko Kato. Recurrent Novae, Classical Novae, Symbiotic Novae, and Population II Novae. In *The Physics of Cataclysmic Variables and Related Objects*, pages 595–604, 2002.

Mariko Kato and Izumi Hachisu. High-velocity jets in recurrent nova outbursts. *The Astrophysical Journal*, 587:39–42, 2003.

Koji S. Kawabata, Youichi Ohyama, Noboru Ebizuka, Tadafumi Takata, Michitoshi Yoshida, Mizuki Isogai, Yuji Norimoto, Akira Okazaki, and Masashi S. Saitou. Low- and Medium-Dispersion Spectropolarimetry of Nova V475 Scuti (Nova Scuti 2003): Discovery of an Asymmetric High-Velocity Wind in a Moderately Fast Nova. *The Astronomical Journal*, 132(1):433–442, 7 2006. ISSN 0004-6256. doi: 10.1086/505040. URL https://iopscience.iop.org/article/10.1086/505040.

G M Kennedy, M C Wyatt, B Sibthorpe, G Duchêne, P Kalas, B C Matthews, J S Greaves, K Y L Su, and M P Fitzgerald. 99 Herculis: host to a circumbinary polar-ring debris disc. *Mon. Not. R. Astron. Soc*, 421:2264–2276, 2012. doi: 10.1111/j.1365-2966.2012.20448.x. URL https://academic.oup.com/mnras/article-abstract/421/3/2264/1077552.

U Kolb and I Baraffe. Brown dwarfs and the cataclysmic variable period minimum. Technical report, University of Leicester, 1999. URL https://academic.oup.com/mnras/article-abstract/309/4/1034/1084865.

A. Kramida, Yu. Ralchenko, and J. Reader. NIST Atomic Spectra Database (version 5.2), 2014. URL https://www.nist.gov/pml/atomic-spectra-database.

Joachim Krautter, Charles E. Woodward, Michael T. Schuster, Robert D. Gehrz, Terry J. Jones, Kunegunda Belle, A. Evans, S. P. S. Leyers, Sumner Starrfield, James Truran, and Matthew A. Greenhouse. Hubble Space Telescope NICMOS Observations of Classical Nova Shells. *The Astronomical Journal*, 124(5):2888–2898, 11 2002. ISSN 00046256. doi: 10.1086/343833. URL https://iopscience.iop.org/article/10.1086/343833.

Li Kwan-Lok. ATel #13868: Fermi-LAT detection of the naked-eye classical nova MGAB-V207, 2020. URL http://www.astronomerstelegram.org/?read=13868.

L. Izzo, P. Molaro, E. Aydi, P. Bonifacio, G. Cescutti, M. Della Valle, E. Harvey, M. Hernanz, G. Schiulaz, and P. Selvelli. ATel 14048: UVES observations of Nova Reticuli 2020 during Minimal Science Operations show it is entering the nebular phase. *The Astronomer's Telegram*, 14048(1), 2020. URL http://www.astronomerstelegram.org/?read=14048.

Jean-Pierre Lasota and Noam Soker. Why there are no jets from cataclysmic variable stars. In *The Astrophysics of Cataclysmic Variables and Related Objects ASP Conference*, volume 330, 2005.

Kwan-Lok Li, Brian D Metzger, Laura Chomiuk, Indrek Vurm, Jay Strader, Thomas Finzell, Andrei M Beloborodov, Thomas Nelson, Benjamin J Shappee, Christopher S Kochanek, José L Prieto, Stella Kafka, Thomas W.-S. Holoien, Todd A Thompson, Paul J Luckas, and Hiroshi Itoh. A nova outburst powered by shocks. *Nature Astronomy*, 1(10):697–702, 2017. ISSN 2397-3366. doi: 10.1038/s41550-017-0222-1. URL `https://doi.org/10.1038/s41550-017-0222-1`.

Wei-Min Liu and Xiang-Dong Li. Can the Friction of the Nova Envelope Account for the Extra Angular Momentum Loss in Cataclysmic Variables? *The Astrophysical Journal*, 870:22, 2018. doi: 10.3847/1538-4357/aaf19f. URL `https://doi.org/10.3847/1538-4357/aaf19f`.

Knox S. Long and Christian Knigge. Modeling the Spectral Signatures of Accretion Disk Winds: A New Monte Carlo Approach. *The Astrophysical Journal*, 579:725–740, 2002. ISSN 0004-637X. doi: 10.1086/342879.

P. Martin, G. Dubus, P. Jean, V. Tatischeff, and C. Dosne. Gamma-ray emission from internal shocks in novae. *Astronomy & Astrophysics*, 612:A38, 4 2018. doi: 10.1051/0004-6361/201731692. URL `https://doi.org/10.1051/0004-6361/201731692https://www.aanda.org/10.1051/0004-6361/201731692`.

Wes McKinney. Data Structures for Statistical Computing in Python. Technical report, AQR Capital Management, 2010. URL `http://conference.scipy.org/proceedings/scipy2010/pdfs/mckinney.pdf`.

Wes Mckinney. Data Structures for Statistical Computing in Python. *PROC. OF THE 9th PYTHON IN SCIENCE CONF*, pages 56–61, 2010.

Dean B McLaughlin. SPECTRAL STAGES OF NOVAE. *The Astrophysical Journal*, 95:428–436, 1942.

Dean B McLaughlin. A Model of a Nova Outburst. Technical report, Maria Mitchell Observatory, 1950. URL `http://articles.adsabs.harvard.edu/cgi-bin/nph-iarticle_query?1950PA.....58...50M&data_type=PDF_HIGH&whole_paper=YES&type=PRINTER&filetype=.pdf`.

Dean B. McLaughlin. The spectra of novae. *Vistas in Astronomy*, 2:1477–1485, 1956. ISSN 00836656. doi: 10.1016/0083-6656(56)90076-9.

Dominic McLoughlin, Katherine M Blundell, and Steven Lee. Classical Nova Carinae 2018: Discovery of circumbinary iron and oxygen. *Monthly Notices of the Royal Astronomical Society*, 494(1):743–749, 2020. doi: 10.1093/mnras/staa651/5801033. URL https://academic.oup.com/mnras/advance-article-abstract/doi/10.1093/mnras/staa651/5801033.

Dominic McLoughlin, Katherine M Blundell, Steven Lee, and Chris McCowage. The precessing jets of classical nova YZ Reticuli. *Monthly Notices of the Royal Astronomical Society*, 503(1):704–714, 5 2021a. doi: 10.1093/mnras/stab581/6153877. URL https://academic.oup.com/mnras/advance-article/doi/10.1093/mnras/stab581/6153877.

Dominic McLoughlin, Katherine M Blundell, Steven Lee, and Chris McCowage. The onset of jets in classical novae. *Monthly Notices of the Royal Astronomical Society*, 505(2):2518–2529, 6 2021b. ISSN 0035-8711. doi: 10.1093/mnras/stab1364. URL https://academic.oup.com/mnras/article/505/2/2518/6296682.

A. Mehner, W. J. De Wit, J. H. Groh, R. D. Oudmaijer, D. Baade, T. Rivinius, F. Selman, H. M.J. Boffin, and C. Martayan. VLT/MUSE discovers a jet from the evolved B[e] star MWC 137. *Astronomy and Astrophysics*, 585(A81):1–11, 1 2016. ISSN 14320746. doi: 10.1051/0004-6361/201527180.

B D Metzger, G Martínez-Pinedo, S Darbha, E Quataert, A Arcones, D Kasen, R Thomas, P Nugent, I V Panov, and N T Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *Monthly Notices of the Royal Astronomical Society*, 406(4):2650–2662, 8 2010. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.16864.x. URL https://doi.org/10.1111/j.1365-2966.2010.16864.x.

P Molaro, L Izzo, E Mason, P Bonifacio, and M Della Valle. Highly enriched Be in the ejecta of Nova Sagittarii 2015 No. 2 (V5668 Sgr) and the Galactic Li origin. *MNRAS*, 463:117–121, 2016. doi: 10.1093/mnrasl/slw169.

J D Monnier, R K Barry, W A Traub, B F Lane, R L Akeson, S Ragland, P A Schuller, H Le Coroller, J.-P Berger, R Millan-Gabet, E Pedretti, F P Schloerb, C Koresko, N P Carleton, M G Lacasse, P Kern, F Malbet, K Perraut, M J Kuchner, and M W Muterspaugh. No expanding fireball: resolving the recurrent nova RS Ophiuchi with infrared interferometry. Technical report, University of Michigan, 2006. URL http://msc.caltech.edu.

Y Osaki. Dwarf nova outbursts: a unification theory. In Emmi Meyer-Hofmeister and Henk Spruit, editors, *Accretion Disks — New Aspects*, pages 92–102, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg. ISBN 978-3-540-68715-3.

Joseph Patterson. Accretion-Disk Precession and Substellar Secondaries in Cataclysmic Variables. *Publications of the Astronomical Society of the Pacific*, 113: 736–747, 2001.

M. Pavana, A. Raj, T. Bohlsen, G. C. Anupama, Ranjan Gupta, and G. Selvakumar. Spectroscopic and geometrical evolution of the ejecta of the classical nova ASASSN-18fv. *Monthly Notices of the Royal Astronomical Society*, 495:2075–2087, 4 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1219.

Fabian Pedregosa, Vincent Michel, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Jake Vanderplas, David Cournapeau, Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Bertrand Thirion, Olivier Grisel, Vincent Dubourg, Alexandre Passos, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. URL `http://scikit-learn.sourceforge.net`.

Sebastian Perez and Katherine M Blundell. SS 433's circumbinary ring and accretion disc viewed through its attenuating disc wind. *Mon. Not. R. Astron. Soc*, 408:2–8, 2010. doi: 10.1111/j.1365-2966.2010.16638.x. URL `https://academic.oup.com/mnras/article-abstract/408/1/2/1053848`.

Geraldine J Peters, Luqian Wang, Douglas R Gies, and Erika D Grundstrom. The hot companion and circumbinary disk of the Be star HR 2142. *The Astrophysical Journal*, 828, 2016. doi: 10.3847/0004-637X/828/1/47. URL `http://nova.astro.umd.edu/`.

William H. Press and Saul A. Teukolsky. Savitzky-Golay Smoothing Filters. *Computers in Physics*, 4(6):669, 1990. ISSN 08941866. doi: 10.1063/1.4822961. URL `http://scitation.aip.org/content/aip/journal/cip/4/6/10.1063/1.4822961`.

M Price-Whelan, B M Sipőcz, H M Günther, P L Lim, S M Crawford, S Conseil, D L Shupe, M W Craig, N Dencheva, A Ginsburg, J T Vanderplas, L D Bradley, D Pérez-Suárez, M De Val-Borro, T L Aldcroft, K L Cruz, T P Robitaille, E J Tollerud, C Ardelean, T Babej, Y P Bach, M Bachetti, A V Bakanov, S P Bamford, G Barentsen, P Barmby, A Baumbach, K L Berry, F Biscani, M Boquien, K A Bostroem, L G Bouma, G B Brammer, E M Bray, H Breytenbach, H Buddelmeijer, D J Burke, G Calderone, J L Cano Rodríguez, M Cara, J V M Cardoso, D Crichton, D D'avella, C Deil, É Depagne, J P Dietrich, A Donath, M Droettboom, N Earl, T Erben, S Fabbro, L A Ferreira, T Finethy, R T Fox, L H Garrison, S L J Gibbons, D A Goldstein, R Gommers, J P Greco, P Greenfield, A M Groener, F Grollier, A Hagen, P Hirst, D Homeier, A J Horton, G Hosseinzadeh, L Hu, J S Hunkeler, Ž Ivezić, A Jain, T Jenness, G Kanarek, S Kendrew, N S Kern, W E Kerzendorf, A Khvalko, J King, D Kirkby, A M Kulkarni, A Kumar, A Lee, D Lenz, S P Littlefair, Z Ma, D M Macleod, M Mastropietro, C Mccully, S Montagnac, B M Morris, M Mueller, S J Mumford, D Muna, N A Murphy, S Nelson, G H Nguyen, J P Ninan, M Nöthe, S Ogaz, S Oh, J K Parejko, N Parley, S Pascual, R Patil, A A Patil, A L Plunkett, J X Prochaska, T Rastogi, V Reddy Janga, J Sabater, P Sakurikar, M Seifert, L E Sherbert, H Sherwood-Taylor, A Y Shih, J Sick, M T Silbiger, S Singanamalla, L P Singer, P H Sladen, K A Sooley, S Sornarajah, O Streicher, P Teuben, S W Thomas, G R Tremblay, J E H Turner, V Terrón, M H Van Kerkwijk, A De La Vega, L L Watkins, B A Weaver, J B Whitmore, J Woillez, and V Zabalza. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package * The Astropy Collaboration, A. *The Astronomical Journal*, 156:123, 2018. doi: 10.3847/1538-3881/aabc4f. URL https://doi.org/10.3847/1538-3881/aabc4f.

T Prusti, J H J de Bruijne, A G A Brown, A Vallenari, C Babusiaux, C A L Bailer-Jones, U Bastian, M Biermann, D W Evans, L Eyer, F Jansen, C Jordi, S A Klioner, U Lammers, L Lindegren, X Luri, F Mignard, D J Milligan, C Panem, V Poinsignon, D Pourbaix, S Randich, G Sarri, P Sartoretti, H I Siddiqui, C Soubiran, V Valette, F van Leeuwen, N A Walton, C Aerts, F Arenou, M Cropper, R Drimmel, E Høg, D Katz, M G Lattanzi, E K Grebel, A D Holland, C Huc, X Passot, L Bramante, C Cacciari, J Castañeda, L Chaoul, N Cheek, F De Angeli, C Fabricius, R Guerra, J Hernández, A Jean-Antoine-Piccolo, E Masana, R Messineo, N Mowlavi, K Nienartowicz, D Ordóñez-Blanco, P Panuzzo, J Portell, P J Richards, C Barache, C Barata, A Barbier, F Barblan, M Baroni, D Barrado Navascués, M Barros, M A Barstow, U Becciani, M Bellazzini, G Bellei, A Bello García, V Belokurov, P Bendjoya, A Berihuete, L Bianchi, O Bienaymé, F Billebaud, N Blagorodnova, S Blanco-Cuaresma, T Boch, A Bombrun, R Borrachero, S Bouquillon, G Bourda, H Bouy, A Bragaglia, M A Breddels, N Brouillet, T Brüsemeister, B Bucciarelli, F Budnik, P Burgess, R Burgon, A Burlacu,

D Busonero, R Buzzi, E Caffau, J Cambras, H Campbell, R Cancelliere, T Cantat-Gaudin, T Carlucci, J M Carrasco, M Castellani, P Charlot, J Charnas, P Charvet, F Chassat, A Chiavassa, M Clotet, G Cocozza, R S Collins, P Collins, G Costigan, F Crifo, N J G Cross, M Crosta, C Crowley, C Dafonte, Y Damerdji, A Dapergolas, P David, M David, P De Cat, F de Felice, P de Laverny, F De Luise, R De March, D de Martino, R de Souza, J Debosscher, E del Pozo, M Delbo, A Delgado, H E Delgado, F di Marco, P Di Matteo, S Diakite, E Distefano, C Dolding, S Dos Anjos, P Drazinos, J Durán, Y Dzigan, E Ecale, B Edvardsson, H Enke, M Erdmann, D Escolar, M Espina, N W Evans, G Eynard Bontemps, C Fabre, M Fabrizio, S Faigler, A J Falcão, M Farràs Casas, F Faye, L Federici, G Fedorets, J Fernández-Hernández, P Fernique, A Fienga, F Figueras, F Filippi, K Findeisen, A Fonti, M Fouesneau, E Fraile, M Fraser, J Fuchs, R Furnell, M Gai, S Galleti, L Galluccio, D Garabato, F García-Sedano, P Garé, A Garofalo, N Garralda, P Gavras, J Gerssen, R Geyer, G Gilmore, S Girona, G Giuffrida, M Gomes, A González-Marcos, J González-Núñez, J J González-Vidal, M Granvik, A Guerrier, P Guillout, J Guiraud, A Gúrpide, R Gutiérrez-Sánchez, L P Guy, R Haigron, D Hatzidimitriou, M Haywood, U Heiter, A Helmi, D Hobbs, W Hofmann, B Holl, G Holland, J A S Hunt, A Hypki, P Koubsky, A Kowalczyk, A Krone-Martins, M Kudryashova, I Kull, R K Bachchan, F Lacoste-Seris, A F Lanza, J-b Lavigne, C Le Poncin-Lafitte, Y Lebreton, T Lebzelter, S Leccia, N Leclerc, I Lecoeur-Taibi, V Lemaitre, H Lenhardt, F Leroux, S Liao, E Licata, H E P Lindstrøm, T A Lister, E Livanou, A Lobel, W Löffler, M López, A Lopez-Lozano, D Lorenz, T Loureiro, I MacDonald, T Magalhães Fernandes, S Managau, R G Mann, G Mantelet, O Marchal, C Pagani, I Pagano, F Pailler, H Palacin, L Palaversa, P Parsons, T Paulsen, M Pecoraro, R Pedrosa, H Pentikäinen, J Pereira, B Pichon, A M Piersimoni, F-x Pineau, E Plachy, G Plum, E Poujoulet, A Prša, L Pulone, S Ragaini, S Rago, N Rambaux, M Ramos-Lerate, P Ranalli, G Rauw, A Read, S Regibo, F Renk, C Reylé, R A Ribeiro, L Rimoldini, V Ripepi, A Riva, G Rixon, M Roelens, M Romero-Gómez, N Rowell, F Royer, A Rudolph, L Ruiz-Dern, G Sadowski, T Sagristà Sellés, J Sahlmann, J Salgado, E Salguero, M Sarasso, H Savietto, A Schnorhk, M Schultheis, E Sciacca, M Segol, J C Segovia, D Segransan, E Serpell, I-c Shih, R Smareglia, R L Smart, C Smith, E Solano, F Solitro, R Sordo, S Soria Nieto, J Souchay, A Spagna, F Spoto, U Stampa, I A Steele, H Steidelmüller, C A Stephenson, H Stoev, F F Suess, M Süveges, J Surdej, L Szabados, E Szegedi-Elek, D Tapiador, F Taris, G Tauran, M B Taylor, R Teixeira, D Terrett, B Tingley, S C Trager, C Turon, A Ulla, E Utrilla, G Valentini, A van Elteren, E Van Hemelryck, M van

Leeuwen, M Varadi, A Vecchiato, J Veljanoski, T Via, D Vicente, S Vogt, H Voss, V Votruba, S Voutsinas, G Walmsley, M Weiler, K Weingrill, D Werner, T Wevers, G Whitehead, Ł Wyrzykowski, A Yoldas, M Žerjal, S Zucker, C Zurbach, T Zwitter, A Alecu, M Allen, C Allende Prieto, A Amorim, G Anglada-Escudé, V Arsenijevic, S Azaz, P Balm, M Beck, H-h Bernstein, L Bigot, A Bijaoui, C Blasco, M Bonfigli, G Bono, S Boudreault, A Bressan, S Brown, P-m Brunet, P Bunclark, R Buonanno, A G Butkevich, C Carret, C Carrion, L Chemin, F Chéreau, L Corcione, E Darmigny, K S de Boer, P de Teodoro, P T de Zeeuw, C Delle Luche, C D Domingues, P Dubath, F Fodor, B Frézouls, A Fries, D Fustes, D Fyfe, E Gallardo, J Gallegos, D Gardiol, M Gebran, A Gomboc, A Gómez, E Grux, A Gueguen, A Heyrovsky, J Hoar, G Iannicola, Y Isasi Parache, A-m Janotto, E Joliet, A Jonckheere, R Keil, D-w Kim, P Klagyivik, J Klar, J Knude, O Kochukhov, I Kolka, J Kos, A Kutka, V Lainey, D LeBouquin, C Liu, D Loreggia, V V Makarov, M G Marseille, C Martayan, O Martinez-Rubi, B Massart, F Meynadier, S Mignot, U Munari, A-t Nguyen, T Nordlander, P Ocvirk, A Olias Sanz, P Ortiz, J Osorio, D Oszkiewicz, A Ouzounis, M Palmer, P Park, E Pasquato, C Peltzer, J Peralta, F Péturaud, T Pieniluoma, and E Pigozzi. The Gaia mission. *Astronomy & Astrophysics*, 595(A1):1–36, 2016. doi: 10.1051/0004-6361/201629272. URL http://www.cosmos.esa.int/gaia.

Alon Retter. Constraining jet theories using nova outbursts. *The Astrophysical Journal*, 615:125–128, 2004.

Alon Retter. The Transition Phase in Classical Novae. In *ASP Conference Series*, pages 232–235. University of Sydney, 2006.

V. A. R. M. Ribeiro, U. Munari, and P. Valisa. Optical morphology, inclination and expansion velocity of the ejected shell of Nova Monocerotis 2012. *The Astrophysical Journal*, 768(1):49, 4 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/768/1/49. URL https://iopscience.iop.org/article/10.1088/0004-637X/768/1/49.

D. H. Roberts, J. Lehar, and J. W. Dreher. Time Series Analysis with Clean - Part One - Derivation of a Spectrum. *The Astronomical Journal*, 93(4):968, 1987. ISSN 00046256. doi: 10.1086/114383.

Thomas P Robitaille, Erik J Tollerud, Perry Greenfield, Michael Droettboom, Erik Bray, Tom Aldcroft, Matt Davis, Adam Ginsburg, Adrian M Price-Whelan, Wolfgang E Kerzendorf, Alexander Conley, Neil Crighton, Kyle Barbary, Demitri Muna, Henry Ferguson, Frédéric Grollier, Madhura M Parikh, Prasanth H Nair, Hans M Günther, Christoph Deil, Julien Woillez, Simon Conseil, Roban Kramer, James E H Turner, Leo Singer, Ryan Fox, Benjamin A Weaver, Victor Zabalza, Zachary I Edwards, K Azalee Bostroem, D J Burke, Andrew R Casey, Steven M Crawford, Nadia Dencheva, Justin Ely, Tim Jenness, Kathleen Labrie, Lian Lim, Francesco Pierfederici, Andrew Pontzen, Andy Ptak, Brian Refsdal, Mathieu Servillat, and Ole Streicher. Astrophysics Astropy: A community Python package for astronomy. *A&A*, 558:33, 2013. doi: 10.1051/0004-6361/201322068. URL http://www.iausofa.org.

Michael P. Rupen, Amy J. Mioduszewski, and Jennifer L. Sokoloski. An Expanding Shell and Synchrotron Jet in RS Ophiuchi. *The Astrophysical Journal*, 688(1):559–567, 11 2008. ISSN 0004-637X. doi: 10.1086/525555. URL https://iopscience.iop.org/article/10.1086/525555.

E. Santamaría, M. A. Guerrero, G. Ramos-Larios, J. A. Toalá, L. Sabin, G. Rubio, and J. A. Quino-Mendoza. Angular Expansion of Nova Shells. *The Astrophysical Journal*, 892(1):60, 3 2020. ISSN 1538-4357. doi: 10.3847/1538-4357/ab76c5. URL https://iopscience.iop.org/article/10.3847/1538-4357/ab76c5.

K. B. Schmidt, L. Wisotzki, T. Urrutia, J. Kerutt, D. Krajnović, E. C. Herenz, R. Saust, T. Contini, B. Epinat, H. Inami, and M. V. Maseda. Three dimensional optimal spectral extraction (TDOSE) from integral field spectroscopy. *Astronomy & Astrophysics*, 628(A91):1–22, 6 2019. ISSN 0004-6361. doi: 10.1051/0004-6361/201935857.

M R Schreiber, M Zorotovic, and T P G Wijnen. Three in one go: consequential angular momentum loss can solve major problems of CV evolution. *MNRAS*, 455:16–20, 2016. doi: 10.1093/mnrasl/slv144. URL https://academic.oup.com/mnrasl/article-abstract/455/1/L16/2589619.

Greg J. Schwarz, Jan Uwe Ness, J. P. Osborne, K. L. Page, P. A. Evans, A. P. Beardmore, Frederick M. Walter, L. Andrew Helton, Charles E. Woodward, Mike Bode, Sumner Starrfield, and Jeremy J. Drake. Swift X-ray observations of classical novae. II. the super soft source sample. *Astrophysical Journal, Supplement Series*, 197(2), 12 2011. ISSN 00670049. doi: 10.1088/0067-0049/197/2/31.

A. W. Shafter. The galactic nova rate revisited. *The Astrophysical Journal*, 834(2): 196, 1 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/834/2/196.

Michael M. Shara, Trisha Doyle, Tod R. Lauer, David Zurek, Edward A. Baltz, Attay Kovetz, Juan P. Madrid, Joanna Mikołajewska, J. D. Neill, Dina Prialnik, D. L. Welch, and Ofer Yaron. A Hubble Space Telescope Survey for Novae in M87. II. Snuffing out the Maximum Magnitude–Rate of Decline Relation for Novae as a Non-standard Candle, and a Prediction of the Existence of Ultrafast Novae. *The Astrophysical Journal*, 839(2):109, 4 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/aa65cd. URL https://iopscience.iop.org/article/10.3847/1538-4357/aa65cd.

Michael M. Shara, Dina Prialnik, Yael Hillman, and Attay Kovetz. The Masses and Accretion Rates of White Dwarfs in Classical and Recurrent Novae. *The Astrophysical Journal*, 860(2):110, 6 2018. ISSN 1538-4357. doi: 10.3847/1538-4357/aabfbd.

Steven N Shore. Spectroscopy of novae - a user's manual. Technical report, Universita di Pisa, 2012.

A. Skopal. Multiwavelength modeling the SED of supersoft X-ray sources. II. RS Ophiuchi: From the explosion to the SSS phase. *Elsevier*, 36:128–138, 2015. ISSN 13841076. doi: 10.1016/j.newast.2013.12.005.

A J Slavin, T J O'Brien, and J S Dunlop. A deep optical imaging study of the nebular remnants of classical novae. *353S Mon. Not. R. Astron. Soc*, 276:353–371, 1995.

J. L. Sokoloski, S. J. Kenyon, B. R. Espey, Charles D. Keyes, S. R. McCandliss, A. K. H. Kong, J. P. Aufdenberg, A. V. Filippenko, W. Li, C. Brocksopp, Christian R. Kaiser, P. A. Charles, M. P. Rupen, and R. P. S. Stone. A "Combination Nova" Outburst in Z Andromedae: Nuclear Shell Burning Triggered by a Disk Instability. *The Astrophysical Journal*, 636(2):1002–1019, 1 2006. ISSN 0004-637X. doi: 10.1086/498206. URL https://iopscience.iop.org/article/10.1086/498206.

J L Sokoloski, M P Rupen, and A J Mioduszewski. Uncovering the nature of nova jets: a radio image of highly collimated outflows from RS Ophiuchi. *The Astrophysical Journal*, 685:137–140, 2008.

K. V. Sokolovsky. ATel #13900: NuSTAR detection of Nova Reticuli 2020 = MGAB-V207, 2020. URL http://www.astronomerstelegram.org/?read=13900.

Kirill Sokolovsky, Elias Aydi, Laura Chomiuk, and Adam Kawash. ATel #14043: Super-soft X-ray emission of Nova Reticuli 2020, 2020a. URL http://www.astronomerstelegram.org/?read=14043.

Kirill V Sokolovsky, Koji Mukai, Laura Chomiuk, Raimundo Lopes de Oliveira, Elias Aydi, Kwan-Lok Li, Elad Steinberg, Indrek Vurm, Brian D Metzger, Adam Kawash, Justin D Linford, Amy J Mioduszewski, Thomas Nelson, Jan-Uwe Ness, Kim L Page, Michael P Rupen, Jennifer L Sokoloski, and Jay Strader. X-ray spectroscopy of the γ-ray brightest nova V906 Car (ASASSN-18fv). *Monthly Notices of the Royal Astronomical Society*, 497(3):2569–2585, 9 2020b. ISSN 0035-8711. doi: 10.1093/mnras/staa2104. URL https://academic.oup.com/mnras/article/497/3/2569/5873688.

Pei Songpeng. ATel #14067: NICER observations of Nova Ret 2020, 2020a. URL http://www.astronomerstelegram.org/?read=14067.

Pei Songpeng. ATel #14205: Coronal Line Emission in YZ Reticuli (Nova Reticuli 2020), 2020b. URL http://www.astronomerstelegram.org/?read=14205.

H C Spruit and Ronald E Taam. Circumbinary disks and cataclysmic variable evolution. *The Astrophysical Journal*, 548:900–907, 2001. URL https://iopscience.iop.org/article/10.1086/319030/pdf.

K Z Stanek, Thomas W-S Holoien, C S Kochanek, J V Shields, T A Thompson, L Chomiuk, J Strader, B J Shappee, J L Prieto, S Dong, and M Stritzinger. ATel #11454: ASAS-SN Discovery of a Possible, Very Bright Galactic Nova ASASSN-18fv, 2018. URL http://www.astronomerstelegram.org/?read=11454.

Richard J. Strope, Bradley E. Schaefer, and Arne A. Henden. Catalog of 93 nova light curves: Classification and properties. *Astronomical Journal*, 140(1):34–62, 2010. ISSN 00046256. doi: 10.1088/0004-6256/140/1/34.

Victor Szebehely. Stability of planetary orbits in binary systems. Technical report, University of Texas, 1980.

Akito Tajitsu, Kozo Sadakane, Hiroyuki Naito, Akira Arai, Hideyo Kawakita, and Wako Aoki. The $ˆ{7}$Be II Resonance Lines in Two Classical Novae V5668 Sgr and V2944 Oph. *The Astrophysical Journal*, 818(191):1–9, 2016. doi: 10.3847/0004-637X/818/2/191. URL http://arxiv.org/abs/1601.05168%0Ahttp://dx.doi.org/10.3847/0004-637X/818/2/191.

A. R. Taylor, R. J. Davis, R. W. Porcas, and M. F. Bode. VLBI observations of RS Oph – a recurrent nova with non-spherical ejection. *Monthly Notices of the Royal Astronomical Society*, 237(1):89–91, 3 1989.

J A Toalá, M A Guerrero, E Santamaría, G Ramos-Larios, and L Sabin. Extended X-ray emission from the classic nova DQ Her – on the possible presence of a magnetized jet. *Monthly Notices of the Royal Astronomical Society*, 495 (4):4372–4379, 7 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1502. URL https://academic.oup.com/mnras/article/495/4/4372/5850382.

Dean M Townsley and Lars Bildsten. Classical novae as a probe of the Cataclysmic Variable population. Technical report, University of California, 2005. URL http://physics.open.ac.uk.

Peter G Tuthill, John D Monnier, Nicholas Lawrance, William C Danchi, Stan P Owocki, and Kenneth G Gayley. The prototype colliding-wind pinwheel WR 104. Technical report, University of Sydney, 2007.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R.J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silterra, James T. Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-Baeza. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3): 261–272, 3 2020. ISSN 15487105. doi: 10.1038/s41592-019-0686-2.

Brian Warner. *Cataclysmic Variable Stars*. Cambridge University Press, 12 1995. ISBN 9780521412315. doi: 10.1017/CBO9780511586491. URL `https://www.cambridge.org/core/product/identifier/9780511586491/type/book`.

Jerrick Wee, Nadejda Blagorodnova, Bryan Edward Penprase, Jett Pierce Facey, Taiga Morioka, Hank Corbett, Brad N Barlow, Thomas Kupfer, Nicholas M Law, Jeffrey K Ratzloff, Ward S Howard, Ramses Gonzalez Chavez, Amy Glazier, Alan Vasquez Soto, and Takashi Horiuchi. Multiwavelength Photometry and Progenitor Analysis of the Nova V906 Car. *The Astrophysical Journal*, 899:162, 2020. doi: 10.3847/1538-4357/aba3cc. URL `https://doi.org/10.3847/1538-4357/aba3cc`.

R E Williams, M Hamuy, M M Phillips, S R Heathcote, Lisa Wells, and M Navarrete. The evolution and classification of post-outburst novae spectra. *The Astrophysical Journal*, 376:721–737, 1991.

Robert Williams. Episodic Mass Transfer: A Trigger for Nova Outbursts? *Proceedings of the International Astronomical Union*, 7(S281):88–95, 2011. ISSN 1743-9213. doi: 10.1017/S1743921312014755. URL `https://www.cambridge.org/core/product/16888970600A683E73545E20FC404438`.

Robert Williams and Elena Mason. Novae ejecta as colliding shells. *Astrophysics and Space Science*, 327(2):207–217, 6 2010. ISSN 0004640X. doi: 10.1007/s10509-010-0318-x.

Robert Williams, Elena Mason, Massimo Della Valle, and Alessandro Ederoclite. Transient heavy element absorption systems in novae: episodic mass ejection from the secondary star. Technical report, Space Telescope Science Institute, 2008. URL `https://iopscience.iop.org/article/10.1086/590056/pdf`.

Robert E Williams. The Formation of Novae Spectra. *The Astronomical Journal*, 104 (2):725–733, 1992. URL `http://articles.adsabs.harvard.edu/pdf/1992AJ....104..725W`.

Robert E. Williams. Extinction, ejecta masses, and radial velocities of novae. *The Astrophysical Journal*, 426(426):279–287, 5 1994. ISSN 0004-637X. doi: 10.1086/174062. URL `http://adsabs.harvard.edu/doi/10.1086/174062`.

Monica Zorotovic and Matthias R. Schreiber. Cataclysmic variable evolution and the white dwarf mass problem: A Review. *Advances in Space Research*, 9 2019. ISSN 02731177. doi: 10.1016/j.asr.2019.08.044. URL `https://linkinghub.elsevier.com/retrieve/pii/S0273117719306453`.