

Orbits and Outflows of Emission-line Stars



David Grant
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2021

Statement of Originality

The work presented in this thesis was undertaken at the Department of Astrophysics at the University of Oxford between October 2017 and August 2021 under the supervision of Professor Katherine Blundell OBE. I hereby declare that no part of this thesis has been submitted in support of another degree, diploma, certificate or other qualification at the University of Oxford or elsewhere. Except where otherwise stated, or where reference is made to the work of others, the work in this thesis is entirely my own.

The work presented in Chapter 4 is based on the paper Grant et al. (2020) (Monthly Notices of the Royal Astronomical Society, Volume 494, p.17-35), for which I am the lead author. However, Section 4.4.5 is based on the simulations I ran for the publication Porter et al. (2021a), (Monthly Notices of the Royal Astronomical Society, Volume 501, p.5554-5574), for which I am the second author. The work presented in Chapter 5 is based on the paper Grant & Blundell (2021) (Monthly Notices of the Royal Astronomical Society, Volume 509, p.367-379), for which I am the lead author. The work presented in Chapter 6 and 7 are based on papers to be imminently submitted for peer review. This work was carried out with regular consultation with Prof. Blundell and Dr. James Matthews.

Acknowledgements

First and foremost, I would like to thank my supervisor, Professor Katherine Blundell. You have been the best mentor throughout my time in Oxford: encouraging and brilliant day-to-day, and understanding and adaptive through the tough times. Every meeting of ours made me feel more positive and full of ideas. The Global Jet Watch is a triumph and it has been a privilege to have been a part of your group.

Thinking back, I am grateful to a long list of inspirational people over the course of my education. At school it was Mr Shaw, at Pembroke it was Alfons, Bill, and Eduardo, and as a research assistant it was the amazing then-post docs, and now soon-to-be professors. Especially to James Matthews, I could not have been luckier to have been given the desk adjacent to you. Peering over at your work filled me with the aspiration to pursue astrophysics. It feels like only yesterday I was faking it with you and Rebecca at a grant-proposal coffee. Also, shout out to Aprajita, Chris, Dimitra, Gus, Ryo, Sam, and Kit.

Outside of physics, the last few years have been made particularly memorable by a handful of great friends. To all my housemates from Bullingdon, Hurst, Hobson, and Hill View, thank you for all the good times, like floating down the river, the xmas party, and making the best of the pandemic together. Also, to my friends from the volleyball club, I will always fondly remember biking around late in Oxford's orange glow, arduous away trips on the bus, the tower of bags in PT, and playing in the hail at the beach courts. Jørgen, I am not sorry for always keeping you up too long.

Finally, to my parents and brother, I am eternally grateful for your support and for fostering an environment in which I could become anything I wanted. Additionally, big thanks to the tiny Katie and Salem.

Abstract

The optical spectra of some stars exhibit spectral lines in emission, which is characteristic of them having powerful radiatively driven winds. These stars are known as emission-line stars. In this thesis, I study the orbits and outflows of the most luminous and massive emission-line stars: Luminous Blue Variables and Wolf-Rayet stars.

Orbital parameters are typically determined by fitting Keplerian models to radial velocity measurements. For stars with powerful winds, the broad emission lines are formed far out in the stellar wind, encoding dynamical information from an extensive volume of outflowing gas, in a complex and stratified circumstellar environment, which may lead to systematics in the observables. I tackle these problems twofold. (1) I formulate a novel semi-analytical model which encapsulates both the star's orbital motion and the propagation of the wind. (2) I devise a statistical framework for validating the performance of algorithms, and utilise a marginalised Gaussian process to account for residual systematics. I benchmark our models with synthetic data and apply them to several prototypical emission-line binaries, namely Eta Carinae, GG Carinae, WR 140, and WR 133. I find our models identify a consistent set of orbital parameters, independent of the emission line used, and given validation on benchmark datasets, the estimated parameter distributions and dynamical masses are more accurate than those previously calculated.

The Global Jet Watch affords unprecedented coverage of Eta Carinae throughout apastron, revealing the distinct dynamics of the emitting versus the absorbing components of the He I profiles. Cognisant of the known Helium ionisation maps, I model these absorption velocities as tracers of the geometry of the post-shock gas in the colliding winds. I find the modelling is concordant with the panchromatic observations and simulations of Eta Carinae and demonstrate how the model can be used to probe the latitudinal dependence of stellar mass loss. Finally, I describe the Global Jet Watch's photometric datastream, and detail the innovation of new algorithms and the resulting light curves.

Contents

1	Introduction	1
1.1	Emission-line stars	2
1.2	Binary stars	4
1.3	The Global Jet Watch	5
1.4	Thesis overview	8
2	Stellar outflows	9
2.1	Types of mass loss	9
2.2	Wind acceleration	11
2.3	Single star flows	17
2.3.1	Wind velocity profiles	18
2.3.2	Numerical simulations	19
2.4	Colliding flows	25
2.4.1	High energy radiation	28
2.5	Summary	29
3	Stellar orbits	30
3.1	Types of orbits	30
3.2	Keplerian motion	34
3.2.1	Radial velocities	36
3.2.2	Astrometric positions	39
3.3	Algorithms for velocity extraction	45
3.3.1	Bisector half-maximum	47
3.3.2	Cross-correlation	48
3.3.3	Multi-Gaussian	50
3.4	Summary	54

4	Uncovering the orbital dynamics of stars hidden inside their powerful winds	55
4.1	Introduction	55
4.1.1	Eta Carinae: the extreme case	56
4.1.2	RMC 140b and GG Carinae: the typical cases	58
4.2	Observations	59
4.2.1	Radial velocities: the Balmer series	62
4.3	Modelling	64
4.3.1	Keplerian motion	64
4.3.2	Convolutional Keplerian motion	68
4.4	Discussion	81
4.4.1	CKM time-scales	81
4.4.2	The outflows of Eta Carinae	83
4.4.3	The orbital dynamics of Eta Carinae	87
4.4.4	Application to RMC 140b	90
4.4.5	Application to GG Carinae	92
4.5	Conclusions	94
5	Probabilistic orbits and dynamical masses of emission-line binaries using marginalised Gaussian processes	97
5.1	Introduction	97
5.2	A benchmark dataset	98
5.2.1	Evaluation criterion	100
5.3	A benchmark algorithm	101
5.3.1	Mean models	102
5.3.2	Gaussian processes	103
5.3.3	Bayesian regression	105
5.4	Setting the standard	107
5.5	Discussion	111
5.5.1	The performance impact of additional data	111
5.5.2	Application to prototypical systems	113
5.6	Conclusions	125

6	Tracing the colliding winds of Eta Carinae in He I	127
6.1	Introduction	127
6.2	Observations	129
6.2.1	The He I line profiles	129
6.3	Tracing the Colliding winds	131
6.3.1	He I absorption dynamics	131
6.3.2	The line formation regions of Helium	134
6.3.3	A geometrical model for the colliding winds	136
6.4	Discussion	142
6.4.1	A latitudinally dependent wind	142
6.4.2	Model caveats	145
6.5	Summary and conclusions	145
7	Mahalanobis ensembles of co-trending light curves for correcting spatially-correlated photometric artefacts	147
7.1	Introduction	147
7.2	The Algorithm	149
7.2.1	Proof of concept I: symmetric Gaussians	151
7.2.2	Proof of concept II: Mahalanobis ensembles	153
7.2.3	Proof of concept III: intensity correlations	156
7.2.4	Optimization	157
7.3	Synthetic benchmarks	159
7.3.1	Data synthesis	159
7.3.2	Synthetic tests	161
7.4	The Global Jet Watch MC2 data stream	166
7.4.1	The observing campaign	167
7.4.2	Global Jet Watch tests	168
7.5	Conclusions	174
8	Epilogue	175
8.1	Future directions	179
8.2	Open-source software	181
	Bibliography	182

List of Figures

1.1	Hertzsprung–Russell diagram of massive star evolutionary tracks. . .	3
1.2	A schematic of the Global Jet Watch network.	6
1.3	Example Global Jet Watch spectra of emission-line stars.	7
2.1	Schematic of equipotential surfaces illustrating RLOF.	11
2.2	Wind-centred coordinates for computing the line acceleration integral.	15
2.3	Wind velocity profiles for line-driven winds.	18
2.4	Synthesised spectra for P Cgyni and an O-star.	22
2.5	P Cgyni line profile formation schematic.	23
2.6	Diagnostics from the P Cygni CMFGEN simulation.	24
2.7	A schematic of a colliding-wind binary.	26
3.1	A schematic of the two-body problem for binary stars.	31
3.2	A schematic of the different conic sections and orbital types.	32
3.3	A comparison of orbital energies.	33
3.4	A schematic of the orbital elements.	36
3.5	The dependence of radial velocity curves on the orbital parameters.	38
3.6	The dependence of astrometric orbits on the orbital parameters.	41
3.7	Example of the measured quantities from relative astrometry.	42
3.8	Synthetic spectra for benchmarking velocity extraction algorithms.	46
3.9	The bisector algorithm for radial velocity extraction.	48
3.10	The cross-correlation algorithm for radial velocity extraction.	49
3.11	The multi-Gaussian algorithm for radial velocity extraction.	50
3.12	BIC scores for the multi-Gaussian algorithm.	51
3.13	Benchmarking radial velocity extraction algorithms.	52
4.1	Multi-Gaussian fits to Eta Carinae’s Balmer lines near to apastron.	60

4.2	Multi-Gaussian fits to Eta Carinae’s Balmer lines near to periastron.	61
4.3	Keplerian models independently fit to Eta Carinae’s Balmer series. . .	66
4.4	Trends in Keplerian parameters for Eta Carinae’s Balmer series. . . .	67
4.5	CMFGEN inner boundary and clumping tests for Eta Carinae.	71
4.6	CMFGEN wind speed and mass-loss rate tests for Eta Carinae. . . .	71
4.7	Line-formation kernels for Eta Carinae.	73
4.8	Convolutional Keplerian motion modelling schematic.	74
4.9	Eta Carinae’s CKM corner plot.	76
4.10	Eta Carinae’s Balmer series co-fit by the CKM model.	77
4.11	Eta Carinae’s H-delta lines fit by the CKM model.	78
4.12	CKM timescales analysis.	82
4.13	Comparison of the forces acting on Eta Carinae’s primary wind. . . .	86
4.14	The dynamical mass contours of Eta Carinae.	89
4.15	Line-formation kernels for RMC 140b.	91
4.16	Forward models of CKM in RMC 140b.	91
4.17	Line-formation kernels for GG Carinae.	93
4.18	CKM model fit to GG Carinae’s optical emission-lines.	93
5.1	A visual representation of Gaussian processes for orbital modelling. .	105
5.2	Probabilistic orbital algorithms synthetic case study 1.	109
5.3	Probabilistic orbital algorithms synthetic case study 2.	110
5.4	Marginalised $\mathcal{GP}$ algorithm results for Eta Carinae.	115
5.5	Marginalised $\mathcal{GP}$ algorithm results for GG Carinae.	118
5.6	Marginalised $\mathcal{GP}$ algorithm results for WR 140.	121
5.7	Marginalised $\mathcal{GP}$ algorithm results for WR 133.	124
6.1	Selected He I 7065 Å spectral line profiles.	130
6.2	Multi-Gaussian fits to Eta Carinae’s He I 7065 Å line profiles. . . .	132
6.3	Dynamics extracted from Eta Carinae’s He I 7065 Å line profiles. . .	133
6.4	Helium ionisation schematic of Eta Carinae.	135
6.5	Exploration of the geometrical model for the colliding winds.	138
6.6	Best fit of the geometrical model for the colliding winds.	141
6.7	Three-dimensional renderings of the model’s hyperboloid surfaces. . .	144

7.1	Examples of spatially-correlated photometric artefacts.	148
7.2	A visual proof of concept of the 1D Mahalanobis algorithm.	152
7.3	A visual proof of concept of the 2D Mahalanobis algorithm.	155
7.4	Optimization tests of the Mahalanobis algorithm.	158
7.5	Light-curve NRMSD vs source magnitude for synthetic datasets. . . .	163
7.6	Corrected light curves for selected synthetic transients.	164
7.7	Map of the sources detected in the GJW MC2 datastream.	167
7.8	Light curve NRMSD vs source magnitude for GJW datasets.	170
7.9	Light curves for selected GJW transients in the SS433 field.	171
7.10	Light curves for selected GJW transients in the GG Carinae field. . .	172

List of Tables

2.1	Stellar atmosphere model parameters for P Cgyni and an O-star. . . .	21
3.1	The three-dimensional orbital elements of a binary star system. . . .	43
3.2	Calculable dynamical mass quantities given the available data. . . .	44
4.1	Eta Carinae's orbital parameters derived from CKM.	80
5.1	Probabilistic orbital algorithms synthetic results, white noise dataset.	107
5.2	Probabilistic orbital algorithms synthetic results, polynomial dataset.	108
5.3	Probabilistic orbital algorithms synthetic results, GP dataset. . . .	108
5.4	Marginalised $\mathcal{GP}$ algorithm results for Eta Carinae.	114
5.5	Marginalised $\mathcal{GP}$ algorithm results for GG Carinae.	117
5.6	Marginalised $\mathcal{GP}$ algorithm results for WR 140.	120
5.7	Marginalised $\mathcal{GP}$ algorithm results for WR 133.	123
6.1	Geometric model results for the colliding winds of Eta Carinae. . . .	140
7.1	Mahalanobis algorithm synthetic results, datasets 1 and 2.	161
7.2	Mahalanobis algorithm synthetic results, datasets 3 and 4.	162
7.3	Mahalanobis algorithm synthetic results, intensity correlations test. .	166

Chapter 1

Introduction

When a photon – a quantum of light – is radiated from a star, it can travel indefinitely across the vast expanse of space. Photons like these start their journeys as a product of nuclear fusion in the cores of stars, diffusing outwards in a random walk, until the optical pathway ahead becomes sufficiently rarefied, and escape is possible. Key to astrophysics is that these photons encode information about the conditions from their final interaction point with the star. In this way, we are able to probe the physics of these stellar laboratories from millions and billions of light years away, and owing to the finite speed of this information flow, c , our analysis looks backwards in time.

At both ground-based and space-based observatories, astronomers intercept photons measuring their distributions in wavelength (spectroscopy), position (astrometry), and polarisation (polarimetry). Further to this, time-series observations have become integral to the analysis of stellar systems, enabling astronomers to search for, and characterise, their time-evolving behaviour. Optical time-series analysis is booming today, with satellites such as the Transiting Exoplanet Survey Satellite (TESS; Jenkins et al. 2016), surveying hundreds of thousands of light curves of potential exoplanet hosting stars, and the Vera Rubin Observatory’s Legacy Survey of Space and Time (LSST; LSST Science Collaboration et al. 2009) which is due to begin surveying all types of stellar systems across the entire sky within the next two years, and is expected to generate up to about 10 million alerts of variable phenomena per night. Paramount to the success of these campaigns is the innovative work into creating algorithms that distill and discern the astrophysical and statistical information.

This thesis is chiefly concerned with novel time-series analysis of stellar systems. Briefly, stars are formed from collapsing clouds of gas where the gravitational energy

is converted into heat and, once a sufficient temperature and pressure is reached, the fusion of hydrogen is initiated within the core. Any further collapse is then halted, as the inwards gravitational force becomes balanced with the outwards force due to the pressure gradient, known as hydrostatic equilibrium. Whilst a star burns hydrogen in its core, it is said to be on the main sequence, a befitting name as most stars spend the majority of their lives in this phase. The length of a star's life and how it evolves after the main sequence is largely determined by its initial mass, with higher mass stars burning their fuel faster and brighter. The mass-luminosity relationship for main sequence stars is approximately $L \propto M^4$, and this results in a relationship for the main sequence lifetimes of $t \propto M^{-3}$.

1.1 Emission-line stars

The types of stellar system that are of particular interest to our work are those that are comprised of emission-line stars. An emission-line star is categorised by its optical spectrum exhibiting emission lines. For the purpose of this thesis, we concentrate on the most luminous and massive of these stars; specifically, Luminous Blue Variables (LBVs) and Wolf-Rayet (WR) stars. In these systems, the presence of emission lines is indicative of powerful stellar winds outflowing from a central star. However, emission lines may have other origins such as stellar flares, pulsations, discs, mass exchanges and accretion flows (Kogure & Leung 2010).

Stellar winds are driven prominently by spectral lines absorbing the photon momentum from the radiation of these extremely luminous stars (Castor et al. 1975). The accelerated gas from these massive stars may reach terminal velocities between $\sim 100\text{--}3000 \text{ km s}^{-1}$ and the mass-loss rate may range from $\dot{M} \sim 10^{-9} \text{ M}_{\odot} \text{ yr}^{-1}$, for the weak winds of late O-type and early B-type stars (Chlebowski & Garmany 1991; Puls et al. 2008), to $\dot{M} \sim 10^{-4} \text{ M}_{\odot} \text{ yr}^{-1}$ for the powerful winds of LBVs and WR-stars (Crowther 2007; Smith 2014). These outflows are observable in UV and optical spectra, due to photons being both emitted from within the outflow and absorbed or scattered from our line of sight. This results in the well-known P Cygni profile, showing blue-shifted absorption relative to the main emission: a telltale sign for the presence of a stellar wind.




Figure 1.1: Hertzsprung–Russell diagram of massive star evolutionary tracks according to the models of Ekström et al. (2012).

LBVs and WR-stars both represent post main-sequence phases of a star’s lifetime, in which they are losing a considerable amount of mass through their stellar winds. For the majority of massive stars, $\sim 8\text{--}25\,M_{\odot}$, when they evolve off the main sequence the general paradigm is that they increase in radius, cool to lower effective temperatures, and start fusing progressively heavier elements in their cores. These evolved stars are known as red supergiants (RSGs). However, for more luminous and massive stars, $\gtrsim 25\,M_{\odot}$, they may evolve along a different track, remaining hot, and experiencing periods of extreme mass loss. LBVs and WR-stars are thought to be stars observed during these extreme phases. Following Crowther (2007) and Maeder & Meynet (2010), each of these evolutionary tracks can be summarised:

$$\sim 8 - 25\,M_{\odot} : \text{O} \longrightarrow \text{RSG} \longrightarrow \text{SN}, \quad (1.1)$$

$$\gtrsim 25\,M_{\odot} : \text{O} \longrightarrow \text{LBV} \longrightarrow \text{WR} \longrightarrow \text{SN}. \quad (1.2)$$

These tracks are depicted as a function of stellar effective temperature and luminosity in Figure 1.1. The LBV phase occurs after the O-star exits the main-sequence to shed

large amounts of mass, resulting in the outer layers of hydrogen being mostly stripped away leaving a WR-star. These tracks represent the essence of the original proposals by Conti (1976); however, there are several variations that are worth noting. For example, the RSG phase may be followed by the star temporarily evolving back to hotter temperatures in a blue loop. As for the more massive track, it may be complicated by different WR phases: hydrogen-rich and hydrogen-poor WN (strong lines of helium and nitrogen), WC (helium and carbon), and WO (helium and carbon and oxygen) (Crowther 2007). Understanding mass loss and linking the different phases of stellar evolution together is still an active area of research, all the more so when binarity is included.

Once an evolved star has finished utilising Helium as its primary fuel, it starts to burn heavier elements in turn. Once iron is reached, nuclear fusion may no longer be able to support the core against gravity, potentially leading to a core-collapse supernova (SN). The outer layers of the star may be violently blown into the interstellar medium, along with heavy elements, and the remnant left behind may form a neutron star or a black hole. These stars, therefore, represent a major source of feedback into their host galaxies, by way of radiation and chemical enrichment, from stellar winds and SNe.

1.2 Binary stars

Stars are often found in binary or multiple-star systems, unlike our Sun which is just a single star system. In fact, massive O-stars, the main-sequence progenitors to LBVs and WRs, are almost exclusively formed in systems comprised of two or more stellar components (Sana et al. 2014; Moe & Di Stefano 2017). Binary stars are often classified by the method in which they are discovered. These classes are as follows: spectroscopic binaries, detected from the Doppler shifting of spectral lines; eclipsing binaries, detected from repeated drops in measured brightness as the stars pass in front of one another; and visual binaries, detected from relative or absolute position changes in the plane of the sky in interferometric images. Binary systems can be detected by one of several of these methods, and if a system is both a spectroscopic and visual binary, for example, the dynamical masses of the stars can be calculated.

For massive stars evolving and expanding in binaries, they will eventually start to interact with each other if the orbit is not too wide. These inevitable interactions can significantly alter the observables, as well as the physics, and hence the evolutionary tracks of stellar systems (Sana et al. 2012). The main process that occurs is the exchange of mass and angular momentum, although in some cases the stars may even merge. The many processes that occur in binary systems mean that the light we measure encodes information that is multi-faceted. For example, for binary stars with powerful winds, or that are exchanging mass, the orbital velocities encoded in the spectra become entangled with emission from, and motion of, the outflowing gas. These types of scenario make inferring the fundamental properties of binaries, and thus predicting their evolutionary paths, challenging.

1.3 The Global Jet Watch

In order to accomplish the research aims of this thesis, we make use of the Global Jet Watch datastreams. The Global Jet Watch is a network of five remotely controlled observatories separated in longitude to enable round-the-clock observing. The observatories are situated in South Africa, Chile, the east and west coasts of Australia, and India. In Figure 1.2, we display a sketch of the observatories, instrumentation, telemetry, and datastreams. Full descriptions of the observatories and instruments are detailed in Blundell et al. (in prep) and Lee et al. (in prep).

Each observatory has a fibre-fed spectrograph, from a 0.5m telescope, that collects mid-resolution optical data in the wavelength range $\sim 5800 - 8400 \text{ \AA}$. The resolving power is $R = \lambda/\Delta\lambda \approx 4000$, which corresponds to the minimum difference in wavelengths that can be distinguished, $\Delta\lambda$, at a given wavelength, λ . The spectra are reduced by a bespoke data-reduction pipeline which includes dark correction, flat-field correction, and wavelength calibration derived from Thorium Krypton frames. In addition to the spectrographs, each observatory has two photometry cameras that piggyback on the main telescope. These cameras benefit from the pointing, tracking, and guiding of the telescope but have separate optical pathways. The cameras capture 5.6 by 4.5 degree images of the sky onto cooled 2749 by 2199 pixel CCDs. Each camera is equipped with a filter, one centred at 656 nm and the other at 736 nm, and they operate in tandem. These cameras take predominately 300 second exposures,




Figure 1.2: A schematic of the Global Jet Watch observatories, instrumentation, telemetry, and datastreams.




Figure 1.3: Example Global Jet Watch spectra of several emission-line stars: P Cygni, Eta Carinae, and GG Carinae. Top panel displays the most prominent emission line, H-alpha, while in the bottom panel the spectra have been normalised by their local continuum and we have zoomed in on all the lines weaker than H-alpha.

but shorter modes are available for particularly bright targets of interest. All observatories also collect meteorological, electrical, and surveillance data to support the observing operations.

At the remote observatory servers the data are compressed before being made available on the Global Jet Watch network. In Oxford there is an ecosystem of dedicated servers for streaming the latest data from each observatory, storing backups, and reporting real-time metrics about the observatory conditions and operations. After the data are decompressed at the Oxford computing hub it is triaged for reduction depending on the data type. Final data products include 1d spectroscopic data (see examples in Figure 1.3) and calibrated light curves. The superb time-resolution, in combination with either agile or sustained observing modes, makes the Global Jet Watch an ideal resource for studying the orbits and outflows of emission-line stars.

1.4 Thesis overview

In this thesis we aim to study the most luminous and massive emission-line stars, especially those that are found in binaries. Through analysing several key case studies, we innovate new methods for characterising the fundamental parameters of emission-line stars. We apply these methods to infer the properties of our case studies and highlight their wider applicability and significance.

The outline of this thesis is as follows. In Chapters 2 and 3, we introduce the theory for describing the outflows and the orbits of massive stars. We include descriptions of the numerical methods for synthesising and modelling the spectra of these systems, as well as current state-of-the-art algorithms for detecting and characterising their orbits. In Chapters 4 and 5, we tackle the systematic deviations away from Keplerian motion observed in the orbits of emission-line binaries, first with a physical model, and second with a statistical model. Here, we utilise the famous LBV Eta Carinae, the B[e] supergiant GG Carinae, and several WR-stars: WR 140, WR 133, and RMC 140. In Chapter 6, we investigate the colliding winds of Eta Carinae. We use observations of Helium to trace the projected velocities from the colliding-wind interface and formulate a geometric model for characterising the orientation and outflows present. In Chapter 7, we switch data types from spectroscopy to photometry. We present an algorithm for correcting spatially-correlated artefacts in the images of a newly commissioned datastream, enabling the improved photometric analysis of hundreds of thousands of stars. Finally, in Chapter 8 we summarise our findings.

Chapter 2

Stellar outflows

In this chapter we introduce the theory for describing the outflows of massive stars. The derivations follow the seminal descriptions of Castor (1970; 1974), Castor et al. (1975), and Cranmer & Owocki (1995). We then cover the numerical methods for synthesising and modelling the spectra of these systems. In particular, we look at the differences between single stars and binary stars, with regard to the conditions and morphology of their stellar winds.

2.1 Types of mass loss

Stellar mass loss occurs in a variety of different forms. The primary steady mode of mass loss comes from stellar winds. These winds are fairly constant outflows of gas from the surface of a massive star. The mass-loss rate for a spherically symmetric, stationary wind is

$$\dot{M} = 4\pi r^2 \rho(r) v(r), \quad (2.1)$$

where ρ is the density of the gas, v is the wind velocity, r is the radius from the star, and stationary means that the wind is not time-variable. Typical values for the mass-loss rate from these massive stars range from $\dot{M} \sim 10^{-9} \text{ M}_{\odot} \text{ yr}^{-1}$, for the weak winds of late O-type and early B-type stars (Chlebowski & Garmany 1991; Puls et al. 2008), to $\dot{M} \sim 10^{-4} \text{ M}_{\odot} \text{ yr}^{-1}$ for the powerful winds of LBVs and WR-stars (Crowther 2007; Smith 2014). Observationally mass-loss rates can be measured from the H-alpha line luminosity (Leitherer 1988) as well as excess radio and far infrared emission (Wright & Barlow 1975; Panagia & Felli 1975). However, modern spectral synthesis codes

are now commonly used to find the mass-loss rates by fitting synthetic spectra to observations spanning a significant range of wavelengths.

Stellar winds, although usually stationary, can exhibit time-dependent variability. For example, the S Doradus outbursts of LBVs represent a visual brightening, as the spectral energy distribution shifts from the UV to visual wavelengths, on timescales of months or years. This variability may be due to these stars increasing their mass-loss rates, resulting in larger optical depths in the winds, and creating pseudo photospheres (Davidson 1987; Humphreys & Davidson 1994). This explanation has however been questioned. Spectroscopic data observed during these outbursts reveal that the mass-loss rates only increase by $\sim 40\%$, but a factor of $\sim 4\text{--}12$ is required to reach the minimum effective temperature observed from these pseudo photospheres (De Koter et al. 1996). Therefore, these outbursts may not be responsible for huge amounts of additional mass loss, but rather the photospheric radius of the star is simply inflating, and then decreasing, similar to a pulsation (Smith 2014).

A more pronounced class of outburst, called giant eruptions, is also attributed to the LBVs. These events are very rare and irregular. They produce hugely enhanced mass-loss rates spanning several years, such as Eta Carinae’s giant eruption in the 1840s, where it is estimated to have ejected in excess of $45 M_{\odot}$ of material over about a decade (Morris et al. 2017). For many LBVs, the only way we know of their giant eruptions is from historical light curve data (eg. Smith & Frew 2011; who pieced together hand written data from the 19th century), or the present day circumstellar shells: the fossils of the past outflows (eg. Currie et al. 1996). There are a few examples of observed giant eruptions from nearby galaxies (eg. Van Dyk & Matheson 2012) that were originally thought to be supernovae. Hence, these systems that experience giant eruptions are sometimes called supernova imposters.

Aside from mass-loss processes pertaining to single stars, there are several mechanisms in binary systems that considerably disrupt these single-star outflows. Most important is Roche-lobe overflow (RLOF). RLOF is initiated when the expanding photospheric radii of massive stars reaches the surface in which orbiting material is no longer gravitationally bound to themselves. A schematic of this occurrence is shown in Figure 2.1. The mass then transfers through the L1 Lagrangian point to the companion star. The timing and rate of RLOF is largely determined by the or-




Figure 2.1: Schematic of equipotential surfaces (black lines) relative to the radii (orange shapes) and separations of detached (left) and semi-detached (right) systems. In semi-detached systems Roche-lobe overflow is initiated through the L1 Lagrangian point, with mass transferring from the primary to the companion star.

bital separation and mass ratio. The onset may begin when the primary is on the main sequence (Case A), H-shell burning (Case B), or post core-He burning (Case C), with the mass transfer rate increasing from Case A to Case C, and as the mass ratio moves further from unity (Podsiadlowski et al. 1992; Smith 2014). Although the rates vary widely, it is typical for values to be on the order of $10^{-3} \text{ M}_{\odot} \text{ yr}^{-1}$. Given that binary and multiple-star systems are more prevalent at high stellar masses (Mason et al. 2009; Sana et al. 2014), and the high rates of probable mass transfer, binary interaction may actually dominate the evolution of these systems (Sana et al. 2012). However, for systems with long orbital periods, and thus wide separations, they may never interact. For example, a massive RSG may expand to $\sim 100\text{s}$ of $R_{\odot}$, and if its orbital period was on the order of $\sim 100\text{s}$ of days, depending on its mass, it may never overflow its Roche lobe. For the remainder of this chapter we focus our attention on stellar winds, in both the single and binary star contexts.

2.2 Wind acceleration

The wind outflowing from a star is subject to several components of acceleration. The total acceleration along the radially flowing streamlines, g_r , can be written as the sum of

$$g_r = -g + g_{\text{rad}}^{\text{C}} + g_{\text{rad}}^{\text{L}}, \quad (2.2)$$

where g is the gravitational acceleration, $g_{\text{rad}}^{\text{C}}$ represents the acceleration from continuum driving, and $g_{\text{rad}}^{\text{L}}$ represents the acceleration from line driving. All three of these forces can be further deconstructed for either an isolated star, or a star within a binary. Let us consider the binary case. The gravitational acceleration term is straightforward to calculate. Considering the flow along the line between the two stars, the gravitational acceleration of the gas is

$$g = -\frac{Gm_1}{r^2} + \frac{Gm_2}{(r - D_{\text{sep}})^2}, \quad (2.3)$$

where m_1 and m_2 are the masses of the primary and companion star respectively, D_{sep} is the binary separation distance, G is the universal gravitational constant, and r is the radius from the star.

The second and third terms on the right-hand side of Equation 2.2 depend on the radiation field throughout the winds. The fundamental quantity of radiative transfer is the specific intensity

$$I_\nu = \frac{dE}{dA d\Omega dt d\nu \cos \theta}, \quad (2.4)$$

which characterises the energy flow, dE , of a light ray passing through a surface area dA , from a solid angle $d\Omega$, at an angle perpendicular to the surface θ , over a given time dt , and frequency $d\nu$. Using this definition, we can write the equation that governs the change in specific intensity along a path, the radiative transfer equation:

$$\frac{dI_\nu}{ds} = -\kappa_\nu \rho I_\nu + j_\nu, \quad (2.5)$$

where ds is a path element, κ_ν is the mass absorption coefficient, and j_ν is the emission coefficient. κ_ν is comprised of both absorption plus scattering processes and is often referred to as the opacity. In this derivation we denote $[\kappa_\nu] = \text{cm}^2\text{g}^{-1}$. The radiation field from one star is able to transfer momentum from its photons to its wind in both the continuum, through the scattering of electrons, and in the spectral lines, through predominately the absorption and scattering of UV metal lines (Lucy 1970). The radiative acceleration at a point, r , in the radiation field is given by

$$\mathbf{g}_{\text{rad}} = \frac{1}{c} \oint \int_{\nu=0}^{\infty} \kappa_\nu I_\nu(\mathbf{r}, \hat{\mathbf{n}}) \hat{\mathbf{n}} d\Omega d\nu, \quad (2.6)$$

where $\hat{\mathbf{n}}$ is a unit vector in the direction of the radiation flow.

For the continuum acceleration, assuming the wind is optically thin to continuum radiation, which is a good approximation for O and B stars, we can solve Equation 2.6 for the single star case, finding

$$g_{\text{rad}}^{\text{C}} = \frac{\sigma_e}{c} \oint I^{\text{C}}(\hat{\mathbf{n}}) \hat{\mathbf{n}} d\Omega = \frac{\sigma_e F}{c}, \quad (2.7)$$

where σ_e is the specific electron opacity, I^{C} is the continuum intensity from the stellar surface, and F is the continuum bolometric flux. For binary stars, along the line between the two stars, the continuum acceleration of the gas is

$$g_{\text{rad}}^{\text{C}} = \frac{\sigma_e L_1}{4\pi c r^2} - \frac{\sigma_e L_2}{4\pi c (r - D_{\text{sep}})^2}, \quad (2.8)$$

where L_1 and L_2 are the continuum luminosities of the primary and companion stars, respectively. Let us compare the continuum acceleration to the gravitational acceleration. The maximum luminosity, beyond which the radiative force in the continuum will overcome gravity, is known as the Eddington luminosity, defined as

$$L_{\text{Edd}} = \frac{4\pi G m c}{\sigma_e}, \quad (2.9)$$

where m is the mass of the star. We can rearrange this equation for the so-called Eddington factor, $\Gamma = \sigma_e L / (4\pi c G m)$, which represents the fraction of L_{Edd} obtained in the stellar wind. Following the calculations of Russell (2013), utilising typical values for an O-star's luminosity, $L = 3 \times 10^5 L_{\odot}$, mass, $m = 30 M_{\odot}$, and assuming a completely ionised gas and hydrogen mass fraction of $X = 0.7$, we find these stars only achieve values in the range of $\Gamma = 0.26$. Hence, it is not sufficient for the continuum opacity alone to be responsible for driving the stellar winds.

Line opacity provides a far stronger source of acceleration, relative to the continuum, in the winds of massive stars. For lines accelerated by the radiation field from a single star it may seem that this would not be the case. In an ionised gas the vast majority of atoms do not have bound electrons and the radiation field can only interact with lines at their specific transition wavelengths, severely limiting the flux available to drive the wind. However, the accelerating nature of the wind results in a single line transition being progressively Doppler shifted to redder wavelengths, relative to the star's radiation field as it flows outwards. Consequently, every line transition can interact with a much greater stellar flux and the result is a net stronger line driving force.

Following Castor et al. (1975; hereafter known as CAK theory) and the work of Cranmer & Owocki (1995), let us consider the wind acceleration from a single line, i , in a single star, which is given by

$$\mathbf{g}_{\text{rad}}^{\text{L},i} = \frac{\kappa_{\text{L},i} \nu_i v_{\text{th}}}{c} \oint \mathbf{I}^{\text{C}}(\hat{\mathbf{n}}) \left(\frac{1 - e^{-\tau_{\nu,\text{Sob}}}}{\tau_{\nu,\text{Sob}}} \right) \hat{\mathbf{n}} d\Omega, \quad (2.10)$$

where $\kappa_{\text{L},i}$ is the mass absorption coefficient for a specific line transition, ν_i is the line's central frequency, v_{th} is the thermal speed of the gas, and $\tau_{\nu,\text{Sob}}$ is the Sobolev optical depth. The Sobolev optical depth is defined as

$$\tau_{\nu,\text{Sob}} = \kappa_{\text{L},i} \rho L_{\text{Sob}} = \kappa_{\text{L},i} \rho(\mathbf{r}) \frac{v_{\text{th}}}{\hat{\mathbf{n}} \cdot \nabla(\hat{\mathbf{n}} \cdot \mathbf{v}(\mathbf{r}))}, \quad (2.11)$$

where $\rho(\mathbf{r})$ and $\mathbf{v}(\mathbf{r})$ are the density and velocity of the wind at location r in the wind, and L_{Sob} is the Sobolev length. This length characterises the overlap between the thermally-broadened resonance zones of the lines, and the overall Doppler shift of the lines from within the accelerating, expanding wind. In particular, the Doppler shifts of the lines dominate, and thus there is only one resonant velocity per light ray path. This Sobolev approximation is key to localising and solving the optical depth integral that leads to the derivation of Equation 2.10.

Next, the total line acceleration is calculated from the sum of the accelerations from an ensemble of individual lines. The key insight from CAK theory is that the line list can be approximated as a power law, which itself can be formulated as a number distribution of lines, N , as a function of line opacity. This is written as

$$\frac{dN}{d\kappa_{\text{L}}} = \frac{1}{\kappa_0} \left(\frac{\kappa_{\text{L}}}{\kappa_0} \right)^{\alpha-2}, \quad (2.12)$$

where κ_0 is defined via

$$k = \frac{\Gamma(\alpha)}{1-\alpha} \left(\frac{\kappa_{v_{\text{th}}}}{c} \right) \left(\frac{\kappa_0}{\sigma_e} \right)^{1-\alpha}. \quad (2.13)$$

In these equations $\Gamma(\alpha)$ is the complete gamma function, and k and α are the standard CAK parameters. k is related to the number of lines contributing to the line acceleration, whilst α determines the distribution of optical depths within the ensemble of lines and takes a value in the interval $0 < \alpha < 1$. For $\alpha = 0$, all the lines are considered optically thin, and for $\alpha = 1$, all the lines are considered optically thick.




Figure 2.2: The wind-centred spherical coordinate system utilised in computing the line acceleration integral for geometries involving a finite stellar disc. $\mathbf{g}_{\text{rad}}$ is the net acceleration of the gas situated at the origin, accelerated by the radiation flow in the direction $\hat{\mathbf{n}}$. The symbols z , θ' , and ϕ' represent the radial distance, azimuthal angle, and polar angle respectively. R is the radius of the star.

The overall line acceleration from a single star can now be formally written, after quite some algebra, as a solvable integral with the form

$$\mathbf{g}_{\text{rad}}^{\text{L}} = \sum_{\text{lines}, i} \mathbf{g}_{\text{rad}}^{\text{L}, i} = \int_0^\infty \mathbf{g}_{\text{rad}}^{\text{L}, i} \left(\frac{dN}{d\kappa_{\text{L}}} \right) d\kappa_{\text{L}}, \quad (2.14)$$

$$= \frac{\sigma_e^{(1-\alpha)} k}{c} \oint \mathbf{I}^{\text{C}}(\hat{\mathbf{n}}) \left(\frac{\hat{\mathbf{n}} \cdot \nabla(\hat{\mathbf{n}} \cdot \mathbf{v})}{\rho(\mathbf{r}) v_{\text{th}}} \right)^\alpha \hat{\mathbf{n}} d\Omega, \quad (2.15)$$

$$= \frac{\sigma_e^{(1-\alpha)} k}{c} \oint I^{\text{C}}(\theta', \phi') \left(\frac{\hat{\mathbf{n}} \cdot \nabla(\hat{\mathbf{n}} \cdot \mathbf{v})}{\rho(\mathbf{r}) v_{\text{th}}} \right)^\alpha \hat{\mathbf{n}} \sin \theta' d\theta' d\phi'. \quad (2.16)$$

In Equation 2.16, we have chosen spherical coordinates centred on the position in the wind we are considering, and we have orientated the z -axis along the radial direction from the star to this position. This coordinate system, in the wind-centred frame, is illustrated in Figure 2.2. These integrals, despite looking like quite a handful, are actually fairly tractable, especially in the case of some simplified geometries. In the simplest case where the radiation field is emanating from a point source, the solution neatly becomes $\mathbf{g}_{\text{rad}}^{\text{L}} = \mathbf{g}_{\text{rad}}^{\text{C}} k t_{\text{CAK}}^{-\alpha}$, where $t_{\text{CAK}} = (\sigma_e / \kappa_{\text{L}, i}) \tau_{\nu, \text{Sob.}}$. This concise result was the intention behind the original CAK theory: acceleration from an ensemble of lines can be written as the continuum acceleration modified by a force multiplier $\mathfrak{M}(t_{\text{CAK}}) = k t_{\text{CAK}}^{-\alpha}$. The values of $\mathfrak{M}$ at incremental depths in the wind are tabulated from calculations involving extensive line lists and after making several assumptions, such as local thermodynamic equilibrium, illumination by a Planck function, and

the Sobolev approximation. The parameters k and α are then found by fitting a power law to $\mathfrak{M}(t_{\text{CAK}})$ (eg. Abbott 1982), although the most recent line lists deviate from a power law sufficiently such that alternative functions may provide better fits (Lattimer & Cranmer 2021).

While the point star approximation provides a reasonable approximation to the line acceleration at large radii, since then, more realistic geometries that provide more accurate solutions close to the star have been developed (eg. Pauldrach et al. 1986). Here we consider the radiation field from a finite stellar disk with a linear limb darkening law. In this case it is helpful to define a change of variables, $\mu' = \cos \theta'$. For a purely radially streaming wind, the projected velocity gradient becomes

$$\hat{\mathbf{n}} \cdot \nabla(\hat{\mathbf{n}} \cdot \mathbf{v}(\mathbf{r})) = \mu'^2 \frac{\partial v_r}{\partial r} + (1 - \mu'^2) \frac{v_r}{r}, \quad (2.17)$$

and the bolometric intensity can be written as

$$I^{\text{C}}(r, \theta', \phi') = \frac{L}{4\pi R^2} D(\mu', r), \quad (2.18)$$

where the limb darkening function, $D(\mu', r)$, follows

$$D(\mu', r) = \begin{cases} 0, & -1 \geq \mu' < \mu_* \\ (2 + 3\mu'')/4\pi, & \mu_* < \mu' < +1 \end{cases}. \quad (2.19)$$

In these equations, θ' is the wind-centred angle, θ'' is the limb angle, θ_0 is the star-centred angle, and R is the radius of the star, each of which is illustrated in Figure 2.2. Additionally, $\mu_*(r) = \sqrt{1 - R^2/r^2}$ and $\mu'' = \cos \theta'' = \sqrt{(\mu'^2 - \mu_*^2)(1 - \mu_*^2)}$. Using this framework we can now write down the radial wind acceleration from an ensemble of lines, for a single star, in terms of the following integral:

$$\mathbf{g}_{\text{rad}}^{\text{L}} = \frac{2\pi\sigma_e^{(1-\alpha)}k}{c(\rho(r)v_{\text{th}})^\alpha} \left(\frac{L}{4\pi R^2} \right) \oint_{\mu_*}^{+1} D(\mu', r) \left(\mu'^2 \frac{\partial v_r}{\partial r} + (1 - \mu'^2) \frac{v_r}{r} \right)^\alpha \mu' d\mu'. \quad (2.20)$$

It is common in the literature to use the succinct point star approximation, but with a finite disc correction factor, $K = g_{\text{rad}}^{\text{L}}(\text{finitedisc})/g_{\text{rad}}^{\text{L}}(\text{pointsource})$. For binary stars, along the line between the two stars, the line acceleration of the gas becomes

$$g_{\text{rad}}^{\text{L}} = \frac{\sigma_e}{c} (\mathfrak{M}_1(t_{\text{CAK}})F_1K_1 - \mathfrak{M}_2(t_{\text{CAK}})F_2K_2), \quad (2.21)$$

where $\mathfrak{M}_1$ and $\mathfrak{M}_2$, F_1 and F_2 , and K_1 and K_2 are the CAK force multipliers, bolometric fluxes, and finite disc correction factors of the primary and companion stars,

respectively. Note that the force multipliers are in general different between the two stars, and their values will affect the coupling between the two stars' radiation fields and winds, which can impact the resulting outflows (Parkin et al. 2011).

The integral in Equation 2.20 cannot in general be solved analytically, and so we use numerical Gaussian integration to compute the line acceleration. The acceleration resulting from gravity, the continuum (Equation 2.8), and an ensemble of lines (Equation 2.20) can be input into Equation 2.2 to compute the overall acceleration of a stellar wind. Further complexities can be built in to the formalism, such as the effect of rapid rotation on the stellar disc geometry and velocity streamlines (Cranmer & Owocki 1995; Owocki et al. 1996), the effects of companion radiation fields (Stevens & Pollock 1994), and tidal forces from close-by companions (Hirai & Mandel 2021). The resulting acceleration can be used to study the dependence of the mass-loss rates of stars on the stellar parameters such as luminosity, temperature, ionisation, and metallicity (eg. Vink et al. 1999; Mokiem et al. 2007). The acceleration integral can also be directly incorporated into simulations of massive stars to study the wind hydrodynamics near to stellar companions (eg. Parkin et al. 2011) or black holes (Hirai & Mandel 2021). In Section 4.4.2.3, we utilise this integral to analyse the strength of the line acceleration at various radii in the wind of Eta Carinae relative to its stellar companion.

2.3 Single star flows

After deriving the equations for computing the wind acceleration in massive stars, we now consider the resulting outflow driven by these forces. In the case of a single star, the wind momentum equation is written as

$$v \frac{dv}{dr} = -\frac{1}{\rho} \frac{dp}{dr} + g_r, \quad (2.22)$$

where g_r is from Equation 2.2 and p is the pressure. This equation is simply an extension of hydrostatic equilibrium, in which the pressure stratification in a star balances the gravitational acceleration inwards, but we have included additional terms for the outwards radiative acceleration and the resulting velocity gradient in the wind. We limit our analysis to time-steady winds, and therefore omit any time dependence from the momentum equation.




Figure 2.3: Wind velocity profiles for line-driven winds. The left and right panels show the dependence of $v(r)$ on the terminal wind velocity and β , respectively.

Using the definition for the Eddington factor, $\Gamma = \sigma_e L / (4\pi c G m)$, the mass-loss rate (Equation 2.1), the point star solution to the line acceleration, and the equation of state, $p = \rho c_s^2$, where c_s is the isothermal sound speed, the wind momentum equation can be written as a differential equation in terms of the wind velocity:

$$\left(1 - \frac{c_s^2}{v^2}\right) r^2 v \frac{dv}{dr} = -Gm(1 - \Gamma) + 2c_s^2 r - r^2 \frac{dc_s^2}{dr} + C_{\text{CAK}} \left(r^2 v \frac{dv}{dr}\right)^\alpha, \quad (2.23)$$

where

$$C_{\text{CAK}} = \Gamma G m k \left(\frac{4\pi}{\sigma_e v_{\text{th}} \dot{M}} \right)^\alpha. \quad (2.24)$$

Equation 2.23 is simplified under the assumption that gas pressure does not play a significant role in highly supersonic winds (Gayley 2000), and in this limit $c_s \rightarrow 0$. From the remaining terms there exists a unique analytical solution for the wind velocity profile.

2.3.1 Wind velocity profiles

The solution to the wind velocity equation (Equation 2.23) has the form

$$v_{\text{CAK}}(r) = v_\infty \left(1 - \frac{R}{r}\right)^{1/2}, \quad (2.25)$$

where $v_\infty = \sqrt{\alpha/(1 - \alpha)} v_{\text{esc}}$ and $v_{\text{esc}} = \sqrt{2Gm(1 - \Gamma)/R}$. More generally, this solution is known as a β -law and is written

$$v(r) = v_0 + (v_\infty - v_0) \left(1 - \frac{R_*}{r}\right)^\beta. \quad (2.26)$$

where v_0 is the initial velocity of the wind. Recall that the $\beta = 0.5$ solution relied on the point star approximation for line acceleration. When considering corrections from a finite stellar disc, the wind accelerates more slowly and to higher terminal velocities. As a result, a value of $\beta = 0.8$ is found for O-type stars (Pauldrach et al. 1986). Further to this, larger values of β are often considered when modelling the line profiles of massive stars. In Figure 2.3 we display various wind velocity profiles as per Equation 2.26. In the left-hand panel we show how $v(r)$ depends on the terminal wind velocity, and in the right-hand panel we show how $v(r)$ depends on β . In all cases we find a monotonically increasing wind velocity that asymptotically reaches its terminal velocity. After 100 stellar radii the wind velocity is within 1% of its maximum for $\beta = 1$.

In addition to describing the wind velocity profile, the CAK theory also leads to a mass-loss rate prediction. The solution to Equation 2.23 necessitates a certain form for the constant C (Castor et al. 1975):

$$C_{\text{CAK}} = \frac{1}{\alpha} \left(\frac{\alpha}{1-\alpha} Gm(1-\Gamma) \right)^{1-\alpha}. \quad (2.27)$$

Through equating this with the formula for C in Equation 2.24, the mass-loss rate can be obtained from

$$\dot{M}_{\text{CAK}} = \frac{4\pi Gm}{\sigma_e v_{\text{th}}} \alpha (1-\alpha)^{(1-\alpha)/\alpha} k^{1/\alpha} \Gamma^{1/\alpha} (1-\Gamma)^{-(1-\alpha)/\alpha}, \quad (2.28)$$

assuming knowledge of all the other stellar parameters. Once again these values are modified when considering the corrections from a finite stellar disc, causing a decrease by a factor ~ 2 in this theoretical mass-loss rate (Pauldrach et al. 1986).

2.3.2 Numerical simulations

Now that we have derived the theory for the structure of stellar winds we can utilise this framework, along with a suitable radiative transfer code, to create synthetic models for the observables of massive stars. Through comparing these models with observed spectra, it is possible to infer many of the fundamental properties of massive stars. However, computing consistent models of both the radiation field and wind stratification is a challenging task. While CAK theory is broadly supported by observations of O and B-stars, the highly stratified winds of emission-line stars

undermine some of the assumptions of the CAK formalism. In particular, photons propagating through the winds of emission-line stars will undergo multiple scattering events. Hence, more realistic treatments have been developed to model the winds of these high mass-loss rate stars, such as Monte Carlo (Abbott & Lucy 1985; Vink et al. 1999; 2000; 2001) and hydrodynamic (Sander et al. 2017) methods. However, most codes avoid these complexities, and do not compute the wind driving, instead requiring prescribed mass-loss rates and velocity fields, within which the radiative transfer models are computed.

To synthesise the emergent spectrum we must tackle a coupled problem within the atmospheres of stars: the radiation field depends on the emission and extinction of the medium, while the emission and extinction of the medium, both themselves, depend on the radiation field. Accordingly, state-of-the-art programs are employed to solve this problem. In this work we make use of the non-local thermodynamic equilibrium, fully line-blanketed stellar atmosphere code **CMFGEN** (Hillier & Miller 1998; Hillier & Lanz 2001). However, in practice any suitable stellar atmosphere code, allowing for non-LTE effects and spherical geometry, could be used and the results should be independent of this choice.

CMFGEN solves the radiative transfer equation in the comoving frame, subject to the constraints of statistical and radiative equilibrium. For statistical equilibrium we require

$$\sum_{j \neq i} = N_i(R_{ij} + C_{ij}) = \sum_{j \neq i} = N_j(R_{ji} + C_{ji}), \quad (2.29)$$

where i and j are individual levels in each model atom for all species present, N_x is the number of electrons populating a given level, and R and C are the transition rates between levels due to radiative and collisional processes, respectively. For radiative equilibrium we require

$$\int_0^\infty \kappa_\nu (J_\nu - S_\nu(T)) d\nu = 0, \quad (2.30)$$

where J_ν is the mean intensity and S_ν is the source function. At each depth in the stellar atmosphere, the temperature, T , must assume a value which satisfies this equation. **CMFGEN** iteratively updates the atomic level populations and temperature structure until a convergence threshold is obtained. After this, the emergent spectrum is computed from the observer's frame.

Model parameters	O-star (Hillier 2011)	P Cygni (Rivet et al. 2020)
L ($L_{\odot}$)	378400	610000
T_{eff} (K)	30000	18700
$\log g$ ($L_{\odot}$)	3.25	2.25
R ($R_{\odot}$)	22.8	75
m ($M_{\odot}$)	33.6	37
$\dot{M}$ ($M_{\odot}\text{yr}^{-1}$)	3.86×10^{-6}	4.0×10^{-5}
f_{∞}	0.1	0.5
f_{vinit} (km s^{-1})	200	20
v_{∞} (km s^{-1})	1877	185
β	0.8	2.3

Table 2.1: Stellar atmosphere model parameters for P Cygni and a typical O-star input to **CMFGEN** to compare the synthetic emergent spectra.

As a demonstrative exercise we compute the synthetic spectrum of the archetypal emission-line star, and LBV, P Cygni, and compare it to a typical O-star. For P Cygni, we follow the model constructed by Rivet et al. (2020), which is itself based on the model of Najarro (2001). For the O-star, we utilise one of the template models from the standard grid of O-star models supplied with the **CMFGEN** distribution (Hillier 2011). The model parameters of these two systems are shown in Table 2.1.

In Figure 2.4, we plot the synthesised spectra of the O-star (top row) and P Cygni (bottom row). The left-hand panels show the emergent spectra at optical wavelengths, and the right-hand panels show the specific morphologies of the He I 5876 Å lines. The O-star spectrum shows absorption lines superimposed onto a continuum distribution that peaks in the UV. These absorption lines form in the photosphere and are Gaussian-shaped having been broadened by thermal and pressure effects in this layer.

In contrast, the P Cygni spectrum shows many emission lines; particularly bright are the Balmer lines that can be seen starting from H-alpha (6563 Å) and becoming incrementally weaker and closer together through the Balmer series to shorter wavelengths. The emission versus absorption difference occurs due to the conditions at




Figure 2.4: Synthesised spectra for P Cygni (bottom row) and a typical O-star (top row) output from **CMFGEN**. Left-hand panels show the emergent spectra at optical wavelengths. Right-hand panels show the specific structures of the He I 5876 Å lines. For comparison the spectra have been scaled such that their maximum values equal one.

the region of the lines' formation. Whilst for absorption, the radiation field passes through the relatively cooler outer layers of the star, in the emission-line case, the stellar radiation heats the dense surrounding wind. This leads to conditions in the wind that cause the source terms to dominate over the sink terms in the radiative transfer equation, at the wavelengths of specific line transitions. As a result, these lines appear in emission.

These emission lines have profiles that originally got their name from this star, P Cygni. In the bottom-right panel, we view the He I 5876 Å line, and include two additional comparison spectral profiles at increasing terminal wind velocities. The




Figure 2.5: Schematic of the formation of P Cygni line profiles. The black arrows indicate the outflowing wind and the white arrows show photon scattering events at various locations in the wind. The wind’s colour corresponds to the line-of-sight projected velocities, which are associated with the formation of each part of the spectral line profiles. The P Cygni profile is the summation of the profiles from the spherical emission region and the absorption along the line-of-sight.

overall P Cygni line profile is comprised of an emission component, from photons scattered into the line of sight from around the star, and a blue-shifted absorption component, from the absorption of photons passing through the expanding gas along our line of sight to the star. A schematic of this process is shown in Figure 2.5. The extent of the blue-shifted absorption indicates the velocity profile of the wind, and can be used to infer the terminal wind velocity.

In Figure 2.6, we display some further diagnostics from the **CMFGEN** simulations for P Cygni. In the left-hand panel we show the volume filling factor and wind velocity as a function of radius. Here we see the β -velocity law derived above used in practice. As for the volume filling factor, this is a standard way of including clumping in stellar atmosphere codes: an effect that develops from the intrinsic instability in line driving (MacGregor et al. 1979; Owocki et al. 1988), which has been observationally substantiated Ebbets (1982); Kaper et al. (1996); Oskinova et al. (2006). Clumping is modelled by introducing optically thin clumps within a void inter-clump medium.




Figure 2.6: Diagnostics from the P Cygni CMFGEN simulation. Left-hand panel shows the volume filling factor, f_v , and wind velocity as a function of radius. Right-hand panel shows the continuum (electron) optical depth, τ_{es} .

The over-density inside the clumps is determined by a so-called clumping factor, f_{cl} , or volume filling factor, f_v , defined by

$$f_{cl} = \langle \rho^2 \rangle / \langle \rho \rangle^2, \quad (2.31)$$

$$f_v = f_{cl}^{-1}, \quad (2.32)$$

where ρ is the density and the angle brackets denote volume averaging. In CMFGEN, the volume filling factor is parameterised in terms of the clumping at infinity, f_∞ , and the velocity at which clumping becomes important, $f_{v_{init}}$. The recognition that winds are clumpy is particularly important for determining the correct mass-loss rates, as diagnostics such as UV P Cygni absorption ($\propto \rho$) and H-alpha emission ($\propto \rho^2$) are affected by the higher densities present. In the right-hand panel of Figure 2.6, we plot the optical depth in the continuum, τ_{es} . Here, we see the value of $\tau_{es} = 2/3$ is crossed beyond a single stellar radius. This indicates that the powerful wind has moved the effective photosphere out into the wind. We also note the flattening of the rate at which the optical depth decreases at larger radii, as the wind velocity starts to asymptote.

Overall, synthesising spectroscopic models, like those demonstrated here, is a key method for determining the properties of single stars. For binary stars, the spectra

would ideally be disentangled, if possible, and then individually analysed with a stellar atmosphere code (eg. Shenar et al. 2019).

2.4 Colliding flows

In binary star systems where both stars drive powerful stellar winds, these winds collide forming shock fronts at the location where their wind momenta balance. These systems, known as colliding-wind binaries, complicate the single-star outflows described in Section 2.3, introducing three-dimensional morphological considerations. This may introduce variability into the UV and optical line profiles during a binary’s orbit (Stevens 1993; Szostek et al. 2012).

In Figure 2.7 we sketch the geometrical picture of a colliding-wind binary. Following Stevens et al. (1992), the location of the apex of the colliding-wind interface is calculated from the one-dimensional balance of wind momentum fluxes (or ram pressures) along the binary axis:

$$\rho_1 v_1^2 = \rho_2 v_2^2, \quad (2.33)$$

where the subscripts, 1 and 2, represent the primary and companion stars, respectively. By defining the stellar separation along the binary axis, D , and using the definition for mass-loss rate from Equation 2.1, we can write

$$\frac{\dot{M}_1 v_1(r_1)}{r_1^2} = \frac{\dot{M}_2 v_2(r_2)}{r_2^2}, \quad (2.34)$$

which becomes

$$r_1 = \frac{D}{1 + \sqrt{\eta_{\text{wm}}}}, \quad (2.35)$$

through defining the momentum ratio

$$\eta_{\text{wm}} = \frac{\dot{M}_2 v_{\infty,2}}{\dot{M}_1 v_{\infty,1}}. \quad (2.36)$$

Here, r_1 is the distance to the apex of the colliding-wind interface from the primary star. Note that these equations assume that the wind has reached its terminal velocity before colliding, which is a valid assumption for wide binaries.

The apex of the colliding-wind interface is rounded, and the interface curves outwards towards the star with the weaker wind. At large radii this interface asymptotes




Figure 2.7: A schematic of a colliding-wind binary. The yellow lines indicate the primary and companion winds' shock fronts, which are separated by a contact discontinuity (black line). θ is the asymptotic half-opening angle. ϕ is the angle between the line of sight and the shock-cone intersection. The blue, green, and red lines represent contours of constant primary wind velocity that has been projected onto the line of sight, v_{LOS} , in fractions of the terminal velocity v_{∞} . This figure is adapted from Szostek et al. (2012).

to a steady opening angle, resembling a cone-like geometry. This shock cone is parameterised in terms of its asymptotic half-opening angle, θ , as seen in Figure 2.7. The value of θ depends on the type of shocks. For radiative shocks, Canto et al. (1996) calculated θ by considering only the momentum flux balance, resulting in the formula

$$\eta_{\text{wm}} = \frac{\tan \theta - \theta}{\tan \theta - \theta + \pi}. \quad (2.37)$$

However, for adiabatic shocks the heat generated is converted purely into work done by the gas flowing within the shock cone. Gayley (2009) calculated θ for several different regimes. When there is no post-shock mixing and each wind returns adiabatically

to its original terminal speed, the formula is

$$\eta_{\text{wm}} = \tan^4 \left(\frac{\theta}{2} \right). \quad (2.38)$$

The angles from this equation are slightly wider, at fixed η_{wm} , than for radiative shocks. And when considering regimes in which the post-shock gas can mix, the adiabatic shock angles become slightly wider again (Gayley 2009).

In addition to the location and angles associated with the shock-cone, in Figure 2.7 we also show how the colliding winds impact the primary wind. For companions with weaker, more rarefied winds, the shock cone forms a cavity within the outflows of the primary star. This can significantly impact the formation of emission-lines, especially lines that may be formed at large radii such as H-alpha. Shown are the contours of constant line-of-sight velocities in the outflowing wind of the primary star. As the companion orbits, it will sweep out different portions of these velocities, leading to variability in the line profiles. We also label the angle, ϕ , made between the line of sight and the contact discontinuity: the interface that separates the post-shock gas of each star. The angle and location of this intersection is important to the formation of absorption lines, and we make use of these ideas in Chapter 6.

There are several further complexities to note which are not shown in our simple schematic of a colliding-wind binary in Figure 2.7. First, the orbital motion results in the shock-cone experiencing a Coriolis force within its rotating frame of reference. The Coriolis effect acts to deflect the shock cone from being symmetric about the binary axis, and instead the shock-cone lags behind the orbital motion. The extent to which this impacts the geometry of the colliding winds depends on the orbital parameters of the system, where in general faster rotation and larger length scales produce greater deflections. Second, the flow is not necessarily smooth. The density contrast and velocity shear across the contact discontinuity lead to Rayleigh–Taylor and Kelvin–Helmholtz instabilities, respectively. Third, when the separation between binary stars is small, the radiation field of one star may affect the velocity profile of the other star’s wind by radiative inhibition (Stevens & Pollock 1994) or braking (Gayley et al. 1997). Radiative inhibition reduces the initial acceleration of the other star’s stellar wind, whilst radiative braking suddenly decelerates the wind as it approaches the companion. In addition, depending on the ratio of luminosities and masses, the

counterveiling effect of gravity may increase the wind velocities and mass-loss rate (Stevens & Pollock 1994).

2.4.1 High energy radiation

The presence of the shocks at the colliding-wind interface may act as an additional source of radiation. The shocks compress and heat the gas, leading to the production of X-rays (Prilutskii & Usov 1976; Cherepashchuk 1976; Luo et al. 1990), and non-thermal emission (Pollock 1987; White & Chen 1994; Hamaguchi et al. 2018), as well as dust (Williams 2008).

The standing shocks cause an instantaneous jump in conditions. If we define the pre-shock, upstream conditions in terms of the density, ρ_u , velocity, v_u , and temperature, T_u , then the corresponding jump in these variables results in the post-shock, downstream conditions of $\rho_d > \rho_u$, $v_d < v_u$, and $T_d > T_u$. In particular, as the wind kinetic energy is converted into heat the temperature can reach millions of degrees, and therefore, will emit thermal Bremsstrahlung and line emission as X-rays. These X-rays can be utilised to identify colliding-wind binaries as well as characterise the outflows present. For example, extensive modelling of Eta Carinae’s X-ray light curve has enabled estimates of the companion’s mass-loss rate and wind velocity, despite not having been observed directly (Pittard & Corcoran 2002; Okazaki et al. 2008; Parkin et al. 2011; Russell et al. 2016). We note that X-rays can also be produced from shocks in the winds of single stars that develop from the intrinsic instability in line driving (MacGregor et al. 1979; Owocki et al. 1988; Oskinova 2016).

As the free electrons in the shocked gas experience accelerations from the ions present, and emit electromagnetic dipole radiation, the electrons slow down, the gas decreases in energy, and hence cools. This radiative cooling can alter the morphology of the shocks, as previously described by Equations 2.38 and 2.37, and so it is important to know if the shocks are in the radiative (isothermal) or adiabatic regime. Following Stevens et al. (1992), a ratio for the cooling time of the shocked gas verses the escape time from the intershock region can be defined:

$$\chi = \frac{t_{\text{cool}}}{t_{\text{esc}}} \approx \frac{v_8^4 d_{12}}{\dot{M}_{-7}}, \quad (2.39)$$

where v_8 is the wind velocity in units of 1000 km s^{-1} , d_{12} is the distance to the contact discontinuity in units of 10^7 km , and $\dot{M}_{-7}$ is the mass-loss rate in units of $10^{-7} \text{ M}_{\odot} \text{ yr}^{-1}$. This ratio enables us to determine if the shocks are adiabatic ($\chi \gtrsim 1$) or radiative ($\chi \ll 1$). However, it is worth noting that as the orbital separation changes in eccentric binaries the cooling may switch between the different regimes, with wider binaries having more adiabatic shocks.

As mentioned above, the shocks of colliding-wind binaries may also create non-thermal emission. At the shock interface, high energy cosmic rays are generated through a process known as Fermi acceleration or diffusive shock acceleration (Bell 1978a;b), although this requires particles to have energies that are already much greater than thermal, known as the injection problem. The Fermi-accelerated particles may then collide with ambient material, emitting non-thermal X-rays and γ -rays. Only recently have observations of this type been confirmed for colliding-wind binaries, such as for Eta Carinae (Hamaguchi et al. 2018).

2.5 Summary

In this chapter we have introduced the theory for describing the outflows of massive stars. Starting from a picture of a single star with a wind accelerated radially from a point source, by a combination of continuum and line driving, we showed analytical results for the wind velocity profile and mass-loss rate. We further generalised this formalism to include the effects of a finite stellar disc with a linear limb-darkening law. Using the derived wind structures, we described the numerical methods for solving the radiative transfer equation, to synthesise the emergent spectra of emission-line stars. In particular, we compared simulated spectra of P Cygni to that of a typical O-star.

Moving on to binary systems, we presented the morphology of their colliding winds. We presented equations for characterising the location and angle of the colliding-wind interface, as well as considerations for the adiabatic or radiative cooling regimes of the post-shock gas.

Chapter 3

Stellar orbits

In this chapter we introduce the mathematical framework for describing the orbits of binary stars. We then cover the observations and standard algorithms used for detecting and characterising these orbits. The derivations follow the detailed and thorough descriptions of Aitken (1964) and Murray & Dermott (2000). Whilst we construct our equations with a binary star system in mind, the physics applies to all scales from massive stars to exoplanets.

3.1 Types of orbits

Let us consider two stars, viewed as point particles, where the only forces acting are their mutual gravitational attraction. This is known as the two-body problem in classical mechanics. The solution to the two-body problem was first described by Newton (1687) where he stated the universal law of gravitation. Applying this law for two stars we can write

$$\mathbf{F}_1 = m_1 \ddot{\mathbf{r}}_1 = \frac{+Gm_1m_2}{r^2} \hat{\mathbf{r}}, \quad (3.1)$$

$$\mathbf{F}_2 = m_2 \ddot{\mathbf{r}}_2 = \frac{-Gm_1m_2}{r^2} \hat{\mathbf{r}}, \quad (3.2)$$

where G is the universal gravitational constant, m_1 and m_2 are the masses of the stars, $\ddot{\mathbf{r}}_1$ and $\ddot{\mathbf{r}}_2$ are the accelerations of the stars, $\mathbf{F}_1$ and $\mathbf{F}_2$ are the forces acting on the stars, $\hat{\mathbf{r}}$ is a unit vector pointing from the primary to the companion star, and the subscripts 1 and 2 denote the primary and companion star, respectively. These vectors are illustrated in Figure 3.1. Next, let us consider the acceleration of the




Figure 3.1: A schematic of the two-body problem where we have defined the geometric setup of the mutual gravitational forces, $\mathbf{F}$, and position vectors, $\mathbf{r}$, of the primary and companion stars.

companion star relative to the primary star, $\ddot{\mathbf{r}}$. By combining equations 3.1 and 3.2 we have

$$\ddot{\mathbf{r}} = \ddot{\mathbf{r}}_2 - \ddot{\mathbf{r}}_1 = \frac{-G(m_1 + m_2)}{r^2} \hat{\mathbf{r}}. \quad (3.3)$$

Now we recast this equation in polar co-ordinates, r and θ , shift the origin to the location of the primary star, and set the polar angle relative to an arbitrary reference direction in the plane of the stars. In this new co-ordinate system, if we only consider the $\hat{\mathbf{r}}$ terms – the motion along the line joining the two stars – we can write

$$\ddot{r} - r\dot{\theta}^2 = \frac{-G(m_1 + m_2)}{r^2}. \quad (3.4)$$

From this expression we want to eliminate the time dependence, so that we end up with a function, $r(\theta)$. To achieve this we temporarily make the substitution $u = 1/r$. We also make use of the fact that the total angular momentum, defined as the cross product of the position and momentum vectors, $\mathbf{J} = \mathbf{r} \times \mathbf{p}$, is a constant of the motion. More specifically, for motion in a plane (ie. the orbital plane in our case) we have $J = rp = rvm$; and it follows that the specific angular momentum can be written as $h = J/m = rv = r^2\dot{\theta}$ where the circumferential velocity component, $v = r\dot{\theta}$, was substituted. Altogether we can now write the following derivatives of our substitution

$$\frac{du}{d\theta} = \frac{d}{dt} \left(\frac{1}{r} \right) \frac{dt}{d\theta} = \frac{-\dot{r}}{r^2\dot{\theta}} = \frac{-\dot{r}}{h}, \quad (3.5)$$

$$\frac{d^2u}{d\theta^2} = \frac{-1}{h} \frac{\dot{r}}{dt} \frac{dt}{d\theta} = \frac{-\ddot{r}}{h\dot{\theta}} = \frac{-\ddot{r}}{h^2u^2}. \quad (3.6)$$




Figure 3.2: A schematic of the different conic sections which correspond to various geometries of the orbit followed by the companion: circle (yellow), ellipse (orange), parabola (pink), and hyperbola (purple). The black circle denotes the primary star.

Using the above results we can rewrite Equation 3.4 as

$$\frac{d^2u}{d\theta^2} + u = \frac{G(m_1 + m_2)}{h^2}, \quad (3.7)$$

which is a second order linear differential equation. This equation is known as Binet's equation and has the general solution:

$$u = \frac{G(m_1 + m_2)}{h^2} [1 + e \cos(\theta - \delta)], \quad (3.8)$$

where e and δ are the amplitude and angular constants of the integration, respectively. At this point we can convert back to r which results in an equation of the form

$$r = \frac{l}{1 + e \cos(\theta - \delta)}, \quad (3.9)$$

where $l = h^2/G(m_1 + m_2)$. Equation 3.9 describes a set of curves in polar co-ordinates, known as conic sections, where one focus is at the origin. A conic section is the curve created by the intersection of a plane with the surface of a cone. The constants e and l take on the geometric quantities known as the eccentricity and semi-latus rectum, respectively. In Figure 3.2, we illustrate the different conic sections which correspond to various eccentricities: circle ($e = 0$), ellipse ($0 < e < 1$), parabola ($e = 1$), and hyperbola ($e > 1$). Overall, we have shown that by considering only the gravitational force, as an inverse-square law, and constant total angular momentum, that orbiting




Figure 3.3: An orbital energy diagram comparing the effective potential U_{eff} , and its components, to discern the corresponding orbital geometry. Values presented are for the Earth ($m_2 = m_{\oplus}$) orbiting the Sun ($m_1 = M_{\odot}$). The colours correspond to the conic sections and orbital geometries illustrated in Figure 3.2.

stars will move along conic sections. For cases when $e < 1$, we have derived a result equivalent to what Kepler found empirically, and would later become known as his first law.

The type of conic section an orbiting star follows depends on the initial conditions of the system. Let us now consider the energy equation for a companion orbiting about its primary star

$$E = \frac{1}{2}m_2\dot{r}^2 + \frac{J^2}{2m_2r^2} + U(r), \quad (3.10)$$

where E/m_2 is the effective Hamiltonian of the system and $U(r)$ is its gravitational potential energy. We define an effective potential, $U_{\text{eff}}(r)$, such that

$$E = \frac{1}{2}m_2\dot{r}^2 + U_{\text{eff}}(r). \quad (3.11)$$

The first term on the right-hand side of Equation 3.11 is always positive and so it follows that $U_{\text{eff}}(r) < E_{\text{tot}}$. Through using this inequality we can gain insights into the allowed geometries for a given system. Let us consider the energy diagram in the right panel of Figure 3.3. Here, as an example we have set up the energy considerations for the Earth's orbit around the Sun. Depending on the value of E , the Earth is energetically restricted to radii bounded by the $U_{\text{eff}}(r)$ function, to regions below each horizontal line. As a result, 4 regimes emerge:

- $E = E_{\min}$. This is the minimum value of the total orbital energy, where $\dot{r} = 0$, and corresponds to bound circular motion. This radius is at 1 au as expected for the Earth's orbit around the Sun on an almost circular orbit.
- $E_{\min} < E < 0$. This corresponds to bound orbits with elliptical geometries. The orbiting body can move between a minimum and maximum radius, corresponding to the periastron and apastron of the orbit.
- $E = 0$. This is a special case of an unbound orbit with a parabolic geometry. The orbiting body will fly-past the primary star, reach a minimum separation, and then just escape the gravitational pull moving out to larger separations indefinitely. Note that the escape velocity is defined as $v_{\text{esc}} = (2Gm_1/r)^{1/2}$, which is the exact speed at which you get parabolic orbits.
- $E > 0$. This is the general case of unbound orbits following hyperbolic geometries. Similarly to the parabolic orbit, the orbiting body will fly-past the primary star, reaching a minimum orbit, and then escape. The higher the orbiting body's initial velocity the more open (flat) the hyperbolic orbit will become.

These energy considerations are applicable to many types of systems, from planets orbiting their host stars to massive binary stars as well as compact objects. In particular, for binary star systems that experience sudden mass loss or random kick velocities, such as those produced during SNe explosions, the orbital dynamics may be greatly affected. Mass loss from a SN explosion reduces the absolute value of the potential energy more than it reduces the orbital kinetic energy, leading to binaries becoming unbound (Hills 1983). This process is important for our understanding of so-called runaway stars with large peculiar velocities (Blaauw 1961) and stellar evolution tracks that may or may not lead to gravitational wave mergers (Belczynski et al. 2002).

3.2 Keplerian motion

Now that we have derived the geometry of one star's orbit around another, we can describe the observables through which we measure these orbits. For the remainder of

this chapter we limit ourselves to the bound orbits: the orbits about which this thesis is primarily concerned. For elliptical orbits $l = a(1 - e^2)$ where a is the semi-major axis. The semi-minor axis, b , is related via the eccentricity $e = \sqrt{1 - b^2/a^2}$.

At this point we need to define a co-ordinate system that is anchored in three dimensional space and has a reference point in time. First, we define the true anomaly, $\nu = \theta - \delta$, which is the angle between the periastron direction and the position of the star on the orbit. The true anomaly is calculated from the orbital phase, ϕ , via the mean anomaly, M , and eccentric anomaly, E , related by

$$2\pi\phi = M = E - e \sin E, \quad (3.12)$$

$$\tan\left(\frac{\nu}{2}\right) = \sqrt{\frac{1+e}{1-e}} \tan\left(\frac{E}{2}\right). \quad (3.13)$$

The mean anomaly tracks the fraction of the orbit that has elapsed in time since the last periastron. Starting from the mean anomaly we compute the eccentric anomaly using Equation 3.12. This is known as Kepler's equation and is a transcendental equation, so it has no closed form. Therefore, we solve Kepler's equation iteratively using the Newton–Raphson root-finding algorithm

$$f(E) = E - e \sin E - M, \quad (3.14)$$

$$E_{n+1} = E_n - \frac{f(E_n)}{f'(E_n)} = E_n - \frac{E - e \sin E - M}{1 - e \cos E_n}. \quad (3.15)$$

The algorithm iterates until $f(E) < \epsilon$ where ϵ is the desired accuracy. Usually setting $E_0 = M$ provides a good starting point.

Using the elliptical expression for l and the true anomaly we update Equation 3.9 for bound orbits to have the form

$$r = \frac{a(1 - e^2)}{1 + e \cos(\nu)}. \quad (3.16)$$

An alternative form of this equation exists by substituting a for the period P and masses of the stars via Newton's generalisation of Kepler's third law, $a^3/P^2 = G(m_1 + m_2)/(4\pi^2)$. We can now use Equation 3.16 for measuring orbiting stars via two different observing modes: velocities along the line of sight and positions in the plane of the sky.

To fully specify the stars' orbits we further define several geometric quantities that are illustrated in Figure 3.4. There are two coordinate systems defined: (x, y, z)




Figure 3.4: A schematic of the orbital elements and co-ordinate systems for transforming between the orbital plane, (x, y, z) (shown in pink), and the plane of the sky, (X, Y, Z) (shown in black). Also shown is the direction of the observer and north, the direction of the companion's orbit, the position of the ascending node, the position of periastron, and the angles for the longitude of the ascending node, Ω , inclination, i , the argument of periastron, ω , and the true anomaly, ν .

with respect to the orbital plane with x in the direction of periastron, and (X, Y, Z) with respect to the plane for the sky with X in the direction of North. The longitude of the ascending node, Ω , is defined as the angle in the plane of the sky from the North direction to where the orbital plane intersects with the plane of the sky, on the node where the companion is moving towards the observer, the ascending node. The inclination, i , is the angle between the orbital plane and the plane of the sky. The convention is that $0 < i < 90$ degrees for counterclockwise orbits on the plane of the sky from the observer's perspective, and $90 < i < 180$ degrees for clockwise orbits. The argument of periastron, ω , is the angle in the orbital plane, along the direction of the orbit, from the ascending node to the direction of periastron. This angle is sometimes referred to as the longitude of periastron in the literature.

3.2.1 Radial velocities

The radial velocities we measure are the Keplerian velocities projected onto our line of sight. Radial velocities are calculated from the Doppler shifts extracted from spectral lines using time-series spectroscopy. To calculate the line of sight orbital velocities of a primary star about its companion we have to move into the centre-of-mass frame.

This involves the transformation

$$a_1 = \frac{m_2}{m_1 + m_2} a, \quad (3.17)$$

$$r_1 = \frac{a_1(1 - e^2)}{1 + e \cos(\nu)} = \frac{m_2}{m_1 + m_2} \frac{a(1 - e^2)}{1 + e \cos(\nu)}, \quad (3.18)$$

and its derivative

$$\dot{r}_1 = \frac{a_1(1 - e^2)\dot{\nu}e \sin \nu}{(1 + e \cos(\nu))^2} = \frac{r_1\dot{\nu}e \sin \nu}{a_1(1 - e^2)} = \frac{m_2}{m_1 + m_2} \frac{r\dot{\nu}e \sin \nu}{a(1 - e^2)}, \quad (3.19)$$

where a_1 is the semi-major axis of the primary's orbit in the centre-of-mass frame. Next let us write the primary's position and velocity using the coordinates in the orbital plane:

$$\mathbf{r}_1 = \begin{pmatrix} r_1 \cos \nu \\ r_1 \sin \nu \end{pmatrix}, \quad (3.20)$$

$$\dot{\mathbf{r}}_1 = \begin{pmatrix} \dot{r}_1 \cos \nu - r_1 \dot{\nu} \sin \nu \\ \dot{r}_1 \sin \nu + r_1 \dot{\nu} \cos \nu \end{pmatrix}. \quad (3.21)$$

Now we want to rearrange Equation 3.21 such that we have an expression for the primary's velocity as a function of the true anomaly only. In order to achieve this, we substitute in expressions for r_2 and $\dot{r}_2$ from Equations 3.18 and 3.19 respectively, and following some algebra we have

$$\dot{\mathbf{r}}_1 = \frac{m_2}{m_1 + m_2} \frac{r^2 \dot{\nu}}{a(1 - e^2)} \begin{pmatrix} -\sin \nu \\ e + \cos \nu \end{pmatrix} = \sqrt{\frac{Gm_2^2}{m_1 + m_2} \frac{1}{a(1 - e^2)}} \begin{pmatrix} -\sin \nu \\ e + \cos \nu \end{pmatrix}, \quad (3.22)$$

where we have eliminated $\dot{\nu}$ by using Kepler's second (equal areas, A , swept out in equal times, $\dot{A} = \text{const}$) and third laws. Lastly, we project this vector onto the line of sight; the Z direction in our plane of the sky co-ordinate system. Using the angles defined in Figure 3.4, this projection takes the form

$$\hat{\mathbf{Z}} = \begin{pmatrix} \sin \omega_1 \sin i \\ \cos \omega_1 \sin i \\ \cos i \end{pmatrix}, \quad (3.23)$$

and results in

$$v_{\text{kep},1} = \dot{\mathbf{r}}_1 \cdot \hat{\mathbf{Z}} = k_1(\cos(\omega_1 + \nu) + e \cos \omega_1) + \gamma, \quad (3.24)$$

where $v_{\text{kep},1}$ is the Keplerian radial velocity of the primary star projected onto our line of sight and we have included γ , the radial velocity offset, which accounts for the




Figure 3.5: Examples of Keplerian radial velocity curves for changes to the orbital parameters: eccentricity, e , argument of periastron, ω , semi-amplitude, k_1 , and radial velocity offset, γ . Only one parameter is varied per panel, see text for additional details.

proper motion of the centre of mass of the binary along our line of sight. ω_1 is the argument of periastron of the primary star, and $\omega_1 = \omega_2 + \pi$, but often just $\omega = \omega_1$ is quoted. k_1 is the radial velocity semi-amplitude of the primary and can be written in various helpful forms:

$$k_1 = \sqrt{\frac{G}{1-e^2}} m_2 \sin i (m_1 + m_2)^{-1/2} a^{-1/2}, \quad (3.25)$$

$$k_1 = \frac{29.78 \text{ km s}^{-1}}{\sqrt{1-e^2}} \frac{m_2 \sin i}{M_\odot} \left(\frac{m_1 + m_2}{M_\odot} \right)^{-1/2} \left(\frac{a}{1 \text{ au}} \right)^{-1/2}, \quad (3.26)$$

$$k_1 = \frac{29.78 \text{ km s}^{-1}}{\sqrt{1-e^2}} \frac{m_2 \sin i}{M_\odot} \left(\frac{m_1 + m_2}{M_\odot} \right)^{-2/3} \left(\frac{P}{1 \text{ yr}} \right)^{-1/3}, \quad (3.27)$$

where the masses can be specified in solar masses, time in years, and semi-amplitude in astronomical units.

In Figure 3.5 we show some example radial velocity curves and their dependence on the key parameters of Equation 3.24. In the top-left panel we show the dependence on e over the interval $0 < e < 0.9$, for a fixed $k_1 = 150 \text{ km s}^{-1}$, $\omega = 270^\circ$, and $\gamma = 0 \text{ km s}^{-1}$. As the eccentricity increases the radial velocities show faster changes around periastron. In the top-right panel we show the dependence on ω over the interval $180^\circ < \omega < 360^\circ$, for a fixed $k_1 = 100 \text{ km s}^{-1}$, $e = 0.45$, and $\gamma = 0 \text{ km s}^{-1}$. For $\omega = 270^\circ$, we view the system directly from the apastron side. As ω increases (decreases) from this value, we see corresponding increases in positive (negative) velocities around periastron. This is a result of the faster orbital velocities at periastron having a greater magnitude when projected on to these lines of sight. In the bottom-left panel we show the dependence on k_1 over the interval $50 \text{ km s}^{-1} < k_1 < 150 \text{ km s}^{-1}$, for a fixed $\omega = 270^\circ$, $e = 0.45$, and $\gamma = 0 \text{ km s}^{-1}$. Here we see that as k_1 increases the amplitude of the velocity variations increases. In the bottom-right panel we show the dependence on γ over the interval $-50 \text{ km s}^{-1} < \gamma < 50 \text{ km s}^{-1}$, for a fixed $\omega = 270^\circ$, $k_1 = 100 \text{ km s}^{-1}$, and $e = 0.45$. Here we see that γ simply introduces a constant offset in the radial velocity curves.

Two additional parameters are also required to compare these radial velocity curves to data. The time of periastron, T_0 , and the period, P , which are used to calculate ϕ in Equation 3.12. This changes the time taken for each orbit and the reference time at which the companion is at its smallest separation to the primary star, ie. periastron.

3.2.2 Astrometric positions

To fully characterise the three-dimensional orbit of a binary star system we require astrometric measurements in combination with radial velocities. Astrometry has two forms: relative and absolute. Relative astrometry is particularly useful for binary star systems as high-spatial resolution, adaptive optics imaging systems can resolve the primary and companion stars of some near-by binaries. Usually Gaussians or point-spread function fitting is performed on each source, along with a known orientation

and pixel scale of the detector, to produce relative measurements of the stellar locations at the epoch of each exposure. For example, the state-of-the-art near-infrared interferometer, Gravity, on the VLT, has a maximum spatial resolution of 2 milliarcseconds. This corresponds to resolving stellar separations of 1 au out to distances of 500 pc. However, its narrow field of view reduces the ability to observe reference stars, and so the measurements are relative.

To calculate the relative positions of the stars we start from Equation 3.16. At each epoch we can write in Cartesian co-ordinates, in the orbital plane, the three-dimensional position of the orbiting star to be

$$x = r \cos \nu, \quad (3.28)$$

$$y = r \sin \nu, \quad (3.29)$$

$$z = 0. \quad (3.30)$$

We then define the transformation required to convert the coordinates from the orbital plane to the plane of the sky. We use two rotation matrices, labelled as $\mathbf{R}_x(\theta)$ for rotation about the x -axis, and $\mathbf{R}_z(\theta)$ for rotation about the z -axis, with elements

$$\mathbf{R}_x(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}, \quad (3.31)$$

and

$$\mathbf{R}_z(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (3.32)$$

We apply three successive rotations

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \mathbf{R}_z(\Omega) \mathbf{R}_x(i) \mathbf{R}_z(\omega) \begin{pmatrix} x \\ y \\ z \end{pmatrix} \quad (3.33)$$

which can be followed graphically in Figure 3.4. First, a rotation about the z -axis in the orbital plane by the argument of periastron. Second, a rotation about the x -axis by the inclination. Third, a rotation about the z -axis in the sky plane by the longitude of the ascending node. As a result the three-dimensional position of the




Figure 3.6: Examples of astrometric orbits for changes to the orbital parameters: eccentricity, e , argument of periastron, ω , longitude of the ascending node, Ω , inclination, i , semi-major axis, a , and parallax π . Only one parameter is varied per panel, see text for additional details.




Figure 3.7: Example of stellar separation, ρ , and position angle, P.A., that would be measured from relative astrometry. The example orbital parameters used were $e = 0.45$, $\omega = 270^\circ$, $\Omega = 0^\circ$, $i = 45^\circ$, $a = 10$ au, and $\pi = 0.43$ mas.

orbiting star in the plane of the sky is

$$X = r(\cos \Omega \cos(\omega + \nu) - \sin \Omega \sin(\omega + \nu) \cos i), \quad (3.34)$$

$$Y = r(\sin \Omega \cos(\omega + \nu) + \cos \Omega \sin(\omega + \nu) \cos i), \quad (3.35)$$

$$Z = r \sin(\omega + \nu) \sin i. \quad (3.36)$$

In Figure 3.6 we show some example astrometric orbits and their dependence on the key parameters of Equation 3.16, 3.34, and 3.35. The overall scale of the visual orbits are set by the parallax, π : the apparent shifts in position of a system relative to background or distant objects. The distance of a system is related to the parallax by $d = 1/\pi$, where $[d]$ = parsecs and $[\pi]$ = arcseconds. In each panel we vary one parameter and fix all the others. The fixed set is $e = 0.45$, $\omega = 270^\circ$, $\Omega = 270^\circ$, $i = 45^\circ$, $a = 10$ au, and we set a system parallax of $\pi = 0.43$ mas. Across the Figure's panels we see how e changes the ellipticity of the projected orbit, ω changes the direction of periastron in the orbital plane, Ω changes the direction of periastron in the plane of the sky, i changes the how the orbit is angled relative to the plane of the sky, and a has a similar effect to π , both changing the observed size of the projected orbit.

In practice, relative astrometry data are usually in the form of stellar separation, ρ , and position angle, P.A., defined as the angular offset of the companion relative

Orbital element	Star	Range	Units
Elements of the system			
Period	both	$0 < P$	days
Time of periastron	both	$0 < T_0$	JD
Eccentricity	both	$0 < e < 1$	
Argument of periastron	primary	$0 < \omega < 2\pi$	radians
Elements of the radial velocity orbit			
Semi-amplitude	primary	$0 < k_1$	km s^{-1}
Semi-amplitude	secondary	$0 < k_2$	km s^{-1}
Radial velocity offset	primary	$-\infty < \gamma_1 < \infty$	km s^{-1}
Radial velocity offset	secondary	$-\infty < \gamma_2 < \infty$	km s^{-1}
Elements of the astrometric orbit			
Semi-major axis	both	$0 < a$	arcseconds
Inclination	both	$0 < i < \pi$	radians
Longitude of the ascending node	both	$0 < \Omega < 2\pi$	radians

Table 3.1: A breakdown of the orbital elements to fully specify the three-dimensional orbit of a binary star system.

to the North celestial pole. In Figure 3.7, we show an example orbit as measured in ρ and P.A. The P.A spans $-\pi < \text{P.A.} < \pi$ and so there is a discontinuity where the curve crosses the edge of this range.

Now that we have derived models for both radial velocity and relative astrometry data, we have a mathematical framework in place that enables us to model data and infer the orbital parameters of binary star systems. A summary of the model parameters is shown in Table 3.1. The parameters required to specify an orbital model depend on the availability of data. In the simplest case, for single-lined spectroscopic binaries (SB1) we have one radial velocity curve, and therefore we can measure the parameter set $\{P, T_0, e, \omega, k_1, \gamma_1\}$. If our spectroscopic data shows two distinct stellar components, known as a doubled-lined spectroscopic binary (SB2), then we can measure the parameter set $\{P, T_0, e, \omega, k_1, k_2, \gamma_1, \gamma_2\}$. In the case of relative astrometry

Data available	Dynamical mass quantity calculable
SB1	minimum mass $m_2 \sin i$, if m_1 is estimated
SB2	mass ratio q and minimum masses $m_1 \sin^3 i$ and $m_2 \sin^3 i$
Relative astrometry	total mass $m_1 + m_2$
SB1 + Relative astrometry	individual masses m_1 and m_2
SB2 + Relative astrometry	individual masses m_1 and m_2
Absolute astrometry	individual masses m_1 and m_2

Table 3.2: Calculable dynamical mass quantities given the availability of radial velocity and astrometric data.

data, we can measure the parameter set $\{P, T_0, e, \omega, a, i, \Omega\}$. When we have both radial velocity and astrometry data we can measure all the parameters in the set $\{P, T_0, e, \omega, k_1, k_2, \gamma_1, \gamma_2, a, i, \Omega\}$.

From the aforementioned parameter sets, we may be able to derive other parameters of interest: most notably the distance of the system and masses of the stars. In Table 3.2 we summarise the link between the data available and the dynamical mass quantity that is calculable. The SB1 case does not encode enough information about the pair of stars to measure the dynamical masses directly. However, SB1 systems are not without merit. The mass function

$$f = \frac{Pk_1^3}{2\pi G} = \frac{m_2^3 \sin^3 i}{(m_1 + m_2)^2}, \quad (3.37)$$

enables calculation of the minimum mass, $m_2 \sin i = f^{1/3} m_1^{2/3}$, of the unseen companion, in the limiting case where $m_1 \gg m_2$ and there exists an estimate of m_1 by other means. The SB2 case encodes sufficient information to calculate a mass ratio, $q = m_2/m_1$, and a minimum mass, $m_1 \sin^3 i$ and $m_2 \sin^3 i$, for each star which is only dependent on dynamical information. For eclipsing binaries the inclination may be apparent from the system's light curve allowing the computation of actual individual masses. From relative astrometry the summed mass of both stars can be measured. In cases of SB1 or SB2 data combined with astrometry data, the individual masses can be computed. Dynamical masses are essentially model independent, and relative to empirical estimates or spectroscopic and evolutionary models they can produce very accurate masses.

The state of the art in this area is showcased by a couple of recent publications. Richardson et al. (2021) measure the dynamical masses of a nitrogen-rich Wolf-Rayet star, WR 133, using a combined visual and spectroscopic orbit for the first time. They find the dynamical masses to be lower than the masses inferred from the spectral modelling. This highlights the importance of dynamical modelling and how it can perform a crucial role in testing spectral and evolutionary models of interacting binaries. Thomas et al. (2021) collate and expand the astrometric and spectroscopic data for the colliding-wind binary, WR 140. Through these measurements they are able to improve the estimates of the orbital elements and stellar masses. This in turn enabled them to analyse the evolutionary history of the system, finding that the Wolf-Rayet star would have lost the majority of its envelope through stellar winds and only a limited amount of binary interaction. The utility of these dynamical models also extends down to lower mass regimes. For example, near the substellar boundary, dedicated observing campaigns (eg. Dupuy & Liu 2017) have proved effective at probing various evolutionary models and parameter distributions of ultracool dwarfs.

3.3 Algorithms for velocity extraction

Now that we have outlined parameterised models for characterising the orbits of binary systems, we turn our attention to the data from which the models will learn the orbital solutions. In most cases, the orbital parameters are determined using radial velocity data, enabling the motion of the star to be constrained along the line of sight. Radial velocity data are obtained by algorithmically calculating the Doppler shifts of spectral lines from time-series spectroscopy. In order to put good constraints on the orbital parameters, good coverage of the system’s orbit is required. In particular, for highly-eccentric stars it is important to schedule the majority of observations at the time of periastron passage, in order to observe the full extent of the radial velocity excursions.

To extract velocity data from spectra we must first connect how a given line profile relates to the velocity of interest. This complexity arises since line profiles are never a Dirac delta function, encoding a singular velocity value, but rather, a broadened, multifaceted profile that needs disentangling. In the simplest case we have absorption lines that are often simple symmetrical functions. In more complex




Figure 3.8: Synthetic time-series spectra for benchmarking radial velocity extraction algorithms. Each panel displays a subset of the 200 spectra per dataset for the He I 6678 Å line. The blue dashed lines represent the injected primary orbital motion. The red and grey dashed lines represent the injected companion orbital motion for the absorption and emission lines in datasets B and C, respectively. The continuum has been removed for all spectra.

cases we have broad emission-lines, formed over large volumes, within complex debris fields of outflowing material, and afflicted by other interloping sources of absorption and emission. In our case we are interested in the orbital motion, and as such we must carefully consider how an algorithm extracts the dynamical information from the spectra.

In order to elucidate the workings of a few different algorithms for radial velocity extraction, we create synthetic datasets to test each algorithm’s ability to recover the injected orbital motion. In particular, we target emission-line stars with P Cygni profiles. We simulate the primary star’s He I line profiles by making use of the same procedure previously described in Section 2.3.2 for the prototypical emission-line star, P Cygni. We select He I 6678 Å as our test line and we re-bin the spectra to a resolution of 0.2 Å. Next, we inject radial velocities by translating the entire spectrum at each of 200 epochs across two orbital periods. The ground truth orbital parameters are $e = 0.2$, $k_1 = 100 \text{ km s}^{-1}$, $\omega = 270^\circ$, and $\gamma_1 = 0 \text{ km s}^{-1}$.

The synthetic spectral line profiles are then duplicated across three datasets which are displayed in Figure 3.8. In dataset A we do not add any further signals, making this the easiest dataset. This is demonstrative of an SB1 system. In dataset B we add a companion’s absorption line, resulting in an SB2 system similar to observations of emission-line stars orbiting O-stars. This line has a peak absorption one-quarter the amplitude of the P Cygni profile and is generated from a narrow, negative Gaussian function. In dataset C we add a companion which exhibits emission. The emission line is generated from a broad, positive Gaussian function with a peak emission one-quarter the amplitude of the P Cygni profile. Lastly, Poisson noise is added to the data, at the level equivalent to having a signal-to-noise ratio of 100 at the peak of the He I line profiles.

3.3.1 Bisector half-maximum

A simple and fast algorithm for extracting radial velocities is to use the bisector at half the maximum flux (BHM) (eg. Martínez Fiorenzano et al. 2005). In Figure 3.9 we display the BHM algorithm applied to some test data from dataset A. The left panel shows the bisector, at a range of different heights, applied to the first emission-line profile in dataset A. In the right panel we show the bisectors for all spectra in dataset




Figure 3.9: Left panel: the bisector algorithm for radial velocity extraction for the first He I line profile in dataset A. The blue line represents the bisector at different fractions of the emission-line height. Right panel: the bisector for all spectra from dataset A, overlaid in the rest frame, for comparison relative to the true velocity shown by the dashed line.)

A, after having been shifted back to the de-orbited rest frame. For P Cygni profiles we find that the velocity of the bisector is dependent on the height at which it is calculated. This is due to the asymmetry in the emission portion of the line, owing to absorption affecting its blue shoulder. Consequently, we can expect systematic offsets to higher velocities, which may be regulated by the absorption depth or interloping signals at different epochs.

3.3.2 Cross-correlation

A popular algorithm for extracting radial velocities is the cross-correlation method (Tonry & Davis 1979). This algorithm involves finding the maximum overlap of a spectral line (or entire section of a spectrum) with a template. More formally the cross-correlation function is written

$$\mathbf{C}_\lambda = \sum_i \mathbf{S}_i \mathbf{T}_{i+\lambda}, \quad (3.38)$$

where $\mathbf{C}_\lambda$ is the cross-correlation function at a specific wavelength shift λ , $\mathbf{S}$ is the observed spectrum, and $\mathbf{T}$ is the template. The velocity shift can either be taken as




Figure 3.10: Left panel: a visual representation of the cross-correlation algorithm for radial velocity extraction for the first He I line profile in dataset A. The blue profiles represent the master template at different wavelength shifts. Right panel: the parabolic fit to the cross-correlation function calculated from the left panel. The dashed line is the derived velocity relative to 0 km s^{-1} .

the maximum of the cross-correlation function or alternatively from a functional fit to the cross-correlation function to compute the maximum at a sub-pixel level.

There are a few options for how to formulate the template. Gaussian functions are often used, especially with simple absorption line data, and work well when many lines are cross-correlated at once. Another method is to generate a spectral model of the line profile using a radiative transfer code, as described in Section 2.3.2, and use the simulated spectrum as the template. Alternatively, as utilised by Shenar et al. (2019) the observations can be used to form the template in an iterative scheme. We implement this method here.

We initially set the first spectral line as the template. We cross-correlate this template with each spectrum finding the cross-correlation functions. We fit parabolas to the cross-correlation functions to derive the velocity shifts. Next, we shift all spectra by these velocities and median stack them to create a master line profile. We then use this master profile as the new template and repeat the cross-correlation to derive the final velocities.

In Figure 3.10 we display an example of this cross-correlation algorithm applied to the first spectral line in dataset A. In the left panel we show a visual representation




Figure 3.11: The multi-Gaussian algorithm for radial velocity extraction. Both panels include the first He I line profile from dataset A, however in right panel the signal-to-noise ratio is lowered from 100 to 20. The individual Gaussian components are shown in grey, and the residuals of the total fit are shown in the lower panels.

of how the master template is shifted over the observed spectrum. In the right panel we plot the cross-correlation function at each velocity shift, and the parabolic fit for deriving the velocity shift for this example spectrum.

3.3.3 Multi-Gaussian

Multi-Gaussian is an algorithm which finds its use when the line profiles are particularly complex, enabling the decomposition of individual line profiles into their constituent components. Gaussians are used as the base function because they are mathematically uncomplicated; and moreover, they often well represent the appearance of simpler line profiles. As an example, this line-fitting algorithm has been previously applied to the complex Galactic system SS433 whose multiple modes of mass-loss are elucidated by consideration of its constituent Gaussian components (Blundell et al. 2008; 2011).

Here, we implement the multi-Gaussian fitting algorithm as follows. The contin-




Figure 3.12: Determining the optimal number of Gaussian components in the multi-Gaussian algorithm with the Bayesian Information Criterion (BIC).

uum is removed using a Savitzky–Golay filter (Savitzky & Golay 1964) on regions of the spectrum not dominated by line transitions. We then construct a template, for a specific line profile, consisting of multiple Gaussian functions, each with an associated initial set of parameter guesses and bounds. The resulting model is a non-linear function constructed from the sum of the individual Gaussians. The model parameters are optimized using a Trust Region Reflective algorithm (Branch et al. 1999; Jones et al. 2001) which efficiently explores the multidimensional parameter space in order to minimize the sum of squares.

The templates are built using a combination of human input and statistical testing. Starting from one component, we manually add new components based on visual inspection of the residuals. We repeat this process until we find a template that consistently produces good fits for spectra over a few observing nights. We allow ourselves the freedom to vary the number of components through time and use negative Gaussians where appropriate to model absorption. To evaluate the goodness of fit, we follow Riener et al. (2019) by assessing whether the normalized residuals show a normal distribution, since we expect the errors in the flux bins to be Gaussian and homoscedastic. We perform a two-sided Kolmogorov–Smirnov test (Kolmogorov




Figure 3.13: Benchmarking velocity extraction algorithms: bisector at half the maximum flux (BHM), cross-correlation, and multi-Gaussian. The top panel shows the injected orbital motion, and the panels below show the residuals, for datasets A, B, and C, after this motion has been subtracted from the extracted velocities. The panels on the right-hand side display the histograms for each of the residuals.

1933; Smirnov 1939), where the null hypothesis states that the normalized residuals resemble a normal distribution. If the p -value for a multi-Gaussian fit is greater than 0.0027 (3σ), then we fail to reject the null hypothesis and the fit is considered good. An additional test, often used to decide between competing models, is to evaluate the Bayesian Information Criterion (BIC) (Schwarz 1978). The BIC score balances the likelihood of the data given a model, with the model complexity, to select the most appropriate model.

In Figure 3.11 we display an example of the multi-Gaussian algorithm applied to the first spectral line in dataset A. In the left panel we show the best fit to the spectral line at its original signal-to-noise ratio of 100. In Figure 3.12 we show how

the BIC score for this spectral line varies as the number of Gaussian components are increased (blue line). Here we see the best model, or minimum score, is achieved by a combination of 5 Gaussian components. Using fewer components results in a worse fit to the data, and more components results in a large model complexity for diminishing returns in the quality of the fit. For comparison, in the right panel of Figure 3.11, and with corresponding BIC scores in Figure 3.12 (green line), we display the multi-Gaussian algorithm applied to the same spectral data after it has had its signal-to-noise ratio reduced to 20. In this case there is less information available in the data; and as such, the number of components that can be justified is fewer. The result is an optimal multi-Gaussian fit comprised of 3 Gaussian components.

Once the spectral lines are fit we can scrutinize each Gaussian component for astrophysical meaning. Despite using Gaussians as our base function we do not require dynamical components to be strictly Gaussian. As shown by Blundell et al. (2007), it is possible to take a weighted mean of multiple components (centroid wavelength of each Gaussian weighted by its area) to extract clear signals from a superposition of profile shapes. For our synthetic datasets we take the weighted mean of the narrow emission components as representations of the bulk orbital motion.

We apply all three velocity extraction algorithms – BHM, cross-correlation, and multi-Gaussian – to the entirety of datasets A, B, and C. The results are shown in Figure 3.13 and we quantify the dispersion in the residuals, between the injected orbits and the extracted data, in terms of their standard deviations, σ . For dataset A, the results show all three algorithms performing well, with approximately Gaussian residuals, but different offsets. The cross-correlation algorithm results in the lowest standard deviation ($\sigma = 0.56$), and the BHM ($\sigma = 1.72$) and multi-Gaussian algorithm ($\sigma = 1.92$) performing about equally. For dataset B, the presence of the absorption line causes a loss in performance for the BHM ($\sigma = 2.68$) and cross-correlation ($\sigma = 1.76$) algorithms. In particular, these two algorithms show periodic systematics caused by the absorption line crossing through the emission-line profile. For the multi-Gaussian algorithm ($\sigma = 1.93$) the standard deviation of the residuals remains the same between datasets A and B, as the algorithm is able to disentangle the individual stellar components. For dataset C, all three algorithms show a loss in

performance as the companion’s broad emission line makes extracting the primary star’s motion difficult. In this case, using the multi-Gaussian algorithm ($\sigma = 2.30$) results in the lowest standard deviation in the residuals, relative to the BHM ($\sigma = 3.83$) and cross-correlation ($\sigma = 2.55$) algorithms.

3.4 Summary

In this chapter we have introduced the mathematical framework for describing the orbits of binary stars. Starting from a picture of two bodies moving only under the influence of their gravitational attraction, we showed, via Binet’s differential equation, that the stars’ trajectories follow conic sections. For bound orbits the stars trace out ellipses. We parameterised and projected this motion on to both the line of sight and plane of the sky, for modelling the observables extracted from spectroscopy and interferometry, respectively.

We described three commonly used algorithms for velocity extraction from Doppler shifting spectral lines: bisector at half the maximum flux, cross-correlation, and multi-Gaussian. We created three synthetic datasets, based on the He I line emission-line profile in P Cygni, with different interloping companion lines to highlight the advantages and limitations of each algorithm. We found the cross-correlation algorithm performed the best when the data was uninterrupted by companion signals. However, the multi-Gaussian algorithm performed the best when disentangling multiple dynamical components. Although we have used these algorithms for finding the orbital motion of binary stars, the techniques can, in general, be used to discover the dynamics of many different phenomena.

Chapter 4

Uncovering the orbital dynamics of stars hidden inside their powerful winds

4.1 Introduction

Binary systems that include stars exhibiting powerful winds are pertinent case studies for our overall picture of massive stars. Observationally, these systems show spectral lines predominantly in emission owing to their extended atmospheres and dense winds where the lines are formed (Beals 1929). This, in turn, presents certain diagnostic challenges, especially in the most extreme mass-loss cases where the outflows can obscure measurements of the central star. The optically-thick gas enshrouding the star moves the effective photospheric radius outwards in the flow; and therefore, deducing the obscured system's parameters can be problematic.

The orbital parameters of stars are typically determined by fitting Keplerian models to radial velocity measurements, and the radial velocities are calculated from the Doppler shifts extracted from spectral lines, as presented in Chapter 3. For stars without strong outflows, having absorption line spectra, the orbital motion can be tracked at the stellar surface. Precision on the order of $\sim 10 \text{ ms}^{-1}$ can be achieved with state-of-the-art instrumentation and extraction algorithms (Konacki 2005; Konacki et al. 2010), nearing the magnitude at which intrinsic stellar variability becomes important (Saar et al. 1998).

For stars with powerful winds, the emission-line spectra are manifestly more complex and systematic effects are the dominant source of error. The broad emission lines

are formed far out in the stellar wind (Beals 1929), encoding dynamical information from an extensive volume of outflowing gas. The winds show stratification in the line emission regions dependent on each line’s excitation energy (Bowen 1928; Kuhl 1973; Willis 1982) and opacity (Hillier 1987). In Chapter 3 we considered the radial velocities of spectral lines to be simply a translation of the line profile, for stars with powerful winds this may not be the case.

To overcome these challenges, earlier work has focused on extracting radial velocities from high excitation lines, known to form deep down in the wind, such as He I/II or N IV/V, depending on the spectral type and availability of lines (eg. Shenar et al. 2019). But, even for high excitation lines it is unclear whether the entire line emission region is co-moving along the star’s orbital path. In fact, large discrepancies have been found between the radial velocities extracted from various high-energy lines in WR binaries (Stickland et al. 1984; Shylaja 1986).

It has been suggested that these discrepancies are the result of non-Keplerian velocities present in the extended wind (Shylaja 1987). Radial velocities may also be affected by the wind and radiation of a companion star: the formation of colliding winds can complicate the line profiles (eg. Marchenko et al. 2003; Groh et al. 2012a;b) and the radiation field can influence the ionization structure of the primary star’s wind (Kruip 2011; Madura et al. 2012; Clementel et al. 2015a;b). Consequently, emission lines may contain velocity information which is not purely Keplerian, even if the motion of the stars from which they arise is Keplerian. The degree to which extracted radial velocities deviate from the true Keplerian orbits will depend on the exact nature of the line transition and outflows in which they are emitted.

4.1.1 Eta Carinae: the extreme case

When it comes to testing models of lines formed in extended regions the LBV Eta Carinae is an extreme case study. Characteristic of LBVs, Eta Carinae has an enormous mass-loss rate and relatively slow wind, resulting in the longest wind-flow times to the emission-line regions of almost any emission-line star known today. As a consequence, the observational effects will be at their most pronounced in this system.

Eta Carinae is located in the Trumpler 16 cluster in the Carina nebula (Walborn 1973). Situated at a distance of 2.3 ± 0.1 kpc (Allen & Hillier 1993; Smith 2006) and

generating a luminosity of $L = 5 \times 10^6 L_{\odot}$ (Davidson & Humphreys 1997), it is one of the most luminous sources in our Galaxy. Eta Carinae first captured the attention of astronomers due to its “Giant Eruption” in the 1840s, where it is estimated to have ejected in excess of $45 M_{\odot}$ of material over about a decade (Morris et al. 2017), forming the bipolar Homunculus nebula seen today (Currie et al. 1996). Further quasi-periodic eruptions have been detected, such as the so-called Lesser Eruption in 1890 (Walborn & Liller 1977; Humphreys et al. 1999) and a photometric event in 1941 (de Vaucouleurs & Eggen 1952). These outbursts are responsible for a complex environment of debris surrounding the source, including peculiar objects such as the “Weigelt blobs” or “speckle objects” (Weigelt & Ebersberger 1986; Dorland et al. 2004).

Eta Carinae is a binary system with a stable orbital period of $P = 2022.7 \pm 1.3$ d (Damineli et al. 2008a). The binary scenario was originally proposed by Damineli et al. (1997; 1992 periastron data) and since that time the orbital parameters have been computed using numerous techniques and observational data. Overall, the general consensus is for a highly eccentric (~ 0.9) binary which is supported by models of X-ray light curves (Pittard et al. 1998; Corcoran et al. 2001; Okazaki et al. 2008) and He I/II line equivalent widths (Richardson et al. 2015; Teodoro et al. 2016). However, the exact orbital solution has been hard to establish by virtue of the differences in results obtained from various species and the influence of the companion (Groh et al. 2012a).

The primary star is an LBV with an extremely powerful wind: a mass-loss rate of $8.5 \times 10^{-4} M_{\odot} \text{yr}^{-1}$ and terminal wind velocity of 420 km s^{-1} (Hillier et al. 2001; Groh et al. 2012a; Clementel et al. 2015a). Under these conditions, Hillier et al. (2001) showed that H-alpha could be formed in an extended region between 6 – 60 AU (assuming $R_* = 60 R_{\odot}$) outwards from the primary, and these results have been observationally supported by Wu et al. (2017). As a result, we know the primary star displays spectral lines formed in highly extended regions from the central star.

Details of the companion are far more uncertain because it has never been directly observed; it is a few orders of magnitude less bright than the primary in visible and longer wavelengths, and the companion is not seen in the UV either (Hillier et al. 2006). Estimates for the companion’s wind have been made from modelling X-ray

observations, yielding estimates of a mass-loss rate of $\sim 10^{-5} \text{ M}_{\odot} \text{ yr}^{-1}$ and terminal wind velocity of $\sim 2000 - 3000 \text{ km s}^{-1}$ (Pittard & Corcoran 2002; Parkin et al. 2011).

4.1.2 RMC 140b and GG Carinae: the typical cases

We also investigate RMC 140b (alternative names: BAT99 103 or VFTS 509) and GG Carinae (alternative names: HD 94878 or CPD-59 2855), which are more typical emission-line star systems than Eta Carinae. The aim here is to gauge how large the observable effects may be for lines formed in extended regions, for systems more representative of a larger population of stars, and the ramifications for their calculated orbital parameters.

RMC 140b is located in the Large Magellanic Cloud and at a distance of $49.97 \pm 1.11 \text{ kpc}$ (Pietrzyński et al. 2013). Moffat et al. (1987) first discovered that RMC 140b is a binary system, finding a period of 2.76 d. The system is comprised of an O-star (type O4 V, $L = 5.62 \times 10^5 L_{\odot}$) and a WR star (type WN5, $L = 1.23 \times 10^6 L_{\odot}$) (Evans & Taylor 2011; Shenar et al. 2019). For our analysis, we only concern ourselves with the line emission from the WR-star in RMC 140b due to its powerful wind emission; it has a mass-loss rate of $1.26 \times 10^{-5} \text{ M}_{\odot} \text{ yr}^{-1}$ and terminal wind velocity of 1300 km s^{-1} (Shenar et al. 2019).

GG Carinae is a single-lined spectroscopic binary with a ~ 31 day period (Hernandez et al. 1981; Gosset et al. 1985; Marchiano et al. 2012), situated at a distance of $3.4_{-0.5}^{+0.7} \text{ kpc}$ (Gaia Collaboration et al. 2016; 2018). The primary star is a B[e] supergiant with a mass-loss rate of $2.2 \times 10^{-6} \text{ M}_{\odot} \text{ yr}^{-1}$ and terminal wind velocity of 265 km s^{-1} (Porter et al. 2021a).

In this chapter, we investigate the dynamical information encoded in stellar emission lines, principally for stars with powerful winds where line emission can be highly extended. Eta Carinae, as previously introduced in Section 4.1.1, is our main case study. We present the observations and extraction methods used for computing radial velocities for Eta Carinae in Section 4.2. In Section 4.3, we begin by fitting the radial velocity curves with Keplerian models and show how selecting different Balmer emission lines can result in changes to the resulting orbital parameters. We then describe a new semi-analytical model that aims to reconcile different emission lines' radial

velocity curves with a single set of orbital parameters and apply it to Eta Carinae. In Section 4.4, we discuss several complexities of the model and the newly derived orbital parameters for Eta Carinae. We then create a forward model of RMC 140b to understand how our model may affect the interpretation of other emission-line stars' radial velocity curves. We apply our model to another typical emission-line star, GG Carinae. Finally, in Section 4.5 we summarize our findings.

4.2 Observations

The observations of Eta Carinae used throughout this chapter are from an online open-source data archive¹. We make use of data observed by the Gemini-South Multi-Object Spectrograph (GMOS) on the Gemini-South telescope. The data set contains 310 individual spectra over 33 nights of observing. The data set spans the 2009 periastron of Eta Carinae, starting in 2007 June and ending in 2010 January. The spectra are particularly well sampled at times near periastron, having 252 spectra in the interval $0.95 < \phi < 1.05$, where ϕ is the orbital phase defined by Equations 3.12 and 3.13. The spectra cover a wavelength range of $3557 \text{ \AA} < \lambda < 7546 \text{ \AA}$. The exposure times range from 1 to 180 s to optimize the signal-to-noise ratio for different emission lines and instrumental configurations, and the resolving power varies in the interval $3700 < R < 4400$. For further technical information, see the original work by Mehner et al. (2011) detailing the observing campaign.

The Eta Carinae spectra contain a wealth of information entangled in a multitude of complex line profiles. Given the complexity of the system, the spectra contain light from the central star, companion, extended wind, Weigelt blobs, and scattering from dust in the Homunculus nebula. We are particularly interested in the lines that form in the extended wind of the primary. These lines are known to show variability reflecting the orbital cycles as extensively discussed by previous authors: the He I (Damineli et al. 1997), the He II (Steiner & Damineli 2004; Teodoro et al. 2012; 2016; Davidson et al. 2015), the N II lines (Davidson et al. 2015), H-alpha (Richardson et al. 2010) and a handful of further lines (He I, N II, Na I D, Si II, Fe II) during the 2009 periastron (Richardson et al. 2015).

¹<http://etacar.umn.edu/archive/>




Figure 4.1: Multi-Gaussian fitting of simultaneously-observed Balmer series spectral profiles near to apastron. Each panel indicates the Balmer line transition, the number of emission components (positive Gaussians) and the number of absorption components (negative Gaussians). The top of each panel shows the total model (light blue) and its Gaussian components (light green and orange) vs the observations (black). The Gaussian components attributed to wind emission are highlighted orange. The bottom of each panel shows the residuals of each fit.




Figure 4.2: The same as Figure 4.1, but for periastron.

4.2.1 Radial velocities: the Balmer series

In this study of Eta Carinae, we concentrate our efforts on the Balmer emission lines from H-alpha to H-kappa. These lines serve as an ideal probe into the dynamics of extended emission, thanks to their high signal-to-noise and wide coverage of different spatial scales of emission in the wind.

To extract radial velocities from the Balmer series lines, we employ multi-Gaussian fitting to decompose individual line profiles into their constituent components. The approach of using multiple Gaussian components to track independent variations within complex line profiles in Eta Carinae has been utilized previously (eg. Ruiz et al. 1984; Nielsen et al. 2007). By fitting different line profiles independently we can monitor discrepancies between different line transitions.

For Eta Carinae our multi-Gaussian fitting algorithm is implemented as described in Section 3.3.3. To evaluate the goodness of fit, we follow Riener et al. (2019) by assessing whether the normalized residuals show a normal distribution, since we expect the errors in the flux bins to be Gaussian and homoscedastic. Using a two-sided Kolmogorov–Smirnov test we achieve a 96 per cent success rate (p -value > 0.0027) for the high excitation Balmer lines between H-gamma and H-kappa. For H-alpha and H-beta, we are not able to satisfy the statistical test. The high signal-to-noise ratio of these lines requires many Gaussian components to fit the observed complexity. The model-optimization algorithm is only capable of co-fitting a finite number of free parameters without a substantial reduction in performance. Hence, for templates with greater than ~ 10 Gaussians (ie. ~ 30 free parameters) the statistical test is less instructive. We also note that H-alpha and H-beta may show further systematic effects as discussed in Sections 4.4.2.2 and 4.4.2.4.

We apply our multi-Gaussian fitting algorithm to the Eta Carinae spectra. Example fits are shown for apastron in Figure 4.1 and for periastron in Figure 4.2. The individual Gaussian components (light green and orange) that make up the total fit (light blue) can be seen to vary depending on the line transition and position in the orbital cycle. H-alpha requires 14–16 components, whilst in contrast H-iota only requires 6–8 components, depending on orbital phase. Despite the substantial variation, we find a commonality to the templates: four main emission lines, a narrow low velocity component present from H-alpha to H-delta at apastron and becoming

less strong near periastron, and a very broad component present from H-alpha to H-epsilon that persists across all phases.

Once the spectral lines have been fitted, we leverage the power of multi-Gaussian fitting. By deconstructing the complex line profiles into individual components, we are able to scrutinize each component for astrophysical meaning. We note that despite using Gaussians as our base function we do not require dynamical components to be strictly Gaussian. As shown by Blundell et al. (2007), it is possible to take a weighted mean of multiple components (centroid wavelength of each Gaussian weighted by its area) to extract clear signals from a superposition of profile shapes.

We explore our fitting results and, at first, find no coherent signal for any of the four main emission components individually (orange lines in 4.1 and 4.2). However, after taking their weighted mean we extract clear radial velocity curves, as shown in Figure 4.3. Most notably, we uncover a low-amplitude radial velocity signature in H-alpha. The resulting radial velocity curves for the Balmer lines suggest that the four main emission components, all together, comprise the emission from Eta Carinae’s primary wind.

To estimate the uncertainties, we consider the velocity variance on an intra-nightly basis, assuming the velocities are not subject to real astrophysical changes on daily time-scales, given the orbital period of $P = 2022.7$ d. The final uncertainties for our radial velocity data are then calculated by the pooled variance, s_{RV}^2 , of intra-nightly spectra, defined as

$$s_{\text{RV}}^2 = \frac{\sum_{i=1}^k (n_i - 1) s_i^2}{\sum_{i=1}^k (n_i - 1)}, \quad (4.1)$$

where n_i is the number of observations in a given night and s_i is the intra-nightly sample variance. We could allow the uncertainties to vary in time, but due to the number of observations per night being low we can obtain a better overall estimate of the uncertainties by combining the intra-nightly variance across all times.

To verify the radial velocities output from our multi-Gaussian fitting algorithm, we implement the BHM algorithm on Eta Carinae’s Balmer lines. The BHM output will serve as a simple benchmark to compare our results against.

4.3 Modelling

We model the dynamics of the Balmer lines’ emission (H-alpha to H-kappa) originating from the primary wind of Eta Carinae based on the radial velocities computed in Section 4.2. For all models, we assume an orbital period of $P = 2022.7$ d (Damineli et al. 2008a).

4.3.1 Keplerian motion

As an instructive exercise, we start by assuming the entire line emission region tracks the orbital motion of the primary star. If this is the case, the radial velocities will correspond to Keplerian orbits projected onto our line of sight. Using the standard formalism for a binary system, as derived in Section 3.2.1, this is written as

$$v_{\text{kep}}(t) = k(\cos(\omega + \nu) + e \cos \omega) + \gamma, \quad (4.2)$$

where v_{kep} is the Keplerian velocity projected onto our line of sight, k is the radial velocity semi-amplitude, e is the eccentricity, ω is the argument of periastron, ν is the true anomaly, and γ is the radial velocity offset.

We fit the radial velocity observations of Eta Carinae’s Balmer lines with the Keplerian model using the affine-invariant Markov Chain Monte Carlo (MCMC) algorithm, `emcee`, implemented by Foreman-Mackey et al. (2013). We allow the parameters in the set $\theta = \{T_0, e, k, \omega, \gamma\}$ to vary. The log-likelihood function, applied to each line independently, for the Keplerian model, v_{kep} , compared to the radial velocity data, D_ϕ , and associated uncertainties, σ_ϕ , is

$$\ln P(D | \theta, \sigma) = -\frac{1}{2} \sum_{\phi} \left[\frac{(D_\phi - v_{\text{kep}}(\theta)_\phi)^2}{\sigma_\phi^2} + \ln(2\pi\sigma_\phi^2) \right]. \quad (4.3)$$

The algorithm is run with 128 walkers for 6000 burn in steps and then for a further 6000 sampling steps. In our testing, we find chains of this length result in good convergence [using the autocorrelation time (Foreman-Mackey et al. 2013) as a guiding metric] and well-resolved parameter posteriors.

4.3.1.1 Keplerian results

In Figure 4.3, we show the results of applying the Keplerian model independently to a range of different Balmer lines’ radial velocities. Each panel corresponds to a different model run, where the Balmer line is labelled in the top left. We plot v_{kep} for 100 random samples from the MCMC run against the radial velocity observations extracted using the multi-Gaussian algorithm. The radial velocities are the weighted mean of the four main emission components (orange Gaussians) as shown in Figures 4.1 and 4.2. The Keplerian model is capable of producing good fits to all of the Balmer series lines. However, as expected the parameters derived using this model do not agree between lines.

In Figure 4.4, we plot the optimized parameters from all of the MCMC runs for both the multi-Gaussian (dark blue) and BHM (light blue) methods of radial velocity extraction. We show the parameters’ median values and 16th–84th percentile uncertainties from their posterior distributions. The parameters show correlations across the Balmer series when ordered by transition energy. We find very clear monotonic trends for T_0 , e and k – whilst γ exhibits more complex behaviour. The trends are less noisy in the multi-Gaussian results relative to the BHM results, indicating the superiority of this extraction technique. The trends show that the higher energy Balmer lines have associated radial velocity curves that appear phase shifted later in time, as well as less eccentric and lower in amplitude. These effects have also been observed between different species in Eta Carinae. Nielsen et al. (2007) noted the radial velocities of the Balmer lines appeared offset in phase and reduced in amplitude relative to He I, and Davidson et al. (2000) found their higher series Paschen lines showed larger velocity shifts relative to the lower energy lines. The monotonic trends are also asymptotic: the higher energy lines appear to converge towards a single set of parameter values.

We suggest that the systematic sequencing of these results is a consequence of probing different emission-line regions for each Balmer line. The higher excitation lines are preferentially emitted in denser and hotter regions relative to lower excitation Balmer lines. Hence, lines lower down the Balmer series are emitted in progressively more extended regions further away from the star, with H-alpha being emitted over the largest and most extended volume. For extended line emission, the radial velocity




Figure 4.3: Keplerian models independently fit to the radial velocity observations of Eta Carinae’s Balmer series. Each panel shows a Balmer line, 100 samples from the converged MCMC run (dark blue) and the radial velocities (black) extracted from the multi-Gaussian algorithm and associated uncertainties. The radial velocities are the weighted mean of the 4 main emission components (orange Gaussians) as shown in Figures 4.1 and 4.2.




Figure 4.4: Trends in the parameters output from Keplerian models applied to the multi-Gaussian (dark blue) and BHM (light blue) radial velocity observations of Eta Carinae. The final convolutional Keplerian motion (CKM) model parameters are shown as the red dashed asymptotes (see Section 4.3.2 for details).

measurements will contain dynamical information which is not purely Keplerian. The more extended the line emission, the greater the departure from the true Keplerian motion the observations will imply. Conversely, lines towards the top of the Balmer series, approaching the asymptote in Figure 4.4, are formed at increasingly close regions to the central star, deeper in the wind, and as such the parameters become increasingly representative of the true orbital motion.

4.3.2 Convolutional Keplerian motion

We formulate a semi-analytical model to account for the dynamical information encoded in lines formed in extended regions. The model is a first attempt to encapsulate both the star's orbital motion, and the propagation of the wind by accounting for its time delay to the location where the photons are actually emitted.

To explain the net motion observed from the volume of gas in which an emission line is formed, we simplify the picture of a binary star with a powerful wind to the following paradigm:

- Gas in the stellar wind at radii less than the sonic point co-orbits with the star. Emission from these regions will appear Keplerian.
- Outside of the sonic point, there is no way for gas to respond to the acceleration of the central star via pressure waves. The gas moves outwards and is subject to radiative acceleration.
- At each radius in the stellar wind, the radial velocities of an infinitesimally thin shell average to zero, assuming a steady spherically symmetric wind. However, the shell has an overall velocity equal to the orbital velocity of the star at the time when it was accelerated past the sonic point.
- Therefore, the total stellar wind can be thought to consist of shells, each encoding the Keplerian motion of the star from a time in the past defined by the flow time of the wind from where the sonic point was previously to the location where the photons are emitted.
- As such, the net motion of the total line emission region is an average of the Keplerian velocities over an interval backwards in time which equates to the

flow time of the wind to reach the emission radii. In effect, this is a moving average of the Keplerian velocities.

- Further to this, depending on the line transition, lines are emitted at various strengths over a range of different radii, and thus different flow times. So, the moving average should be weighted by the luminosity of the line transition for a given species at a given radius.

This description is idealized in order to assess its effectiveness, which we explore in Sections 4.4.1 and 4.4.2. Using this framework, we can now define our model. More formally, a weighted moving average of the Keplerian velocities can be written as a convolution

$$v_{\text{ckm},n}(t) = v_{\text{kep}}(t) * \lambda_{L,n}(t) = \int_{-\infty}^{\infty} v_{\text{kep}}(\tau) \lambda_{L,n}(t - \tau) d\tau, \quad (4.4)$$

where $v_{\text{ckm},n}$ is the convolutional Keplerian motion (CKM), v_{kep} is the Keplerian velocity from Equation 4.2, and $\lambda_{L,n}(t)$ is the line-formation kernel for each emission line, n , which weights v_{kep} as a function of wind flow time.

An alternative way to think about CKM is by drawing parallels with the retarded potential in electrodynamics (Griffiths 2007). The stellar wind’s velocity field generated by the orbital motion is similar in principle to the electromagnetic field generated by a time-varying charge distribution. Any changes to the central source take a finite time to propagate outwards – in our case at the flow speed of the wind, and in electrodynamics at the speed of light.

4.3.2.1 Radiative transfer and the line-formation kernels

In order to compute the line-formation kernels, $\lambda_{L,n}(t)$, we require knowledge of a specific line’s emission strength as a function of radius in the stellar wind. We employ the non-local thermodynamic equilibrium, fully line-blanketed stellar atmosphere code, **CMFGEN** (Hillier & Miller 1998; Hillier & Lanz 2001), as introduced in Section 2.3.2. We select **CMFGEN** because of its historical use in modelling Eta Carinae, making our results more comparable to past work. In practice, any stellar atmosphere code could be used and the results of our CKM model should be independent of this choice.

We use **CMFGEN** to find the linear luminosity density (otherwise known as luminosity per unit radius) as a function of radius, $\Lambda_{L,n}(r)$, for each line transition observed

in Eta Carinae. The line-formation kernels are then calculated by converting $\Lambda_{L,n}(r)$ to be a normalized function of wind flow time by integrating over the velocity profile of the wind using the following equations:

$$v_{\text{wind}}(r) = v_0 + (v_\infty - v_0) \left(1 - \frac{R_*}{r}\right)^\beta, \quad (4.5)$$

$$t_{\text{flow}} = \int_{r_s}^r \frac{dr}{v_{\text{wind}}(r)}, \quad (4.6)$$

where v_0 is the initial wind velocity, v_∞ is the terminal wind velocity, R_* is the stellar radius, r_s is the sonic radius, and β is a free parameter describing the steepness of the velocity law. In Equation 4.5, we assume that the wind is accelerated by a standard β -law as first derived by Castor et al. (1975) for line-driven winds and we set $\beta = 1$.

For Eta Carinae, and most stars with powerful winds, this is likely a good approximate form for the conversion. We note, however, it has been suggested that Eta Carinae's wind may be super-Eddington and its velocity profile would therefore be a modified β -law (Owocki et al. 2004). However, in general the line-formation region is extended well beyond the acceleration region of the wind and so the results will be largely insensitive to the velocity prescription.

The main dependence of the $\Lambda_{L,n}(r)$ curves is on a few stellar parameters input into the code. In principle, it is possible to create a multidimensional grid of $\Lambda_{L,n}(r)$ curves, spanning a few key parameters, and then interpolate between the results by including the parameters in our CKM modelling routine. Thereby, it would be possible to match radial velocities and stellar types in one unique solution. Of particular interest are the stellar radius, R_* , mass-loss rate, $\dot{M}$, and the terminal wind velocity, v_∞ . Unfortunately, for Eta Carinae the effect of R_* has a negligible effect on the output. As previously found by Hillier et al. (2001), the extreme mass-loss prevents the accurate determination of R_* . We test the effect of changing the inner boundary of the simulation on $\Lambda_{L,n}(r)$ and reach the same conclusion: the results are insensitive to changes in R_* and thus we cannot calculate an accurate effective temperature for Eta Carinae (see the left panel of Figure 4.5). Likewise, we cannot determine values for $\dot{M}$ and v_∞ because $\Lambda_{L,n}(r)$ is degenerate for these two parameters (see Figure 4.6).

To overcome the difficulties in determining the stellar parameters, we introduce a preliminary step into the modelling routine. Rather than only utilizing CMFGEN's




Figure 4.5: Left panel: a comparison of the relative line contributions for lines H-alpha (black), H-beta (red) and H-iota (orange) for different **CMFGEN** inner boundaries (darker lines indicate smaller radii). Right panel: a comparison of the relative line contributions for H-alpha for different **CMFGEN** volume filling factors (darker lines indicate more clumping).




Figure 4.6: Line-formation kernels for an interpolated 3x3 grid of **CMFGEN** simulations varying $\dot{M}$ and v_∞ . Each kernel is comprised of the linear luminosity density, $\lambda_{L,n}$, as a function of wind flow time.

output line-formation regions, we first constrain the stellar parameters by ensuring a good fit between the synthetic spectrum and the observations. In this way, we are now incorporating the information from the total observed spectrum to lock down R_* , $\dot{M}$ and v_∞ , as well as the radial velocities extracted from a time series of spectra. By forcing our model to fit both a spectral snapshot and the time-domain motion, we constrain the results more.

The final wind parameter that has a significant effect on $\Lambda_{L,n}(r)$ is the clumping. A standard line-driven wind is expected to have strong clumping that develops from the intrinsic instability in line driving (MacGregor et al. 1979; Owocki et al. 1988), and this effect has been observationally substantiated (eg. Ebbets 1982; Kaper et al. 1996; Oskinova et al. 2006). A clumpy wind creates higher densities at larger radii; therefore, line emission is possible in more extended regions from the star than previously allowed in a homogeneous wind (see the right panel of Figure 4.5). For Eta Carinae, we use the same parametrization as Hillier et al. (2001) and Groh et al. (2012a): a volume filling factor, $f_v = 0.1$, which describes the clumping at infinity and $v_{cl} = 100 \text{ km s}^{-1}$ which denotes the wind velocity at which clumping becomes important. We note, however, that owing to the extremely high mass-loss rate in Eta Carinae, the line-driving instability may not be a strong factor. But, if the wind is instead continuum driven then it may be regulated by “porosity”, which reduces the effective opacity in deep layers (Owocki et al. 2004). In either case, we expect overdensities to be present in the wind and we have chosen standard parameter values to account for this effect in the radiative transfer simulation.

For Eta Carinae, we fix $R_* = 60 R_\odot$ and for $\dot{M}$ and v_∞ we use the values found by Groh et al. (2012a; also using **CMFGEN**) of $\dot{M} = 8.5 \times 10^{-4} M_\odot \text{ yr}^{-1}$ and $v_\infty = 420 \text{ km s}^{-1}$. We show the resulting line-formation kernels for H-alpha, H-beta, and H-iota in Figure 4.7.

4.3.2.2 Modelling routine

In Figure 4.8, we present a schematic of the CKM modelling routine. The schematic shows three example lines (H-alpha, H-epsilon, H-iota) to illustrate the difference in output between lines formed in different regions of the wind, but the model can be applied to any emission line or set of lines for any emission-line star. As described




Figure 4.7: Line-formation kernels for an example set of 3 Balmer lines for Eta Carinae. Each kernel is the linear luminosity density (otherwise known as luminosity per unit radius) as a function of wind flow time beyond the sonic point. Flow time is shown in days (top x -axis) and in units of Eta Carinae’s orbital phase (bottom x -axis).

in Section 4.3.2.1, a stellar atmosphere code is employed to fit a synthetic spectrum to observations of a chosen target. By iteratively comparing the stellar atmosphere, model to the observations we identify the best-fitting values for the key stellar parameters such as R_* , $\dot{M}$ and v_∞ . Using the best-fitting model, a set of $\Lambda_{L,n}(r)$ curves are generated: one for each emission line for which we have corresponding radial velocity observations. At this stage, we are modelling the wind with no latitudinal dependence and so the line-formation regions are spherically symmetric. $\Lambda_{L,n}(r)$ is converted to a function of wind flow time, $\lambda_{L,n}(t)$, using equations 4.5 and 4.6, and summed over all directions producing a 1D line-formation kernel for each line. A Keplerian orbital model is created using Equation 4.2 and convolved with each line-formation kernel (shaded box in Figure 4.8). The resulting set of CKM curves are compared to radial velocity observations.

For our case study of Eta Carinae, we fit the CKM curves to the radial velocities extracted from the Balmer lines – H-alpha to H-kappa. As described in Section 4.3.1, we using the affine-invariant MCMC algorithm, `emcee`, implemented by Foreman-




Figure 4.8: CKM model schematic for an example set of 3 Balmer lines. The stellar atmosphere model iteratively fits a synthetic spectrum to an observed spectrum to constrain the key stellar parameters: mass-loss rate, $\dot{M}$, terminal wind velocity, v_∞ and stellar radius, $R_\star$. The best fit stellar atmosphere model outputs a set of line-formation regions which are converted to 1D kernels as a function of wind flow time. The Keplerian orbital model is convolved with each line-formation kernel (shaded box) to produce a set of CKM models to compare to radial velocity observations. By optimising the Keplerian input parameters – time of periastron, T_0 , radial velocity semi-amplitude, k , eccentricity, e , argument of periastron, ω – to best fit the observations we can infer the true orbital motion from extended emission.

Mackey et al. (2013) to optimize the model fitting. We allow the parameters in the set $\theta = \{T_0, e, k, \omega\}$ to vary. We fix γ for each line to be the same as the results in Figure 4.4. By using the γ values calculated from individually optimizing the Keplerian model for each line, we allow the flexibility for γ to be different between lines, whilst reducing the number of free parameters by the number of lines in the radial velocity set. The log-likelihood function for our CKM model, co-fit to all the Balmer lines concurrently, is equal to the sum of log-likelihoods for each line:

$$\ln P(D \mid \theta, \sigma) = \sum_n \ln P(\theta \mid D_n, \sigma_n) \quad (4.7)$$

$$= -\frac{1}{2} \sum_n \sum_\phi \left[\frac{(D_{\phi,n} - v_{\text{ckm},n}(\theta)_\phi)^2}{\sigma_{\phi,n}^2} + \ln(2\pi\sigma_{\phi,n}^2) \right], \quad (4.8)$$

where $v_{\text{ckm},n}$, $D_{\phi,n}$ and $\sigma_{\phi,n}$ are the CKM model, radial velocity data, and associated uncertainties for a given Balmer line, n , respectively. The algorithm is run with 128 walkers for 6000 burn in steps and then for a further 6000 sampling steps to ensure convergence and well-resolved parameter posteriors.

4.3.2.3 CKM results

In Figure 4.9, we show the posterior probability distributions for each of the parameters in the set θ from the MCMC run. The model converges to a unique solution with the 1D histograms showing approximate Gaussian distributions for all of the free parameters. The best-fitting values are $T_0 = 2454848$ (JD), $e = 0.91$, $k = 69 \text{ km s}^{-1}$ and $\omega = 241^\circ$, with small statistical uncertainties on each value. The optimized set of CKM curves (red) can be seen in the top right-hand panel, for lines H-alpha to H-kappa, along with the inferred Keplerian motion (dark blue) of the central star. The CKM velocity profile for H-alpha shows the greatest departure from the Keplerian motion; it is phase shifted later in time, less eccentric and lower in amplitude owing to it having the most extended line-formation kernel. The rest of the Balmer lines show CKM curves becoming asymptotically more similar to the underlying Keplerian motion as the transition energy increases. These results qualitatively emulate the parameter variations found between different lines, when assuming Keplerian motion, as seen previously in Figure 4.4.




Figure 4.9: Bottom left corner: posterior probability distributions of the CKM model's free parameters: time of periastron, T_0 , in Julian days; eccentricity, e ; radial velocity semi-amplitude, k , in km s^{-1} ; and argument of periastron, ω , in degrees. In the 1D histograms the dashed lines indicate the 16th, 50th and 84th percentiles. In the 2D histograms the contours represent the 1σ , 2σ and 3σ statistical uncertainties. Top right panel: set of CKM model curves (red) for the Balmer lines, H-alpha to H-kappa (without γ correction), and the inferred Keplerian motion (dark blue) of the best fit model.




Figure 4.10: CKM model co-fit to the radial velocity observations of Eta Carinae's Balmer series. Each panel shows the Balmer line, 100 samples from the converged MCMC run (red) and the radial velocities (black) extracted from the multi-Gaussian algorithm and associated uncertainties. The radial velocities are the weighted mean of the 4 main emission components (orange Gaussians) as shown in Figures 4.1 and 4.2.




Figure 4.11: CKM best fitting model (red) to the H-delta radial velocities (black) extracted from the multi-Gaussian algorithm and associated uncertainties. Also shown is the inferred Keplerian motion (dark blue) of Eta Carinae and the H-delta line-formation kernel (top left inset panel) as a function of orbital phase.

In Figure 4.10, we compare the best-fitting CKM model against the radial velocity observations extracted using the multi-Gaussian algorithm. In each panel, we plot $v_{n,\text{ckm}}$ for 100 random samples from the MCMC run for a Balmer line, n , labelled in the top left. Overall, the model is able to reproduce the progression in Balmer line radial velocity profiles. The fits are particularly good at times close to periastron for Balmer lines between H-gamma and H-iota. The model struggles to fit the profiles in the wake of periastron, at around $\phi \sim 1.05$, for a few lines. Further to this, the model is not able to fit the extent to which the velocities are modified in H-alpha, and to some extent in H-beta. However, these remaining deviations may be caused by external factors we have not included in the model; in particular, the influence of the companion which will predominantly affect the lines formed in the most extended regions (see Section 4.4.2 for further discussion).

In Figure 4.11, we focus on H-delta to show in detail the constituent parts of the CKM model. H-delta shows line formation with wind flow times up to 6 per cent of an orbital period. As a result, when the Keplerian motion (dark blue) is convolved with this line-formation kernel (top left-hand inset panel), the resulting CKM motion is

modified significantly. In this case, we see the CKM curve produces a good fit to the observations. This plot showcases the capability of our model to uncover significant differences between the underlying orbital motion and the observed motion.

4.3.2.4 CKM cross-validation

After fitting our CKM model to a set of Eta Carinae’s Balmer lines, it serves as an interesting test to fit each line independently and see if the model recovers similar orbital parameters. In the top section of Table 4.1, we present the results of this cross-validation test. For comparison, the parameters from the original run are shown in the bottom section of the table labelled as H-alpha/kappa. Note that the uncertainties are purely statistical and result from the model fitting procedure. More realistic uncertainties will be larger owing to a range of systematic effects present in this complex system (see Section 4.4.2 for further discussion).

We find broadly stable results for the higher excitation Balmer lines between H-gamma and H-kappa. For these lines, the standard deviations in the resulting parameters are 4.3 d, 0.02, 3.7 km s⁻¹, and 13° for T_0 , e , k , and ω , respectively. We identify some slight degeneracies between two pairs of the parameters: later T_0 with increases in ω and higher e with increases in k . These slight degeneracies are also present in the original run, as seen by the elliptic covariance in the 2D histograms of these parameter pairs in Figure 4.9.

As for H-alpha and H-beta, these lines show larger departures from the original parameter results. We know from Figure 4.10 that our CKM model struggles to fit these lines that form in highly extended regions, and so it is not surprising that these lines perform poorly in our cross-validation testing. As previously mentioned, we are not overly concerned by the performance of our CKM model on these two lines due to complicating factors that are present in Eta Carinae (see Section 4.4.2 for further discussion).

Further to testing individual lines, we run our CKM model against a subset of Balmer lines, H-gamma to H-kappa, to analyse the sensitivity of the results to the systematic effects most prominent in H-alpha and H-beta. The parameters from this test are shown in the bottom section of Table 4.1, labelled as H-gamma/kappa. The

Line(s)	T_0	e	k	ω	γ
	(JD +2454800)		(km s ⁻¹)	(deg)	(km s ⁻¹)
H-alpha	27.4 ^{+16.7} _{-15.0}	0.63 ^{+0.03} _{-0.03}	31.7 ^{+1.1} _{-0.1}	247 ⁺⁷ ₋₇	-25.9 ^{+1.8} _{-1.7}
H-beta	54.9 ^{+4.5} _{-4.1}	0.82 ^{+0.02} _{-0.02}	53.0 ^{+2.1} _{-1.9}	254 ⁺⁴ ₋₄	-25.5 ^{+2.0} _{-2.0}
H-gamma	53.8 ^{+2.4} _{-2.3}	0.90 ^{+0.01} _{-0.01}	73.3 ^{+3.4} _{-2.9}	252 ⁺⁴ ₋₄	-16.8 ^{+1.4} _{-1.4}
H-delta	54.1 ^{+1.8} _{-1.8}	0.87 ^{+0.01} _{-0.01}	64.6 ^{+2.1} _{-2.0}	254 ⁺³ ₋₃	-7.9 ^{+1.5} _{-1.4}
H-epsilon	52.4 ^{+0.3} _{-0.5}	0.87 ^{+0.01} _{-0.01}	72.5 ^{+1.3} _{-1.3}	267 ⁺¹ ₋₁	0.5 ^{+1.1} _{-1.1}
H-zeta	41.5 ^{+1.9} _{-1.8}	0.89 ^{+0.01} _{-0.01}	70.5 ^{+2.8} _{-2.6}	249 ⁺³ ₋₄	12.1 ^{+2.5} _{-2.4}
H-eta	45.7 ^{+1.6} _{-1.0}	0.89 ^{+0.01} _{-0.01}	73.0 ^{+2.1} _{-2.0}	235 ⁺³ ₋₂	22.7 ^{+1.6} _{-1.6}
H-theta	48.2 ^{+2.2} _{-2.0}	0.90 ^{+0.01} _{-0.01}	73.9 ^{+5.4} _{-4.3}	244 ⁺⁵ ₋₅	18.3 ^{+3.1} _{-3.1}
H-iota	45.2 ^{+1.2} _{-0.9}	0.90 ^{+0.01} _{-0.01}	72.6 ^{+2.6} _{-2.4}	229 ⁺³ ₋₂	3.8 ^{+1.8} _{-1.8}
H-kappa	46.9 ^{+0.8} _{-0.7}	0.92 ^{+0.01} _{-0.01}	78.8 ^{+3.5} _{-3.1}	228 ⁺² ₋₂	-15.4 ^{+1.4} _{-1.3}
H-alpha/kappa	48.3 ^{+0.4} _{-0.4}	0.91 ^{+0.00} _{-0.00}	69.0 ^{+0.9} _{-0.9}	241 ⁺¹ ₋₁	*
H-gamma/kappa	48.4 ^{+0.4} _{-0.4}	0.89 ^{+0.00} _{-0.00}	69.9 ^{+0.8} _{-0.8}	246 ⁺¹ ₋₁	*

Table 4.1: Top section: cross-validation of CKM fits to individual Balmer lines in Eta Carinae. Bottom section: results from co-fitting to a set of Balmer lines in the interval H-alpha to H-kappa and H-gamma to H-kappa. Parameters: time of periastron, T_0 , eccentricity, e , radial velocity semi-amplitude, k , argument of periastron, ω , and systemic velocity, γ . The systemic velocities used in the fits to multiple lines (*) are the same as in the top section for each line (see text for details). Uncertainties are purely statistical based on the model fitting.

results show only minimal changes from the original run, demonstrating the resilience of the model to systematic effects when using a larger number of lines.

Overall, we find that it is possible to recover similar orbital solutions regardless of the line used (omitting H-alpha and H-beta). The original run parameters are plotted in Figure 4.4, as red dashed asymptotes, indicating the unified orbital solution our CKM model implies for Eta Carinae, in comparison to the variable Keplerian solutions.

4.4 Discussion

We have presented a semi-analytical model, CKM, which takes the first steps towards reconciling the systematic trends observed in different emission lines' radial velocity curves with a consistent orbital solution. In this section, we will discuss the nuances of the model in greater detail, as well as the impact the results have on future work involving Eta Carinae and other stars with powerful winds.

4.4.1 CKM time-scales

The CKM model is built on a simple paradigm intended to encapsulate the dynamics associated with stellar winds out-flowing from binary stars. By formulating the model as a convolution we are able to generate CKM radial velocity curves relatively simply and quickly. This allows us to sample a multidimensional parameter space for the best-fitting solution. However, this description overlooks some physical details; in particular, the star's deviation away from the reference frame in which the wind is launched outwards. This simplification has implications for the radiation field incident on the wind, the gravitational potential felt by the gas, and the calculation of wind flow times to line-formation regions.

To check the magnitude of these effects, we compare two fundamental time-scales associated with the model. First, the orbital time-scale, τ_{orbital} , defined by

$$\|\mathbf{r}(t) + \tau_{\text{orbital}}\dot{\mathbf{r}}(t) - \mathbf{r}(t + \tau_{\text{orbital}})\| = D, \quad (4.9)$$

where $\mathbf{r}$ is the orbital position, $\dot{\mathbf{r}}$ is the orbital velocity, and D is the deviation distance. τ_{orbital} represents the time it takes for the displacement between the star's future position, based on tangential motion versus orbital motion, to grow to a distance, D .




Figure 4.12: Left panel: timescales comparison of τ_{orbital} (dark blue line) against τ_{flow} (brown dashed line) for a distance scale of 10 au. Right panel: timescales comparison of τ_{orbital} (dark blue line) against τ_{accel} (green dashed line) for a distance scale equal to the distance at which the wind has reached $\frac{2}{3}v_\infty$. The period of Eta Carinae is 2022.7 d.

Secondly, the wind flow time-scale, τ_{flow} , is the time it takes for gas to flow from the sonic point out to a given line-formation radius, D . τ_{flow} is calculated by integrating Equation 4.6, resulting in

$$\tau_{\text{flow}} = \frac{R_*(v_\infty - v_0) \ln(R_*v_0 - v_\infty R_* + v_\infty r) + v_\infty r}{v_\infty^2} \bigg|_{r=r_s}^{r=D}. \quad (4.10)$$

In Figure 4.12, we analyse these time-scales across an orbital cycle of Eta Carinae. In the top panel, we compare τ_{orbital} (dark blue line) against τ_{flow} (brown dashed line) for $D = 10 \text{ AU}$: a distance comparable to the radius at which an extended line peaks. We find $\tau_{\text{orbital}} \gg \tau_{\text{flow}}$ at times around apastron, with the time-scales becoming more comparable at periastron due to the highly eccentric orbit. The slight asymmetry in τ_{orbital} is due to the changes in the future orbital path for epochs before or after periastron. The slope to shorter time-scales before periastron is shallower; this indicates that the larger accelerations associated with the upcoming periastron cause larger deviations in less time relative to after periastron. We also find a step change in τ_{orbital} at $\phi = -0.32$ ($= 0.68$), which is an interesting artefact resulting from the sharp changes in motion at periastron. At this phase, it takes a sufficiently long time for D to reach 10 AU, such that periastron has passed and the change in

direction of the star results in the deviation distance actually decreasing momentarily – hence the suddenly longer time-scales. This analysis shows that changes to the gravitational potential and wind flow times due to CKM not accounting for the star’s deviation away from its inertial frame of reference are small.

In the bottom panel of Figure 4.12, we compare τ_{orbital} (dark blue line) against τ_{flow} (green dashed line) for $D = 0.78$ AU: the distance it takes the wind to reach $\frac{2}{3}v_{\infty}$. Again, we find $\tau_{\text{orbital}} \gg \tau_{\text{accel}}$ at times around apastron, with time-scales becoming closer at periastron. The significantly shorter deviation distance, and hence shorter time-scales, used in this panel result in τ_{orbital} appearing more symmetric, relative to the top panel, as the future orbital motion has less time to become a factor. This analysis shows that changes to the radiation field incident on the wind are small; and therefore, the line-formation regions and acceleration of the wind will not be greatly affected.

As a result, the time-scales considered here establish that our formalism of CKM provides a valid description of the physics associated with powerful winds from binary stars, such as Eta Carinae. In order to establish the limits of applicability, investigations into the extent to which the outflowing gas deviates from the updating orbital motion would be interesting; this would necessitate a full numerical model including additional physics, such as Coriolis acceleration.

4.4.2 The outflows of Eta Carinae

4.4.2.1 A latitudinally dependent wind

It may be the case for Eta Carinae, and similar stars, that the mass-loss is not spherically symmetric. Theoretical work by Bjorkman & Cassinelli (1993) originally showed that above the sonic point the wind should curve towards the equatorial plane and form a disc, so long as it is only subject to radial forces and hence angular momentum is conserved. However, when including both the non-radial line-forces and gravity darkening in models of rapidly rotating luminous stars, the formation of a disc was prevented and the mass flux was actually enhanced in the polar directions (Cranmer & Owocki 1995; Owocki et al. 1996). Furthermore, observations of reflections from Eta Carinae’s Homunculus nebula by Smith et al. (2003) showed the Balmer lines

have faster and deeper P Cygni absorption at higher latitudes, strongly suggestive of a polar enhanced wind.

In our modelling, we have used a spherically symmetric wind. It would be possible to include a latitudinal dependence to the wind by calculating the emission-line regions as a function of both radius and polar angle, $\Lambda_{L,n}(r, \theta)$, from a 2D stellar atmosphere simulation. The line-formation kernels would then be computed by summing the line contributions over all directions as a function of wind flow time. We have not included this complexity in our model, in part to reduce the number of free parameters, but also because the effects of a faster and denser wind in the polar directions are degenerate. Hence, any changes to the line-formation kernels will be small.

4.4.2.2 The colliding winds

Eta Carinae is a colliding wind binary (Pittard et al. 1998). In addition to the powerful primary wind, the companion also emits a strong wind with a mass-loss rate of $\sim 10^{-5} M_{\odot} \text{yr}^{-1}$ and terminal wind velocity of $\sim 2000 - 3000 \text{ km s}^{-1}$ (Pittard & Corcoran 2002; Parkin et al. 2011). The winds collide where the ram pressures are equal, forming shock fronts in each wind separated by a contact discontinuity. The higher wind momentum flux of the primary wind results in the colliding wind region being shifted towards, and wrapped in a cone-like geometry around, the companion (Okazaki et al. 2008; Parkin et al. 2011).

The cavity caused by the lower density secondary wind may produce systematic effects on the bulk velocities of emission lines. Madura et al. (2012) estimate the semimajor axis of the orbit is $\sim 15.4 \text{ AU}$ ($\sim 1.5 \text{ AU}$ binary separation at periastron), and therefore, the lower energy Balmer lines, emitted at radii with the greatest overlap with the cavity, will be the most affected.

Given we have calculated $\omega = 241^{\circ}$, the cavity will reduce the velocities away from our line of sight (positive velocities) at periastron. This reduction may be delayed to times soon after periastron because of Coriolis effects on the colliding wind geometry (Walder 1995). As a result, the observed radial velocity curves may show a reduction in positive velocities in the wake of periastron, appearing as a reduction in e , k , and

ω . This may resolve some of the remaining deviations in our CKM model fitting previously found in Figure 4.10 for H-alpha and H-beta.

Outside the cavity, the colliding winds can produce additional features in the line profiles due to emission in the post-shock gas (Groh et al. 2012a;b) and the radiation field can influence the ionization structure of the primary star’s wind (Kruip 2011; Madura et al. 2012; Clementel et al. 2015a;b). These complications can bias the extraction of radial velocities, as was deduced by Nielsen et al. (2007) when using the He I P Cygni absorption as a proxy for orbital motion. In our study, we have used emission from a large volume of gas to track the orbital motion; hence, our results may be less susceptible to some of these systematic effects that only affect specific volumes of the wind on the companion’s side of the system. Nevertheless, they are sources of systematic errors in our results.

4.4.2.3 Radiative inhibition and braking versus gravity

When the separation between binary stars is small, the radiation field of one star may affect the velocity profile of the other star’s wind by radiative inhibition (Stevens & Pollock 1994) or braking (Gayley et al. 1997). Radiative inhibition reduces the initial acceleration of the other star’s stellar wind, whilst radiative braking suddenly decelerates the wind as it approaches the companion. Additionally, depending on the ratio of luminosities and masses, the counterveiling effect of gravity may increase the wind velocities and mass-loss rate (Stevens & Pollock 1994).

Parkin et al. (2009) analysed the influence of these effects for Eta Carinae. They determined that radiative braking does not occur in either the primary or companion’s winds, as the winds collide at their balance point before approaching close enough to the other star. Parkin et al. (2009) found that the primary star significantly inhibits the companion’s wind, but the primary star’s wind is not greatly impacted.

Additionally, we investigate if the companion’s gravitational field may influence the streamlines of the primary wind. Using the theory detailed in Section 2.2 we calculate the gravitational and radiative forces along the line between the stars. We use the same parameters as Parkin et al. (2011; see their table 2). In Figure 4.13 we present these forces at four different orbital phases, $\phi = 0, 0.05, 0.15, 0.5$. It is clear from this plot, by comparing the companion’s gravity (orange lines) and the




Figure 4.13: Comparison of the forces acting on the primary wind along the line between the two stars. Light blue: the primary’s gravity. Orange: the companion’s gravity at four phases, $\phi = 0, 0.05, 0.15, 0.5$. Dark blue: the net gravity of the primary and companion at periastron. Red: the radiative acceleration from both line and continuum driving. Black dashes: the location of the contact discontinuity where the primary’s and companion’s winds collide at each phase.

radiative force driving the primary wind (red line), that the companion’s gravity will not influence the primary wind’s streamlines. Even in very close proximity to periastron, $\phi = 0$, it appears that the companion’s gravity will not act significantly on the primary wind until radii at, and beyond, the contact discontinuity (dashed line).

In our work, we therefore expect the systematic errors as a result of these effects to be small. Moreover, these calculations were performed for periastron and along the line between the binary star components where the effects are maximal. At phases either side of periastron and elsewhere in the primary wind the influence of these effects will be vastly reduced, further reducing the influence on the emission-line velocities we extracted.

4.4.2.4 Time-dependent mass-loss

There is evidence for Eta Carinae’s mass-loss rate increasing immediately following periastron for approximately 80d (Corcoran et al. 2001). An enhanced mass-loss rate causes emission lines to become more extended as the line-formation regions are dependent on opacity (Hillier 1987). The more extended the emission-line regions, the greater the deviations from the true Keplerian motion the radial velocity curves will show. In this case, these deviations will be exacerbated over the ~ 80 d ($0 \lesssim \phi \lesssim 0.04$) following periastron. This may be another reason for the remaining deviations in our CKM model fitting previously found in Figure 4.10.

4.4.3 The orbital dynamics of Eta Carinae

4.4.3.1 A comparison of orbital solutions

By applying our new CKM model to the Balmer lines in the 2009 periastron GMOS data of Eta Carinae, we have calculated the orbital parameters to be $e = 0.91$, $k = 69 \text{ km s}^{-1}$, and $\omega = 241^\circ$.

For comparison, previous studies have used a wide range of observational diagnostics, and epochs, to determine the orbital parameters. Damineli et al. (1997; 1992 periastron data) were the first to propose an orbital solution after taking radial velocity measurements of the Paschen-gamma line, and they calculated the orbital parameters to be $e = 0.63$, $k = 53 \text{ km s}^{-1}$, and $\omega = 286^\circ$. This solution was updated by Davidson (1997; 1992 periastron data) by including the Paschen-delta and He I lines and allowing the time of periastron to vary, finding $e = 0.8$, $k = 65 \text{ km s}^{-1}$, and $\omega = 286^\circ$. Nielsen et al. (2007; 2003 periastron data) used He I P Cygni absorption and calculated $e = 0.9$, $k = 140 \text{ km s}^{-1}$, and $\omega = 270^\circ$. However, they note the absorption velocities may be influenced by changes in the ionization regions of He around periastron, leading to overestimates of k . Kashi & Soker (2016; 2009 & 2014 periastron data) made use of He I and N II lines finding $e = 0.9$, $k = 51 \text{ km s}^{-1}$, and $\omega = 90^\circ$ when assuming the lines form in the companion’s wind, in contradiction to the orientation found by most authors.

Our results show good agreement with the prevailing view of Eta Carinae’s orbital parameters, having a highly eccentric binary orientated with the companion on the far side of the primary at periastron with respect to our line of sight.

4.4.3.2 A comparison of periastron timings

The results from our CKM model fitting also provide an estimate for the timing of the 2009 periastron [start of cycle number 12 as defined by Groh & Damineli (2004)], resulting in a value of $T_0 = 2454848$ (JD).

Previously, Damineli et al. (2008a) derived $T_0 = 2454842.5 \pm 2$ d (JD) by using the disappearance of the narrow component of He I 6678 Å line. Teodoro et al. (2016) derived $T_0 = 2454851.7 \pm 1.3$ d (JD) by modelling the equivalent width of the He II 4686 Å line. Hamaguchi et al. (2014) observed the deep minimum in hard X-rays to occur between 2454843 and 2454859 (JD).

Our result lies between the values derived from the optical studies and in good agreement with the x-ray minimum. From the eclipse/collapse models of the x-ray light curve (see eg. Pittard et al. 1998; Parkin et al. 2009), we may expect the centre of the deep x-ray minimum, 2454851 (JD), to occur at superior conjunction. For our derived orbital elements, we expect periastron to occur ~ 3 d earlier than superior conjunction, which aligns well with our value of $T_0 = 2454848$ (JD).

4.4.3.3 Inferring the masses

The new value we have calculated for the semi-amplitude, $k = 69 \text{ km s}^{-1}$, allows us to estimate the stellar masses if we make assumptions about the system's inclination and Eddington fraction. The semi-amplitude, as shown previously in Equation 3.27, can be written in the form

$$k = \frac{29.78 \text{ km s}^{-1}}{\sqrt{1 - e^2}} \frac{M_B \sin i}{M_\odot} \left(\frac{M_A + M_B}{M_\odot} \right)^{-2/3} \left(\frac{P}{1 \text{ yr}} \right)^{-1/3}, \quad (4.11)$$

where M_A and M_B are the masses of the primary and companion, respectively, and i is the inclination.

First, only assuming values for i we are able to solve Equation 4.11 and plot contours of constant inclination on a diagram of mass ratio as shown in Figure 4.14 (orange lines). Modelling by Madura et al. (2012) of the broad [Fe III] emission had constrained the inclination to the interval $130^\circ < i < 145^\circ$, and we show this confidence interval as the shaded region in Figure 4.14.

To further constrain the masses requires an estimate for the total mass of the system. Following the work by Hillier et al. (2001), the Eddington luminosity for a




Figure 4.14: The masses of Eta Carinae based on $k = 69 \text{ km s}^{-1}$ for the primary's orbit. The orange lines represent contours of constant inclination. The grey shaded area is the confidence interval for the inclination taken from the Madura et al. (2012) study. The black line shows an Eddington contour for the case of the primary at the Eddington limit.

star composed of fully ionized H and He is given by

$$L_{\text{Edd}} = 3.22 \times 10^4 \left(\frac{N(\text{H})/N(\text{He}) + 4}{N(\text{H})/N(\text{He}) + 2} \right) \left(\frac{M}{M_{\odot}} \right) L_{\odot}, \quad (4.12)$$

where M is the Eddington mass of the system and $N(\text{H})/N(\text{He})$ is abundance ratio of hydrogen to helium. Using Equation 4.12, we estimate M for different fractions of the Eddington limit, assuming $N(\text{H})/N(\text{He}) = 10$ and $L = 5 \times 10^6 L_{\odot}$ (Davidson & Humphreys 1997). The Eddington mass may be attributed to either the primary, $M = M_A$, or the total mass of the system, $M = M_A + M_B$. For the former case, a contour is plotted in Figure 4.14 (black line) of the primary at the Eddington limit.

Examining Figure 4.14 for $i = 137.5^\circ$ and $L = L_{\text{Edd}}$, we obtain estimates for the stellar masses. For the case where $M = M_A$ we estimate $M_A = 133 M_{\odot}$ and $M_B = 93 M_{\odot}$, and for the case where $M = M_A + M_B$ we estimate $M_A = 67 M_{\odot}$ and $M_B = 66 M_{\odot}$.

4.4.4 Application to RMC 140b

We investigate the binary system RMC 140b as a secondary case study which is more representative of a larger population of emission-line stars, relative to the esoteric Eta Carinae. In particular, we use our CKM model to produce forward models for the WR star in RMC 140b to understand, in principle, the extent to which extended emission may affect the computed orbital solutions of stars with powerful winds.

The WR star in RMC 140b has literature estimates for its mass-loss rate $\dot{M} = 1.26 \times 10^{-5} \text{ M}_{\odot} \text{ yr}^{-1}$, terminal wind velocity of $v_{\infty} = 1300 \text{ km s}^{-1}$, and orbital parameters $P = 2.7586 \text{ d}$, $e = 0$, $k_a = 137 \text{ km s}^{-1}$, and $k_b = 266 \text{ km s}^{-1}$ (Shenar et al. 2019). We calculate line-formation kernels, in the same way as described in Section 4.3.2.1, for a few canonical lines used in extracting radial velocities: He II 4686 Å, N IV 4058 Å, and N V 4603 Å. The resulting line-formation kernels can be seen in the top panel of Figure 4.16 as a function of wind flow time in days (top axis) and the orbital phase of RMC 140b (bottom axis).

We convolve each line-formation kernel with the literature Keplerian solution to generate a set of CKM curves, plotted in the middle panel of Figure 4.16. The results show substantial deviations between the Keplerian motion and CKM, with the lines formed furthest from the star showing the largest changes. N V shows a ~ 5 per cent change in the time of periastron and a $\sim 3 \text{ km s}^{-1}$ decrease in semi-amplitude. The lower energy line, He II, shows a ~ 15 per cent change in the time of periastron and a $\sim 57 \text{ km s}^{-1}$ decrease in semi-amplitude.

The deviations between our predicted radial velocities for each line and the true Keplerian motion of RMC 140b have important ramifications for future higher accuracy work. Despite using the lines that are commonly adopted for extracting radial velocities we find variations between the resulting motion. This variability highlights the need to account for non-Keplerian effects and carefully consider which emission lines are used when extracting radial velocities to a high degree of accuracy. In the case of RMC 140b, these differences are being driven by the size of the line-formation kernels relative to the orbital period. The short period of $P = 2.7586 \text{ d}$ exacerbates the time-averaging effects of CKM. To test the sensitivity of the deviations on the period we create a set of CKM curves for a range of artificial period lengths between the original 2.7586d and 100d. The results are shown for the He II line in the bottom




Figure 4.15: Line-formation kernels for RMC 140b. Each kernel is the linear luminosity density (otherwise known as luminosity per unit radius) as a function of wind flow time beyond the sonic point. Flow time is shown in days (top x -axis) and in units of RMC 140b's orbital phase (bottom x -axis).




Figure 4.16: Left panel: RMC 140b's CKM deviations from the Keplerian orbit for each line. Right panel: deviations from the Keplerian orbit for variations in period for He II 4686 Å.

panel of Figure 4.16. The CKM curve for the original 2.7586d period is identical between the middle and bottom panels. However, as the period length increases, the deviations between the Keplerian motion and CKM curves incrementally decreases. At a period of ~ 100 d the differences become negligible.

In general, the expected size of the deviations between radial velocity curves and the true Keplerian motion depends on the interplay between the line-formation kernels and the orbit. The more extended the line emission region relative to the orbital time-scales, the larger the deviations. The emission region is more extended for higher mass-loss rate stars, and the wind flow times are longer for stars with slower winds. As a consequence, LBVs showcase the deviations to maximum effect and other types of massive evolved stars are good candidates for observing further cases of this effect.

4.4.5 Application to GG Carinae

We investigate the binary system GG Carinae as our third and final case study in this chapter. Spectroscopic observations of GG Carinae are from the Global Jet Watch, as described in Section 1.3. The dataset of GG Carinae spectra have been accumulating since early 2015, providing extensive phase coverage across many orbital periods. When GG Carinae’s optical emission-lines are fitted by Keplerian orbital models the estimated orbital parameters show systematic trends, correlated with the line energy, in much the same way as we have found for Eta Carinae (Porter et al. 2021a). Therefore, we implement our CKM model in an attempt to improve the orbital parameter determination.

GG Carinae has an estimated mass-loss rate $\dot{M} = 2.2 \times 10^{-6} M_{\odot} \text{yr}^{-1}$ and terminal wind velocity of $v_{\infty} = 265 \text{ km s}^{-1}$ estimated from high-resolution FEROS spectra. These values are in broad agreement with the original values calculated by McGregor et al. (1988). We calculate line-formation kernels, in the same way as described in Section 4.3.2.1, for a few optical lines used for extracting radial velocities: H-alpha, Si II 6347 Å, Si II 6371 Å, He I 5875 Å, He I 6678 Å, and He I 7065 Å. The resulting line-formation kernels for each species can be seen in the top panel of Figure 4.17 as a function of wind flow time in days (top axis) and the orbital phase of GG Carinae (bottom axis). We note that these line-formation kernels are based on the mass-loss rate for GG Carinae, which itself was estimated using the wind momentum–luminosity




Figure 4.17: Line-formation kernels for GG Carinae. Each kernel is the linear luminosity density (otherwise known as luminosity per unit radius) as a function of wind flow time beyond the sonic point. Flow time is shown in days (top x -axis) and in units of GG Carinae's orbital phase (bottom x -axis).




Figure 4.18: CKM best fitting model (red line) to the optical radial velocities (grey) of GG Carinae. Also shown is the inferred Keplerian motion (dark blue dashed line) and time-averaged radial velocities in bins with a width 10% of the orbital period.

relation for early B I-type supergiant stars (Kudritzki & Puls 2000). Additionally, the line-formation regions of these lines can be attributed to the wind of GG Carinae based on extensive modelling of the line profiles (Porter et al. 2021a), variability (Porter et al. 2021b), and dynamics (Porter et al. 2021c) of the system.

We co-fit our CKM model to all 6 of the spectral lines and present the resulting fit in Figure 4.18. Here, we show the Keplerian motion (dark blue dashed line) relative to the CKM fit (red line) for all the lines. Overall, the CKM model substantially improves the consistency of the fit to the radial velocities. Similarly to Eta Carinae there still exists some systematic deviations that vary depending on the line.

The derived orbital parameters for GG Carinae are $P = 31.01$ days, $T_0 = 2452069.4$ (JD), $e = 0.50$, $k = 48.6 \text{ km s}^{-1}$, and $\omega = 339.9^\circ$. These new parameters have important implications for understanding the photometric variability of GG Carinae. The higher eccentricity and shifted time of periastron, relative to earlier estimates by Marchiano et al. (2012), now lead to a natural explanation of the photometric maximum occurring at periastron due to enhanced mass transfer between the stars (Porter et al. 2021a).

4.5 Conclusions

In this chapter we have investigated discrepancies in radial velocity data between spectral lines of varying excitation emitted from binary stars with powerful winds. We have presented a new semi-analytical model, CKM, which encapsulates both the star’s orbital motion and the propagation of the wind, accounting for the time delay to the location where the photons are emitted. Our work is summarised as follows:

- (i) We made use of archival Gemini/GMOS observations of Eta Carinae from across the 2009 periastron. We employed a multi-Gaussian fitting algorithm to decompose the complex line profiles of the Balmer series into their constituent dynamical components and extracted radial velocities from the primary star’s wind emission.
- (ii) As an instructive exercise we fitted the radial velocity observations with Keplerian models, finding monotonic trends in the orbital parameters T_0 , e , and

k when ordering the lines by their transition energy: different lines produce different orbital solutions.

- (iii) To reconcile the discrepancies in the orbital parameters between lines we formulated a semi-analytical model, CKM, which encapsulates the dynamics associated with lines formed in extended regions, encoding the integrated velocity field of the out-flowing gas in terms of a convolution of past motion due to the finite flow speed of the wind.
- (iv) We fitted our CKM model to 10 of Eta Carinae’s Balmer lines, from H-alpha to H-kappa, concurrently. The best-fitting orbital parameters are $T_0 = 2454848$ (JD), $e = 0.91$, $k = 69 \text{ km s}^{-1}$, and $\omega = 241^\circ$. Cross-validation showed stability in these parameters for the higher excitation Balmer lines, but we found systematic effects, most likely due to the influence of the companion, are apparent in H-alpha and H-beta. Overall, this model is a first step towards being able to generate a consistent set of Keplerian parameters independent of the emission line used, and thereby expands the set of lines that can be used for orbital parameter inference.
- (v) We applied our CKM model to a couple of more typical cases: the WR star in the binary RMC 140b, and the B[e] supergiant in the binary GG Carinae. Despite selecting the canonical lines used for extracting radial velocities (He II and N IV/V), our forward model of RMC 140b found deviations may still exist between the published Keplerian orbital parameters and the predicted radial velocities. This is an important correction for future higher accuracy work. We found these deviations are driven by the short orbital period in RMC 140b. For GG Carinae we analysed 6 optical emission-lines and found the lines displayed significant deviations from Keplerian motion. We applied our CKM model simultaneously to all 6 of these lines and estimated improved orbital parameters of $P = 31.01$ days, $T_0 = 2452069.4$ (JD), $e = 0.50$, $k = 48.6 \text{ km s}^{-1}$, and $\omega = 339.9^\circ$.
- (vi) More generally, the magnitude of the deviations between radial velocity curves and the true Keplerian motion depends on the interplay between the line-

formation kernels and the orbit. The more extended the line emission regions relative to the orbital time-scales, the larger the deviations.

Chapter 5

Probabilistic orbits and dynamical masses of emission-line binaries using marginalised Gaussian processes

5.1 Introduction

Accurate inferences of the orbits and masses of multiple-star systems are required in order to fully understand the lives and deaths of stars. For binary systems, well-sampled time-series spectroscopic measurements enable the line-of-sight projected orbital motion to be calculated. When these radial velocity measurements are combined with either interferometric imaging – or light curves in the case of eclipsing binaries – the full three-dimensional orbit can be constrained (eg. Thomas et al. 2021; Richardson et al. 2021). As a result, accurate dynamical masses of the individual stellar components can be computed. However, these results are predicated on the assumption that the radial velocity data well-represent the orbital motion of the system. Whilst this is a fair assumption for many stars, where the spectral lines are formed unencumbered at the stellar surface, for emission-line stars, as we have seen in Chapter 4, the spectral lines may encode velocities that are not purely orbital in their origin. Consequently, radial velocity data for emission-line binaries may often be afflicted by systematic errors with implications for the inferred orbits and dynamical masses.

The sources of systematics are numerous. The broad emission lines are formed in a complex environment of outflowing gas and often in the presence of colliding-wind

excess emission, both of which can alter the line profiles and bias the measurements. Additionally, methods to extract the radial velocities may be imperfect at disentangling the portion of the line profiles that represents solely the Keplerian motion.

The current standard practice for emission-line binaries is to assume radial velocities are ideally extracted and exhibit Keplerian motion plus some Gaussian white noise. Orbital models are then fitted to the data and the parameters of the orbit are estimated. Any poorly fit data are often noted qualitatively and thereafter ignored as the orbital parameters are carried through to the subsequent analysis of the stellar masses and evolutionary state of the systems. Put bluntly: if the orbital models do not explain all of the signals in the data then we are at risk of inferring biased results.

In this chapter we want to create a sea change in the methodology when estimating the orbital parameters of emission-line binaries. In Section 5.2 we synthesise benchmark datasets, with various noise formulations, for testing the robustness of orbital models in single-lined spectroscopic binaries. We make these datasets freely available with the aim of making model validation an easy and standardised practice. In Section 5.3 we describe the current algorithms used within the field, and introduce a new approach to modelling the radial velocities of emission-line binaries using Gaussian processes (GPs). We apply the algorithms to the benchmark datasets and find the GP algorithm serves as the current best method. We encourage other researchers to employ novel techniques to improve upon our scores with these same synthetic datasets. In Section 5.5 we discuss how the results can be affected by having different data available, such as in the case of a double-lined spectroscopic binary or by having access to relative astrometric measurements. We then apply our GP algorithm to four real systems – Eta Carinae, GG Carinae, WR 140, and WR 133 – to assess the impact of our new approach on the estimated orbital parameters and dynamical masses. Finally, in Section 5.6 we summarise our findings.

5.2 A benchmark dataset

We conceive of a data analysis work flow in which the investigator – analysing radial velocity data in our use case – first tests their models on a benchmark dataset, before fitting their models to their observations. In this way the investigator can ensure their model is robust and performant, as well as providing a quantitative measure by which

the best modelling techniques can take precedence in the field. This methodology is inspired by other fields, such as machine learning, in which advancements have long been validated using benchmark datasets.

To this end, we make a first attempt at creating benchmark datasets for the orbits of emission-line binaries. We restrict the data to single-lined spectroscopic binaries (SB1), however we discuss more complex data combinations in Section 5.5. We generate three datasets each having the same number of synthetic binaries, the same distribution of orbital parameters, but with different noise formulations.

The orbital parameters required to fully specify SB1 data are the time of periastron, T_0 , the orbital period, P , the eccentricity, e , the argument of periastron, ω , the primary star’s semi-amplitude, k_1 , and the primary star’s radial velocity offset, γ_1 . We fix $T_0 = 2458932$ (JD), $P = 50$ days, and $\gamma_1 = 0 \text{ km s}^{-1}$ as these parameters only serve to re-scale or translate the data in the time or velocity dimensions. We select 100 observations from one orbital period for every system. The observations are drawn from a beta distribution with shape parameters $\alpha = \beta = 1 - e$ with values corresponding to the orbital phase. The result is a realistic imitation of observing patterns conducted on binary stars, in which observations around periastron become more important as the eccentricity increases, and there is always sufficient orbital coverage to infer all six of the orbital parameters.

The noise added to each dataset is designed to challenge potential models to be robust against the various problems that may be embedded in the radial velocities extracted from emission-line stars. In dataset 1 we add only Gaussian white noise to the velocities. This is the easiest dataset for models to contend with, resembling the case of ideally extracted observations encoding purely orbital motion. The white noise level is fixed to 5 km s^{-1} and so the value given to k_1 serves to adjust the signal-to-noise level. In dataset 2 we add correlated noise to the velocities by injecting randomly generated polynomials. The polynomials are synthesised from a random set of roots that lie within the observational time frame. For each system, the polynomials are adjusted in amplitude by scaling the maximum deviation by a sample from the uniform distribution $0 - 0.5k_1$, and adjusted in their time-variability by sampling an integer number of roots uniformly over the range 1-10. In dataset 3 we again add correlated noise to the velocities, but in this case we inject functions randomly

drawn from a Gaussian process prior. For each system, we tune the amplitude and time-variability of the interloping signal through the squared exponential covariance function (see Equation 5.9). We use this kernel as it produces smoothly-varying functions in time. For each system we sample an amplitude scale factor from the uniform distribution $0 - 0.5k_1$, and adjust the time-variability by sampling the characteristic length scale from the uniform distribution $0.1 - 30$ days. A more thorough definition of a Gaussian process is treated in Section 5.3.

The synthesised systems each have a total of five parameters that are varied across the orbital model and noise formulation. In order to balance the time required to benchmark a model across the datasets with good parameter coverage, we generate $N = 3^5 = 243$ systems for each of the three datasets¹. For each system the data included are the radial velocity measurements, $\mathbf{y} = (y_1, \dots, y_N)^T$, each at a given time and spectral line energy, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T = ((t_1, E_1)^T, \dots, (t_N, E_N)^T)^T$, and with associated uncertainties, measuring the white noise only, $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_N)^T$, as well as the true orbital parameters.

The benchmark datasets are open-source and freely available². The data are provided in a number of easily usable formats and we encourage researchers to validate their current models with the data, and to beat the scores we set in this study with new and ingenious models.

5.2.1 Evaluation criterion

In order to quantify the performance of different models on each benchmark dataset we define an evaluation criterion. The criterion will assess how well the inferences made about the orbital parameters compare to the injected true values. Following similar motivation as Quiñonero-Candela et al. (2005), the desired behaviour of the criterion is such that for two models of equal accuracy but with different uncertainty estimates, the more confident model will score higher (lower) when the prediction is more accurate (inaccurate). To achieve this we compute the density of the predicted posterior distribution at the true value. For each of the assessable parameters this

¹Our models run for approximately 17 minutes per synthetic object on one 2 GHz Intel core i5 processor, or eg. 3.4 hours for an entire dataset on 20 cores assuming ideal parallelisation.

²<https://github.com/DavoGrant/benchmark-datasets-emission-line-binaries>

can be written as

$$p(\mathbf{w} = \mathbf{w}_t \mid \mathbf{y}, \mathbf{X}) = \frac{p(\mathbf{y} \mid \mathbf{X}, \mathbf{w}_t)p(\mathbf{w}_t)}{p(\mathbf{y} \mid \mathbf{X})}, \quad (5.1)$$

where $\mathbf{w}$ are the parameters of the model and $\mathbf{w}_t$ is the injected true parameter vector we wish to recover. The denominator, known as the marginal likelihood, is given by

$$p(\mathbf{y} \mid \mathbf{X}) = \int p(\mathbf{y} \mid \mathbf{X}, \mathbf{w})p(\mathbf{w})d\mathbf{w}. \quad (5.2)$$

The overall score per model per dataset is

$$\text{score} = \text{median}(S), \quad (5.3)$$

where the set, S , is

$$S = \{\ln p(\mathbf{w} = \mathbf{w}_{t,i} \mid \mathbf{y}_i, \mathbf{X}_i) \mid i = 1, 2, \dots, N\}, \quad (5.4)$$

and N is the number of synthetic objects in the dataset. This score is the mean logarithmic density of the predicted posterior distribution at the true value for each of the parameters. Note that the score can be negative and will depend on the scale of each parameter. In practice, sampling methods will be used to estimate the posterior distribution, owing to the difficulty in computing the marginal likelihood integral. We therefore employ kernel density estimation (Scott 2015), with Silverman’s rule of thumb to select the bandwidth (Silverman 2018), to evaluate the score in a non-parametric way.

5.3 A benchmark algorithm

In this section we introduce benchmark algorithms for analysing the orbits of emission-line binaries through Bayesian parameter estimation. We briefly re-state two common procedures used in the literature, as well as a generative algorithm based on Gaussian processes to account for systematic noise. These algorithms can be summarised as follows:

- (i) χ^2 , this algorithm is the most simple in its procedure. It is based on the assumption that the radial velocity data only encode the orbital motion plus some Gaussian white noise. In this case the objective function is proportional to the χ^2 value.

- (ii) Marginalised σ , this algorithm is a simple extension to the first algorithm. It is used when the uncertainties output from radial velocity extraction techniques do not appropriately represent the variance in the data about the orbital motion. An additional free parameter is added in quadrature to the extracted uncertainties in the data to account for the initially poorly estimated uncertainties. This term is sometimes referred to as a jitter in the literature and can be both statistical or systematic in origin.
- (iii) Marginalised $\mathcal{GP}$, this algorithm is comprised of a mean model plus a GP, to model the orbital motion plus any systematics, respectively. The hyperparameters of the GP are marginalised over to best capture the predictive uncertainty in the orbital parameters.

The third benchmark algorithm described above, marginalised $\mathcal{GP}$, takes a generative approach to dealing with the noise present in the radial velocity data. A GP is defined as a collection of random variables, any finite number of which have a joint Gaussian distribution (Rasmussen & Williams 2006). This non-parametric approach to regression lends itself to modelling the unknown noise component in astrophysics datasets. GPs have been used previously, for example, by Brewer & Stello (2009) for stellar oscillations in light curves, by Gibson et al. (2012) for transmission spectroscopy, by Aksulu et al. (2020) for gamma-ray burst afterglows, and by Aigrain et al. (2016) for exoplanet light curves. Here we describe a novel application of GPs to the radial velocity systematics of emission-line binaries.

5.3.1 Mean models

At the heart of our marginalised $\mathcal{GP}$ algorithm we require a parameterised model of the orbiting emission-line stars. For radial velocities extracted from lines that form deep in the wind, the velocities can be described by Keplerian motion. However, in Chapter 4 we found that for lines emitted over extended radii, the observed motion is better characterised by a convolution of past motion accounting for the emission region of a particular spectral line. Consequently, we have two possible orbital models. For lines which form deep in the wind we have the Aitken (1964) standard expression

for radial velocities

$$v_{\text{kep}}(t) = k_1(\cos(\omega + \nu(t)) + e \cos \omega) + \gamma_1, \quad (5.5)$$

or, for lines forming over extended regions in the wind we have the modified formalism that we developed in Chapter 4,

$$v_{\text{ckm},n}(t) = v_{\text{kep}}(t) * \lambda_{L,n}(t) = \int_{-\infty}^{\infty} v_{\text{kep}}(\tau) \lambda_{L,n}(t - \tau) d\tau, \quad (5.6)$$

where in these equations ν is the true anomaly, v_{kep} is the Keplerian velocity projected onto our line-of-sight, and $v_{\text{ckm},n}$ is the convolutional Keplerian motion which depends on the spectral line, n , via the line-formation kernel, $\lambda_{L,n}(t)$.

For the generation and analysis of the benchmark datasets described in Section 5.2 we exclusively use Equation 5.5 as the orbital model, which has the parameter vector

$$\boldsymbol{\theta} = (T_0, P, \omega, k_1, \gamma_1)^T. \quad (5.7)$$

For some observational datasets analysed in Section 5.5.2 we also make use of Equation 5.6. These models are the basis for all three benchmark algorithms, but specifically for the GP algorithm they are referred to as the mean model.

5.3.2 Gaussian processes

We wrap the mean model in a GP framework to model additional systematics; this could be noise or other non-orbital astrophysical signals. More formally this is written

$$f(\mathbf{X}) \sim \mathcal{GP}(v(\mathbf{X}, \boldsymbol{\theta}), k(\mathbf{X}, \boldsymbol{\phi})), \quad (5.8)$$

where v is the mean model, k is the covariance function, $\boldsymbol{\theta}$ is the parameter vector of the mean model, $\boldsymbol{\phi}$ is the hyperparameter vector of the covariance function, and $\mathbf{X}$ contains the times and spectral line energies of the input data as described in Section 5.2.

Through specifying the covariance function we are encoding our assumptions about the nature of the systematics present in our data. We select the squared exponential kernel with the aim of capturing smoothly-varying deviations from orbital motion. For two input points, the covariance is described as

$$k(\mathbf{X}_p, \mathbf{X}_q) = A \exp \left[-\frac{1}{2} \sum_{r=1}^2 \frac{(\mathbf{X}_{pr} - \mathbf{X}_{qr})^2}{l_r^2} \right] + \delta_{pq} \sigma_w^2, \quad (5.9)$$

where A controls the amplitude, l_r controls the characteristic length scale of the correlations in each input dimension, and σ_w represents the white noise. From this expression it is clear that the closer input values are together, the more correlated the target values are expected to be, plus a certain level of white noise. The white noise parameter is further broken down into $\sigma_w^2 = \sigma_a^2 + \sigma_b^2$ where σ_a^2 is the extracted variance of the white noise and σ_b^2 is a free parameter which helps to account for underestimated values of σ_a^2 .

We formulate the covariance in two dimensions to capture systematics that are correlated in both time and spectral line energy. This follows from previous work on the dynamics of emission-line stars that we carried out in Chapter 4 and work by Porter et al. (2021a) in which the deviations from orbital motion were found to have time-variability and a dependence on the line transition from which the radial velocity data was extracted. In particular, the deviations are strongly correlated with the top level energy of the line transition. Physically this correlation may arise from the fact that similar energy lines originate from similar emission volumes in a system. The resulting hyperparameter vector for the covariance function is

$$\phi = (A, l_1, l_2, \sigma_b)^T. \quad (5.10)$$

In Figure 5.1 we display a visual representation of the GP framework described above. The figure is split into three columns with the full two-dimensional representation on the top row, and two one-dimensional slices in energy (15 eV and 25 eV) on the bottom row. The left-hand column shows a sample taken from the GP prior. The middle column shows a sample taken from the GP conditioned on a small amount of input data with no uncertainty. The right-hand column shows the same GP as the middle column but with the mean model – the orbital velocities – added back in. These plots were made with the hyperparameters fixed at $A = 11^2 \text{ km s}^{-1}$, $l_1 = 5^2 \text{ (days)}^2$, $l_2 = 8^2 \text{ (eV)}^2$, and $\sigma_w = 0 \text{ km s}^{-1} \text{ days}$. The plot shows how the squared exponential kernel specifies a prior over smoothly-varying functions, which after being given data, can learn the systematics present in a given system. The hyperparameters specify the nature of these systematics. For example, larger values of A allow the GP to learn larger deviations in velocity from the mean model, and smaller values of l_1 or l_2 result in the GP learning more rapidly changing deviations in time and energy




Figure 5.1: A visual representation of the GP algorithm. The top row displays the model spanning the full two-dimensional grid, whilst the bottom row shows two slices at energies 15 eV (purple line) and 25 eV (green line). The left-hand column shows a sample taken from the GP prior. The middle column shows a sample taken from the GP conditioned on a small amount of input data with no uncertainty. The right-hand column shows the same GP as in the middle column but with the mean model added back in. The GP hyperparameters were fixed at $A = 11^2 \text{ km s}^{-1}$, $l_1 = 5^2 \text{ days}$, $l_2 = 8^2 \text{ eV}$, and $\sigma_w = 0 \text{ km s}^{-1} \text{ days}$.

space, respectively. Note that this is also how we generated benchmark dataset 3, previously described in Section 5.2, through randomly drawing different hyperparameter vectors to add a variety of different systematic problems into the orbital models. For the benchmark datasets the second dimension of the GP is actually redundant as, for simplicity, we focus on SB1 binaries with radial velocity data from either one line only, or lines combined together into one input. However, in Section 5.5 we utilise the energy dimension and explore correlations between separate radial velocity lines.

5.3.3 Bayesian regression

In order to perform Bayesian parameter estimation of the orbital parameters of interest, θ , we first define an objective function. The log-likelihood of the GP model

is

$$\ln p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}, \phi) = -\frac{1}{2} \mathbf{r}^T \mathbf{K}^{-1} \mathbf{r} - \frac{1}{2} \ln |\mathbf{K}| - \frac{m}{2} \ln(2\pi), \quad (5.11)$$

where $\mathbf{r} = \mathbf{y} - \mathbf{v}$ is the residual vector after the data has had the mean model subtracted, and m is the number of measurements. In this expression the first term accounts for the goodness-of-fit: as the residuals become smaller the value increases for a given covariance matrix. The second term acts as a complexity penalty. For perfect correlation (ie. no GP complexity) the determinant equals zero. As the correlation decreases, by shortening the characteristic length scales and/or increasing the amplitude, the determinant increases, the whole term decreases in value, and the extra complexity of the GP becomes less likely. The additional complexity can, however, be justified by a corresponding increase in the goodness-of-fit. Consequently, the GP will only have an impact on the model when the likelihood of the data having been generated solely by the mean orbital model is sufficiently low. In other words, Equation 5.11 is no more than a Gaussian probability density function and the normalisation factor $|\mathbf{K}|^{1/2}$ naturally acts as a complexity penalty.

Following Simpson et al. (2020), we marginalise over both the hyperparameters of the covariance function as well as the free parameters in the mean model in order to best capture the predictive uncertainty in the parameters. We choose uniform priors over $\boldsymbol{\theta}$ and log-uniform priors over ϕ . We use the Gaussian process regression code, `george` (Ambikasaran et al. 2015), and the affine-invariant Markov Chain Monte Carlo (MCMC) algorithm, `emcee`, implemented by Foreman-Mackey et al. (2013) to sample the probability distributions. We estimate the integrated autocorrelation time to ensure we draw sufficient samples for good convergence and to reduce the resulting errors in the posteriors down to the couple-of-percent level. For the other two benchmark algorithms, χ^2 and marginalised σ , the regression is implemented in the exact same way as the GP model except the covariance matrix is always diagonal. The log-likelihood in Equation 5.11 reduces to

$$\ln p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}, \sigma_b) = -\frac{1}{2} \sum_n^m \left[\frac{(y_n - v(\boldsymbol{\theta})_n)^2}{\sigma_w^2} + \ln(2\pi\sigma_w^2) \right], \quad (5.12)$$

where for the χ^2 algorithm the white noise is fixed such that $\sigma_w^2 = \sigma_a^2$, but for the marginalised σ algorithm σ_b is a free parameter and $\sigma_w^2 = \sigma_a^2 + \sigma_b^2$.

Algorithm	Dataset 1: white noise					
	T_0	P	e	ω	k_1	γ_1
χ^2	0.96	-1.30	3.55	2.30	-1.03	-1.35
marginalised σ	0.85	-1.37	3.53	2.28	-1.05	-1.36
marginalised $\mathcal{GP}$	0.83	-1.41	3.51	2.24	-1.12	-1.45

Table 5.1: Results for dataset 1 for each of the three benchmark algorithms. Bold indicates the best scoring algorithm per parameter. The score is defined in Section 5.2.1.

5.4 Setting the standard

In Section 5.2 we introduced three SB1 benchmark datasets, each with a different noise formulation. In Section 5.3 we described two commonly-employed algorithms, χ^2 and marginalised σ , as well as proposing a novel application of a GP based algorithm, marginalised $\mathcal{GP}$. In this section we apply all three algorithms to all three benchmark datasets and report the results, setting the initial best scores to beat. The scoring is based on the criterion described in Section 5.2.1, evaluating the predictive performance of the posterior parameter distribution when compared with the injected true values.

In Tables 5.1, 5.2, and 5.3 we present the results for benchmark datasets 1, 2, and 3, respectively. We find that our marginalised $\mathcal{GP}$ algorithm substantially outperforms the other more traditional algorithms in both datasets 2 and 3. The performance is comparable between these two datasets’ scores indicating that they are set at a similar difficulty level. For dataset 1, all three algorithms perform very similarly, although the χ^2 algorithm narrowly records the best scores. These scores demonstrate how both the marginalised σ and marginalised $\mathcal{GP}$ algorithms (almost) reduce to the χ^2 algorithm in the case of purely white noise.

In Figures 5.2 and 5.3 we display two illustrative examples from the benchmark results. In both Figures the left-hand column shows a system from dataset 1, containing white noise only, whilst the right-hand column shows the corresponding system from dataset 3, now contaminated by systematic noise. In Figure 5.2 the systematics are varying on roughly the same timescale as the orbital period and in Figure 5.3

Algorithm	Dataset 2: polynomial systematics					
	T_0	P	e	ω	k_1	γ_1
χ^2	-1.84	-19.1	-0.77	-3.34	-6.62	-25.1
marginalised σ	-1.49	-6.94	0.59	-1.32	-3.97	-6.64
marginalised $\mathcal{GP}$	0.13	-3.42	2.09	0.89	-3.00	-3.83

Table 5.2: The same as Table 5.1 but for dataset 2.

Algorithm	Dataset 3: $\mathcal{GP}$ systematics					
	T_0	P	e	ω	k_1	γ_1
χ^2	-3.99	-76.7	-0.45	-5.81	-486	-3130
marginalised σ	-1.89	-6.63	1.16	-1.92	-8.48	-24.9
marginalised $\mathcal{GP}$	0.01	-3.71	2.04	0.96	-3.57	-4.49

Table 5.3: The same as Table 5.1 but for dataset 3.

they vary much faster with multiple mean model crossings per orbital period. In the top panels we plot the marginalised $\mathcal{GP}$ models along with their 2σ uncertainties. In the bottom panels we show violin plots of the posterior distributions for three of the orbital parameters, e , ω , and k_1 , relative to the true values (dashed lines). In the regime of purely white noise we see how all three models accurately and confidently predict the true parameter values. The marginalised $\mathcal{GP}$ algorithm does have slightly longer tails in its predictive distributions, which manifests itself as an ever so slightly lower score in Table 5.1. In the regime of systematic noise we can immediately see how the marginalised $\mathcal{GP}$ algorithm predicts more complex and uncertain posterior distributions in each of the orbital parameters. Consequently, this algorithm is far more accurate in recovering the true value, whilst the other two algorithms are over-confident given their inaccuracy.

The GP algorithm outperforms the other algorithms when systematic noise is present in the data. The χ^2 and marginalised σ algorithms return the orbit that minimises their velocity residuals, which is not necessarily representative of the true underlying orbit if the systematics are asymmetrically distributed about the mean model. Conversely, the GP algorithm can generate flexible systematic functions, so




Figure 5.2: SB1 benchmark dataset example object 32/243. The left-hand column presents the object’s radial velocity data and model fits from benchmark dataset 1, containing white noise only. The right-hand column presents the corresponding object from dataset 3, now containing systematic noise on a similar timescale as the orbital period. In the top panels the dashed lines represent the mean models, while the solid purple represents the GP models and their 2σ uncertainties. In the bottom panels the dashed lines indicate the true values for each parameter, alongside their posterior parameter distributions for each of the three benchmark algorithms.




Figure 5.3: SB1 benchmark dataset example object 86/243. The same as Figure 5.2 except the object’s systematics show variability on a much faster timescale.

long as the systematics can be disentangled from the mean model. In general this is the case, as is shown in Figures 5.2 and 5.3, although the GP model will still of course struggle if the mean model plus the systematics result in radial velocities that can be confused with a different mean model parameterisation. But for the most part, the combinations of various orbital models and GP systematics are marginalised over and the resulting parameter posterior distributions incorporate this uncertainty, making the true values more likely to be recovered by this algorithm.

5.5 Discussion

We have presented benchmark datasets and algorithms for the orbits of emission-line binaries. We have found that when systematics are present in the data, simply inflating the uncertainties does not always modify the predicted orbital parameter distributions enough to recover the true values. Therefore, we find utilising our marginalised $\mathcal{GP}$ algorithm is important for predicting accurate orbital parameter distributions. This has important implications for the derived orbits of emission-line stars, their dynamical masses, and the binary evolutionary scenarios that may be inferred from radial velocity data.

The generative nature of the GP algorithm has further advantages, such that we can separate out orbital motion from other potential physical signals lurking in the data. The GP model can then be used to help identify the sources of systematics and further characterise different aspects of the system, rather than sweeping these signals under the carpet by assuming the uncertainties in the data are underestimated.

5.5.1 The performance impact of additional data

In this section we analyse how having additional data – such as radial velocities from different spectral lines or stellar components, or having relative astrometric measurements – can impact the predictive distributions of our marginalised $\mathcal{GP}$ algorithm. To facilitate the discussion we set a baseline score for all the tests using an SB1 object with orbital parameters $P = 50$ days, $T_0 = 2458932$ (JD), $e = 0.3$, $\omega = 270^\circ$, $k_1 = 50 \text{ km s}^{-1}$, and $\gamma_1 = 0 \text{ km s}^{-1}$. The primary star’s radial velocities also include systematics drawn from a GP with a squared exponential kernel, parameterised by $A = 20^2 \text{ km s}^{-1}$ and $l_1 = 2^2 \text{ (days)}^2$.

5.5.1.1 Multiple spectroscopic lines

Often radial velocities are extracted from many spectral lines and combined, or extracted from entire portions of the spectra, resulting in one radial velocity curve for a given binary system. However, it may be the case that different spectral lines display different systematic behaviour; and therefore, it is preferable to treat the systematics from each line individually. To test the effect of utilising multiple spectral lines on our marginalised $\mathcal{GP}$ algorithm we create a test dataset with 10 spectral lines, each having identical underlying orbital parameters according to the baseline object. Each line also has a randomly generated GP function, drawn from the prior distribution of a GP parameterised as per the baseline object.

We run the marginalised $\mathcal{GP}$ algorithm 10 times, starting with one line and incrementally including an additional line each run. In this case we utilise the second dimension of the GP, allowing the model to learn systematic correlations in the energy dimension. We find that the more spectral lines are included in the parameter estimation, the better the score becomes. Once all 10 lines are utilised the score increases by 0.65 on average across each of the 6 orbital parameters.

5.5.1.2 Double-lined spectroscopic binaries

Spectra of emission-line binaries may include lines from both stellar components, known as a double-lined spectroscopic binary (SB2). In this case it is possible to co-fit both radial velocity curves with the same orbital model as most of the parameters are shared. To test the effect of having SB2 data on our marginalised $\mathcal{GP}$ algorithm we create a test dataset of 10 objects. All of the objects have identical underlying orbital parameters as in the baseline object. The primary star’s radial velocities include systematics as in the baseline object. The secondary’s systematics are then varied between the 10 objects, with the timescale fixed at $l_1 = 2^2$ (days)² and the amplitude linearly increasing over the interval $0 \text{ km s}^{-1} < A < 20^2 \text{ km s}^{-1}$. For $A = 0 \text{ km s}^{-1}$ this simulates the scenario when the secondary is, for example, an O-star with easy to extract absorption lines, and therefore little to no systematics in the radial velocities. For $A = 20^2 \text{ km s}^{-1}$ this simulates having another emission-line star where substantial systematics may remain in the data.

We run the marginalised $\mathcal{GP}$ algorithm over these 10 objects. We find that for small or no systematics in the secondary’s radial velocities the score is improved on average by 1.64 across the 6 orbital parameters. As the amplitude of the secondary’s systematics are increased, the score decreases. Once the amplitude of the secondary’s systematics are comparable to the primary’s, the average score has returned to the same as the SB1 case. Overall, we find that the presence of the secondary’s data helps the algorithm to distinguish between the mean orbital model and the systematics, especially when the secondary itself is reasonably unaffected by systematics.

5.5.1.3 Relative astrometry

For some (close by) systems it is possible to observe the relative motion of the binary stars through precise relative astrometry from interferometric images. Through combining the radial velocities and relative astrometric data we can constrain the entire 3D geometry of the orbits. To test the effect of having relative astrometry in addition to SB1 data we create a test dataset of 10 objects, all of which have the same underlying orbital parameters and systematics as the baseline object. In each object we vary the number of astrometric observations available from 2 to 100, where the epochs of the observations are sampled from a uniform distribution over the same time interval as the radial velocities.

We run the marginalised $\mathcal{GP}$ algorithm over these 10 objects. We find even having as little as two astrometric observations helps to improve the average score of the orbital parameter inference by 0.52. The scores continue to improve as the number of astrometric observations increases. Once we have 100 astrometric observations the average score is improved by 1.97. The parameters P , T_0 , e , and ω get the largest boost in performance as they appear directly in the astrometric orbital model, but interestingly the parameters k_1 and γ_1 also receive a modest improvement, probably as an indirect consequence of the other parameter distributions becoming more accurate.

5.5.2 Application to prototypical systems

In this section we apply our best benchmark algorithm, the marginalised $\mathcal{GP}$, to some real datasets of emission-line binaries to compare how our predicted parameter

Eta Carinae's orbit	Chapter 4	marginalised $\mathcal{GP}$
Elements of the system		
P (fixed)	2022.7	2022.7
T_0	$2454848.3^{+0.4}_{-0.4}$	$2454850.1^{+2.3}_{-2.6}$
e	$0.91^{+0.003}_{-0.003}$	$0.90^{+0.01}_{-0.01}$
ω	241^{+1}_{-1}	248^{+4}_{-5}
Elements of the radial velocity orbit		
k_1	$69.0^{+0.9}_{-0.9}$	$66.8^{+4.6}_{-3.9}$
γ_1 (fixed)	-	-

Table 5.4: Results for Eta Carinae comparing literature parameter distributions, using the χ^2 algorithm, and our estimated parameter distributions, using our marginalised $\mathcal{GP}$ algorithm. The uncertainties are the 68% credible intervals. The orbital parameters are the period, P (days), the time of periastron, T_0 (JD), the eccentricity, e , the argument of periastron, ω (degrees), the primary star's semi-amplitude, k_1 (km s^{-1}), and the primary star's radial velocity offset, γ_1 (km s^{-1}).

distributions differ from those previously published and to elucidate systematics that may otherwise go untraced.

5.5.2.1 Eta Carinae

Eta Carinae is a long-period (2022.7 days (Damineli et al. 2008a)), highly eccentric ($e \sim 0.9$) SB1 system situated at a distance of 2.3 ± 0.1 kpc (Allen & Hillier 1993; Smith 2006). The primary star is a luminous blue variable with an exceptionally strong wind; it has a mass-loss rate of $8.5 \times 10^{-4} \text{M}_{\odot} \text{yr}^{-1}$ and terminal wind velocity of 420 km s^{-1} (Hillier et al. 2001; Groh et al. 2012a; Clementel et al. 2015a). The orbital parameters were originally calculated by Damineli et al. (1997) and in Chapter 4 we re-computed the orbital parameters using the Balmer lines and the mean model described in Equation 5.6.

We utilise the same spectral data as in Chapter 4 which is from an open-source online archive for Eta Carinae³. The dataset contains between 57 and 132 radial velocity measurements per Balmer line. The Balmer lines included span H-alpha to

³<http://etacar.umn.edu/archive/>




Figure 5.4: Marginalised $\mathcal{GP}$ algorithm conditioned on Eta Carinae's observed data. The mean model is shown as the black dashed line, the GP is shown at sequential spectral line energies by the coloured lines from H-alpha (purple) to H-kappa (yellow), and each shaded region represents the 2σ uncertainty in the GP. The γ values are arbitrary.

H-kappa with corresponding top-level line-transition energies from 12.09 eV to 13.50 eV.

We first apply the χ^2 algorithm to the data and check that we recover predicted parameter distributions that are consistent with those obtained in Chapter 4. We note that similarly to the fitting in Chapter 4 we fix the orbital period and pre-set the γ values for each spectral line. Next, we apply our marginalised $\mathcal{GP}$ algorithm to the same data and report the results in Table 5.4. We find that the distributions of all the parameters become significantly more uncertain when using our marginalised $\mathcal{GP}$ algorithm.

In Figure 5.4 we display the conditioned GP for Eta Carinae. In this fit we can see the data exhibits obvious systematics, the most prominent of which occur in H-alpha. Despite this the GP is able to fully capture the variability of the data about the mean model. The median characteristic length scale of the marginalised GP kernel is 10.4 days which manifests itself as systematics varying on timescales much faster than the orbital period.

The overall shift to lower values of k_1 , for our predicted distribution relative to the literature values, results in the predicted mass ratio, $q = m_2/m_1$, moving to lower values. Therefore the extent to which the primary star’s mass is greater than the secondary’s may be larger than previously inferred.

5.5.2.2 GG Carinae

GG Carinae is an SB1 system with a ~ 31 day period (Hernandez et al. 1981; Gosset et al. 1985; Marchiano et al. 2012) at a distance of $3.4^{+0.7}_{-0.5}$ kpc (Gaia Collaboration et al. 2016; 2018). The primary star is a B[e] supergiant with a mass-loss rate of $2.2 \times 10^{-6} M_{\odot} \text{yr}^{-1}$ and terminal wind velocity of 265 km s^{-1} (Porter et al. 2021a). We use the orbital parameters calculated by Porter et al. (2021a; Section 4.4.5 of this thesis) using the same modelling techniques as used for Eta Carinae, making use of the mean model described in Equation 5.6.

We utilise the same spectral data as Porter et al. (2021a) which is from the Global Jet Watch network of telescopes (see Section 1.3). The dataset contains between 513 and 718 radial velocity measurements per spectral line. The lines included are H-alpha, Si II 6347 Å, Si II 6371 Å, He I 5875 Å, He I 6678 Å, and He I 7065 Å with

GG Carinae's orbit	Porter et al. (2021a)	marginalised $\mathcal{GP}$
Elements of the system		
P	$31.01^{+0.01}_{-0.01}$	$30.99^{+0.02}_{-0.02}$
T_0	$2452069.36^{+1.30}_{-1.30}$	$2452071.98^{+2.94}_{-2.93}$
e	$0.50^{+0.03}_{-0.03}$	$0.50^{+0.06}_{-0.06}$
ω	$339.87^{+3.10}_{-3.06}$	$334.36^{+6.10}_{-6.10}$
Elements of the radial velocity orbit		
k_1	$48.57^{+2.04}_{-1.87}$	$49.48^{+3.93}_{-3.44}$
γ_1	$-0.72^{+0.36}_{-0.36}$	$-0.41^{+0.47}_{-0.48}$

Table 5.5: The same as Table 5.4 except for GG Carinae.

corresponding top-level line-transition energies from 12.09 eV to 23.07 eV.

We first apply the χ^2 algorithm to the data and check that we recover predicted parameter distributions that are consistent with Porter et al. (2021a). Similarly to Porter et al. (2021a) we pre-set the γ values for each spectral line, but still fit for an overall offset. Next, we apply our marginalised $\mathcal{GP}$ algorithm to the same data and report the results in Table 5.5. We find that the distributions of all the parameters become more uncertain, by approximately a factor of two, owing to our marginalised $\mathcal{GP}$ algorithm properly accounting for the uncertainties.

In the top panel of Figure 5.5 we display a snapshot of the conditioned GP for GG Carinae, and in the bottom panel we show all of the radial velocity data folded by the derived period. In this fit we can see the data exhibits sporadic systematic offsets from the orbital motion. The GP is able to capture the systematics, and in doing so, the predictions for the orbital parameters become more uncertain. The median characteristic length scale of the marginalised GP kernel is 3.9 days which manifests itself as systematics varying on timescales faster than the orbital period.

The distributions from the literature and our prediction from for k_1 are broadly consistent, however the overall shift to higher values of k_1 , for our predicted distribution, results in the predicted mass ratio, $q = m_2/m_1$, moving to higher values. Therefore the extent to which the primary star's mass is greater than the secondary's may be smaller than previously predicted.




Figure 5.5: Top panel: a 147-day snapshot of the marginalised $\mathcal{GP}$ algorithm conditioned on GG Carinae’s observed data. The mean model is shown as the black dashed line, the GP is shown at sequential spectral line energies by the coloured lines from H-alpha (purple) to He I (green), and each shaded region represents the 2σ uncertainty in the GP. The γ values are arbitrary. Bottom panel: all the radial velocity data folded by the predicted period.

5.5.2.3 WR 140

WR 140 is a long-period (7.93 years), highly eccentric ($e = 0.9$) SB2 and visual system situated at a distance of $1.64^{+0.08}_{-0.07}$ kpc (Gaia Collaboration et al. 2016; 2018; Bailer-Jones et al. 2018). The stellar components are a carbon-rich Wolf-Rayet star (WC7pd) and an O-star (O5.5fc). We use the orbital parameters recently calculated by Thomas et al. (2021) as the current literature estimates.

We utilise the same data as Thomas et al. (2021): radial velocity measurements from Marchenko et al. (2003), Fahed et al. (2011), and Thomas et al. (2021), as well as relative astrometric measurements from Monnier et al. (2011) and Thomas et al. (2021). The dataset contains 465, 460, and 11 data points for the Wolf-Rayet radial velocities, the O-star radial velocities, and the relative astrometry, respectively.

We first apply the χ^2 algorithm to the data and check that we recover predicted parameter distributions that are consistent with Thomas et al. (2021). Next, we apply our marginalised $\mathcal{GP}$ algorithm to the same data. In this case the different spectral lines are combined into one radial velocity curve per stellar component and so we make use of only the time dimension in the GP. We report the results in Table 5.6. We find that the distributions of all the radial velocity parameters become more uncertain, as does the semi-major axis. However, the inclination and longitude of the ascending node are fairly consistent owing to their dependence being primarily related to the astrometric fitting procedure which remains unchanged between the two algorithms.

In the top panel of Figure 5.6 we display a snapshot of the conditioned GP for WR 140 for the most recent periastron passage. In the middle panel we show all of the radial velocity data utilised in the parameter estimation. In the close-up view of periastron we observe systematics in the Wolf-Rayet’s data that are reminiscent of the systematics presented in our synthetic example in Figure 5.2. The median characteristic length scale of the marginalised GP kernel is 130 days for the Wolf-Rayet’s data. For the O-star, as expected, the data shows only minor deviations from the orbital model, and the amplitude of the GP kernel reflects this through having a 3.4-times smaller value relative to the Wolf-Rayet’s kernel. Through applying our marginalised $\mathcal{GP}$ algorithm to these data we are able to make more accurate parameter

WR 140's orbit	Thomas et al. (2021)	marginalised $\mathcal{GP}$
Elements of the system		
P	$2895.00^{+0.29}_{-0.29}$	$2896.34^{+0.67}_{-0.64}$
T_0	$2460636.73^{+0.53}_{-0.53}$	$2460640.45^{+1.29}_{-1.33}$
e	$0.8993^{+0.0013}_{-0.0013}$	$0.8980^{+0.0029}_{-0.0030}$
ω_{WR}	$227.44^{+0.52}_{-0.52}$	$228.20^{+1.16}_{-1.18}$
Elements of the radial velocity orbit		
k_{WR}	$75.25^{+0.63}_{-0.63}$	$72.70^{+2.02}_{-2.04}$
k_{O}	$26.50^{+0.48}_{-0.48}$	$26.32^{+1.25}_{-1.29}$
γ_{WR}	$0.26^{+0.32}_{-0.32}$	$0.98^{+1.16}_{-1.14}$
γ_{O}	$3.99^{+0.37}_{-0.37}$	$2.32^{+1.00}_{-1.15}$
Elements of the astrometric orbit		
a	$8.922^{+0.067}_{-0.067}$	$9.011^{+0.117}_{-0.112}$
i	$119.07^{+0.88}_{-0.88}$	$119.09^{+0.89}_{-0.87}$
Ω	$353.87^{+0.67}_{-0.67}$	$354.43^{+0.88}_{-0.87}$

Table 5.6: Results for WR 140 comparing literature parameter distributions, using the χ^2 algorithm, and our estimated parameter distributions, using our marginalised $\mathcal{GP}$ algorithm. The uncertainties are the 68% credible intervals. The orbital parameters are the period, P (days), the time of periastron, T_0 (JD), the eccentricity, e , the argument of periastron of the Wolf-Rayet star, ω_{WR} (degrees), the Wolf-Rayet's and O-star's semi-amplitudes, k_{WR} and k_{O} (km s^{-1}), the Wolf-Rayet's and O-star's radial velocity offsets, γ_{WR} and γ_{O} (km s^{-1}), the semi-major axis, a (mas), the inclination, i (degrees), and the longitude of the ascending node, Ω (degrees).




Figure 5.6: Top panel: a 247-day snapshot of the marginalised $\mathcal{GP}$ algorithm conditioned on WR 140’s observed data. The mean models are shown as the black dashed lines, the GPs are shown for the Wolf-Rayet star (purple) and the O-star (yellow), and each shaded region represents the 2 σ uncertainty in the GPs. The γ values have been subtracted. Middle panel: all the radial velocity data used in the parameter estimation. Bottom panels: the relative astrometric data and the best-fit model, with the Wolf-Rayet star held fixed at the origin. ρ is the stellar separation and P.A. is the position angle.

predictions for WR 140. In this case the resulting parameter distributions, given the full complexity of the data, are more uncertain than previously predicted.

In the bottom panels of Figure 5.6 we display the simultaneous fit to the relative astrometric data. Given the availability of these data we can infer dynamical masses. Our results lead to the Wolf-Rayet star having a mass, $m_{\text{WR}} = 9.90^{+1.08}_{-1.01} M_{\odot}$, and the O-star having a mass, $m_{\text{O}} = 27.41^{+2.50}_{-2.34} M_{\odot}$. These masses are lower than previously estimated by Thomas et al. (2021; $m_{\text{WR}} = 10.31^{+0.45}_{-0.45} M_{\odot}$, $m_{\text{O}} = 29.27^{+1.14}_{-1.14} M_{\odot}$), however the distributions do still have overlap.

5.5.2.4 WR 133

WR 133 is an SB2 and visual system with a ~ 112.8 day period at a distance of $1.86^{+0.08}_{-0.08}$ kpc (Gaia Collaboration et al. 2016; 2021). The stellar components are a nitrogen-rich Wolf-Rayet star (WN5o) and an O-star (O9I). We use the orbital parameters recently calculated by Richardson et al. (2021) as the current literature estimates.

We utilise the same data as Richardson et al. (2021): radial velocity measurements from Underhill & Hill (1994) and Richardson et al. (2021), as well as relative astrometric measurements also from Richardson et al. (2021). The dataset contains 60 radial velocity measurements for each stellar component, and 8 relative astrometric measurements.

We first apply the χ^2 algorithm to the data and check that we recover predicted parameter distributions that are consistent with Richardson et al. (2021). Next, we apply our marginalised $\mathcal{GP}$ algorithm to the same data. The different spectral lines are combined into one radial velocity curve per stellar component and so we make use of only the time dimension in the GP. We report the results in Table 5.7. We find that the distributions of all the parameters are consistent with the literature predictions.

In the top panel of Figure 5.7 we display a snapshot of the conditioned GP for WR 133. We observe little to no systematics in both the Wolf-Rayet’s and O-star’s radial velocity data. In the middle panel we show all of the radial velocity data folded on the predicted period. Here we see how the deviations from the orbital model appear as white noise. As a result the median amplitude of the marginalised GP kernels for

WR 133's orbit	Richardson et al. (2021)	marginalised $\mathcal{GP}$
Elements of the system		
P	$112.780^{+0.036}_{-0.036}$	$112.763^{+0.030}_{-0.032}$
T_0	$2458702.08^{+0.38}_{-0.38}$	$2458701.90^{+0.32}_{-0.30}$
e	$0.3558^{+0.0050}_{-0.0050}$	$0.3578^{+0.0043}_{-0.0042}$
ω_{WR}	$45.3^{+6.1}_{-6.1}$	$47.5^{+6.5}_{-5.4}$
Elements of the radial velocity orbit		
k_{WR}	$32.30^{+3.02}_{-3.02}$	$34.04^{+4.49}_{-4.50}$
k_{O}	$14.63^{+1.51}_{-1.51}$	$13.91^{+1.23}_{-1.30}$
γ_{WR}	$78.1^{+3.0}_{-3.0}$	$79.0^{+3.4}_{-3.5}$
γ_{O}	$-15.09^{+0.48}_{-0.48}$	$-14.00^{+0.88}_{-0.91}$
Elements of the astrometric orbit		
a	$0.7863^{+0.0060}_{-0.0060}$	$0.7838^{+0.0052}_{-0.0043}$
i	$160.44^{+1.86}_{-1.86}$	$161.15^{+1.46}_{-1.55}$
Ω	$171.5^{+6.5}_{-6.5}$	$174.1^{+7.0}_{-5.7}$

Table 5.7: The same as Table 5.6 except for WR 133.




Figure 5.7: Top panel: a 517-day snapshot of the marginalised $\mathcal{GP}$ algorithm conditioned on WR 133’s observed data. The mean models are shown as the black dashed lines, the GPs are shown for the Wolf-Rayet star (purple) and the O-star (yellow), and each shaded region represents the 2σ uncertainty in the GPs. The γ values have been subtracted. Middle panel: all the radial velocity data folded by the predicted period. Bottom panels: the relative astrometric data and the best-fit model, with the O-star held fixed at the origin. ρ is the stellar separation and P.A. is the position angle.

each star are at the $\sim 3 \text{ km s}^{-1}$ level: a barely noticeable amount given the magnitude of the semi-amplitudes.

Through applying our marginalised $\mathcal{GP}$ algorithm to these data we are able to show how, when the noise present is predominately not correlated, the estimated parameters recover the same distributions as the χ^2 algorithms used throughout the literature. The marginalised $\mathcal{GP}$ is therefore safe to use in both cases of white noise and correlated noise. As for the dynamical masses of WR 133, our results lead to consistent predicted mass distributions with Richardson et al. (2021). We find the Wolf-Rayet star has a mass, $m_{\text{WR}} = 8.9_{-2.2}^{+3.4} \text{ M}_{\odot}$, and the O-star has a mass, $m_{\text{O}} = 21.9_{-7.6}^{+10.1} \text{ M}_{\odot}$. As noted by Richardson et al. (2021) the large uncertainties are mainly due to the nearly face-on geometry of the system, and they improve the precision of the masses through fixing the system at the Gaia distance.

5.6 Conclusions

In this chapter we aimed to create a new approach for estimating the orbital parameters of emission-line binaries. Our work is summarised as follows:

- (i) We synthesised benchmark datasets for testing the performance of different algorithms for Bayesian parameter estimation of the orbits of SB1 systems. The noise added to each dataset was designed to challenge potential models to be accurate in the presence of both white noise and problematic systematics. We have made these datasets freely available with the aim of making model validation an easy and standardised practice.
- (ii) We presented a novel application of Gaussian processes to model the radial velocity systematics of emission-line binaries, named marginalised $\mathcal{GP}$. We benchmarked this algorithm, along with current standardised algorithms, on the synthetic datasets. We found that the marginalised $\mathcal{GP}$ algorithm performs significantly better than the standard algorithms when the data includes complex systematics. Additionally, when the data contains no systematics the marginalised $\mathcal{GP}$ algorithm recovered a similar performance to the standardised algorithms.

- (iii) We applied the marginalised $\mathcal{GP}$ algorithm to four prototypical emission-line binaries. For Eta Carinae, GG Carinae, and WR 140 we found obvious systematics in the radial velocity data, and as a result the orbital parameter distributions predicted by our marginalised $\mathcal{GP}$ algorithm are different and more uncertain than the literature distributions. For WR 133 we did not discern any systematics and our predicted distributions are consistent with those in the literature.
- (iv) Overall, given the validation of the marginalised $\mathcal{GP}$ algorithm on the benchmark datasets, we expect the estimated parameter distributions for Eta Carinae, GG Carinae, and WR 140 to be more accurate than those previously calculated. As a direct consequence, the dynamical masses inferred for these systems may need to be updated.

Chapter 6

Tracing the colliding winds of Eta Carinae in He I

6.1 Introduction

Colliding-wind binaries are characterised by two massive stars both driving powerful stellar winds. These winds collide forming shock fronts at the location where their wind momenta balance. In these shocks the gas is compressed and heated, leading to the production of x-rays (Prilutskii & Usov 1976; Cherepashchuk 1976; Luo et al. 1990), non-thermal emission (Pollock 1987; White & Chen 1994; Hamaguchi et al. 2018), as well as dust (Williams 2008). The presence of the shock cone may also modulate the UV and optical line profile variability during the binary’s orbit (Stevens 1993; Szostek et al. 2012).

In this chapter, we take an in-depth look at the outflows of Eta Carinae, one of the most luminous and energetic colliding-wind binaries in the Milky Way. Eta Carinae has a luminosity of $L = 5 \times 10^6 L_{\odot}$ (Davidson & Humphreys 1997) and an observed long-term brightening due to a vanishing natural coronagraph (Damineli et al. 2021). The primary star is an LBV exhibiting one of the strongest winds on record: a mass-loss rate of $8.5 \times 10^{-4} M_{\odot} \text{yr}^{-1}$ and terminal wind velocity of 420 km s^{-1} (Hillier et al. 2001; Groh et al. 2012a; Clementel et al. 2015a). The companion star, despite having not been observed directly, is estimated to have a mass-loss rate and terminal wind velocity of typically $\sim (1-2) \times 10^{-5} M_{\odot} \text{yr}^{-1}$ and $\sim 3000 \text{ km s}^{-1}$, respectively. These are deduced from modelling the x-ray emission from Eta Carinae’s colliding winds (Pittard & Corcoran 2002; Okazaki et al. 2008; Parkin et al. 2011; Russell et al. 2016).

Photoionisation modelling of the spectral variability of Eta Carinae’s Weigelt blobs (Weigelt & Ebersberger 1986) indicated that the companion may be a hot O-star, with an effective temperature between 34,000 and 38,000 K (Verner et al. 2002; 2005). However, as noted by Smith et al. (2018) the companion’s wind parameters may be more consistent with those of a hydrogen-poor WR star. Smith et al. (2018) described a possible formation channel for this present-day companion via a past merger-in-a-triple scenario, and this has been shown to be tenable by the models of Hirai et al. (2021).

Whether the companion is an O-star or a WR star, it is known to generate a significant flux of Helium-ionising photons ($\lambda < 504 \text{ \AA}$, $h\nu > 24.6 \text{ eV}$) which significantly affect the ionisation structure of the system. There has been a substantial effort to map the ionisation structure of Helium within the central 155 au of Eta Carinae. This is because the Helium lines can be used to probe the conditions within the colliding wind interface, as has been utilised for other colliding-wind binaries such as V444 Cygni (Marchenko et al. 1994) and WR 140 (Williams et al. 2021).

As for Eta Carinae, there has been some speculation as to whether the He I lines are formed in the primary wind (Nielsen et al. 2007; Humphreys et al. 2008) or in the shock cone of the colliding winds (Damineli et al. 2008b). Recently, the three-dimensional radiative transfer simulations of Clementel et al. (2015a;b), taking into account the ionisation structure of the primary wind and the effect of the companion’s radiation, resulted in detailed Helium ionisation maps for Eta Carinae. Since then, spectro-interferometric observations, with the infrared K-band beam-combiner GRAVITY at the VLTI, have enabled milliarcsecond resolution observations of the He I 2s-2p ($2.0587 \text{ }\mu\text{m}$) line transition (GRAVITY Collaboration et al. 2018) and these results show good agreement with the Helium ionisation maps of Clementel et al. (2015a). In combination with these Helium ionisation maps, the estimated orbit from Chapters 4 and 5, and orientation (Madura et al. 2012) mean that our line of sight intercepts the shock cone, providing an exceptional opportunity to study the conditions and geometry of Eta Carinae’s colliding winds.

In this chapter, we investigate the geometry of Eta Carinae’s colliding winds through time-series spectroscopic observations of the optical He I line profiles. In Section 6.2 we present observations from the Global Jet Watch’s campaign on Eta

Carinae and show how at different orbital phases the He I 7065 Å absorption varies between one and two components. In Section 6.3 we employ multi-Gaussian fitting to track the dynamics of these emission and absorption components, revealing that one of the absorption component’s velocities deviates substantially from the expected orbital motion. We then describe a geometrical model of the colliding winds which simulates the He I absorption velocities, while being cognisant of the known Helium ionisation maps. We fit this model to the data and tabulate the inferred parameters. In Section 6.4 we discuss several complexities and assumptions of the model. Finally, in Section 6.5 we summarise our findings.

6.2 Observations

The Global Jet Watch is a network of five remotely controlled observatories separated in longitude to enable round-the-clock observing. Each observatory has a fibre-fed spectrograph, from a 0.5 m telescope, that collects mid-resolution ($R \sim 4000$) optical data in the wavelength range $\sim 5800 - 8400$ Å. The spectra are reduced by a bespoke data-reduction pipeline which includes dark correction, flat-field correction, and wavelength calibration derived from Thorium Krypton frames. Further details of the observatories are described in Section 1.3.

The Global Jet Watch’s observing campaign of Eta Carinae began in early 2014 and since has amounted 5858 spectra across 1.3 orbital periods (~ 7 years), spanning the 2014 and 2020 periastrons. These observations have exposure times ranging from 1 second to 300 seconds to capture both the H-alpha and He I lines at acceptable signal-to-noise ratios but avoid saturation. Pertinent to this study, the observations are densely time-sampled throughout the orbital period, not just at periastron, and thereby provide excellent coverage of Eta Carinae’s apastron behaviour.

6.2.1 The He I line profiles

In this chapter, we focus our analysis on the He I line profiles. As such, we select the 3807 Global Jet Watch spectra with exposure times of at least 100 seconds. These spectra are normalised by their local continuum and are barycentric corrected using the `barycorrpy` package (Kanodia & Wright 2018). To further improve the signal-to-noise ratio, especially for the weaker blue-shifted absorption components,




Figure 6.1: Selected He I 7065 Å spectral line profiles. The profiles are colour-coded with brighter yellow indicating the proximity to Eta Carinae’s periastron.

we median stack the spectra into orbital phase bins of width 0.01. These phases bins are computed according to the period, $P = 2022.7$ d, of Damineli et al. (2008a), and the time of periastron, $T_0 = 2454850.1$ (JD) of from Chapter 5.

Initial analysis of the three He I lines in the Global Jet Watch spectra shows that both Helium lines at 5876 Å and 6678 Å suffer blends from other species. As has been previously noted by Richardson et al. (2015), the red wing of the He I 5876 Å line is blended with the complex Na I D doublet and the blue wing of the He I 6678 Å line is blended with a [Ni II] line. As a result, we restrict our analysis to the unencumbered He I 7065 Å line, from which we can glean clear dynamical information from both the emission and absorption components.

In Figure 6.1, we plot six selected spectral line profiles showcasing the changing morphology throughout Eta Carinae’s orbit. In terms of emission we observe the expected variability: a low excitation event at periastron in which the narrow emission component disappears and the bulk velocity of the emission region shifts violently with the extremely eccentric periastron passage (Groh & Damineli 2004; Nielsen et al. 2007; Humphreys et al. 2008; Damineli et al. 2008a;b; Mehner et al. 2010; 2012; Richardson et al. 2015). As for the blue-shifted absorption, we discern two distinct components in these line profiles. The faster component, at a velocity of $\sim -550 \text{ km s}^{-1}$, is present throughout the orbit; however, the second component, found at slower negative velocities, is detected only either side of periastron. In Figure 6.1 this behaviour is apparent, where the two line profiles corresponding to phases $\phi = 0.98$ and $\phi = 0.03$ show only one absorption component, and in all the other line profiles we can clearly distinguish two. We measure this second component to disappear at $\phi = 0.95$ and reappear at $\phi = 0.08$.

6.3 Tracing the Colliding winds

In this section we extract dynamical information from the He I 7065 Å line and formulate a model for interpreting the remarkable behaviour of the absorption components.

6.3.1 He I absorption dynamics

To extract velocities from the He I 7065 Å lines we employ multi-Gaussian fitting to decompose the complex line profiles into their constituent dynamical components.




Figure 6.2: Example multi-Gaussian fitting of He I 7065 Å spectral line profiles near to periastron (top panel) and apastron (bottom panel). The top of each panel shows the total model (light blue line) and its Gaussian components (light green lines for emission and orange lines for absorption) versus the observations (black points). The bottom of each panel shows the residuals of each fit.

Following Nielsen et al. (2007), we construct a template comprised of four Gaussians to model the main emission, one or two negative Gaussians, depending on the epoch, to model the blue-shifted absorption, and a low-amplitude broad Gaussian to model the extended wings of the profile. The Gaussian parameters are optimised using a Trust Region Reflective algorithm (Branch et al. 1999; Jones et al. 2001) which minimises the sum of squares between the model and the data. We find this superposition of Gaussian components results in a model that produces consistently good fits across all of the spectra; this is validated by goodness-of-fit tests, such as assessing whether the normalised residuals show a normal distribution as suggested by Riener et al.




Figure 6.3: Velocities extracted from the multi-Gaussian fits to all of the He I 7065 Å spectral line profiles superimposed on the orbit estimated in Chapter 5. The green data points represent the weighted mean of the four main emission components and the black data points represent the slower of the two absorption components shown in Figure 6.2.

(2019). Example multi-Gaussian fits are displayed in Figure 6.2 for two epochs, one near to periastron and the other near to apastron.

After fitting the line profiles we examine the extracted velocities. For the bulk velocity of Eta Carinae’s wind emission region we take the weighted mean of the four main emission components (centroid wavelength of each Gaussian weighted by its area, see Blundell et al. (2007); Chapter 4). We plot these velocities in Figure 6.3 alongside the centroid velocities of the slower of the two absorption components. Note that the absorption velocities are not present near to periastron as this component is not detected at these times, as was discussed in Section 6.2.1.

In Figure 6.3 it is clear that the emission and absorption components encode dynamical information from distinct regions of Eta Carinae. Here, the dashed lines represent the orbit predicted by our models from Chapter 5, with arbitrary vertical offsets to enable visual comparison with the emission and absorption dynamics. The emission velocities show similar dynamics to the line-of-sight orbital motion of the system. However, the absorption velocities show apparent deviations from the orbital motion. These deviations are most prominent during apastron as the velocities arc upwards to slower values. Understanding these absorption dynamics will now form the focus of the remainder of this chapter.

6.3.2 The line formation regions of Helium

In order to model the dynamics of the He I 7065 Å absorption, we must first consider the Helium line formation regions in Eta Carinae. There has been a substantial effort to map the ionisation structure of Helium within the central 155 au of Eta Carinae. The companion is known to be a source of Helium ionising photons (Verner et al. 2002; 2005; Hillier et al. 2006) which significantly affects the Helium ionisation structure of the system; and hence, any quantities deduced from these lines. Recently, the three-dimensional radiative transfer simulations of Clementel et al. (2015a;b), which take into account the ionisation structure of the primary wind and the effect of the companion’s radiation, resulted in detailed Helium ionisation maps for Eta Carinae. These maps are mass-loss rate dependent, as the density of the primary wind changes how far the companion’s Helium ionising radiation can penetrate, but for our use case we select the simulations conducted for a mass-loss rate of $8.5 \times 10^{-4} \text{ M}_{\odot} \text{ yr}^{-1}$ from the primary star.

In Figure 6.4, we reproduce the schematic of the Helium ionisation map from GRAVITY Collaboration et al. (2018; their figure 14), which is based on the simulations of Clementel et al. (2015a; their figures 7 and 8) for the central 155 au of Eta Carinae. In this schematic we colour each region based on its Helium ionisation level: He⁰⁺ (blue), He¹⁺ (yellow), He²⁺ (grey). The primary wind is predominately He⁰⁺, except for two He¹⁺ regions. The first is the small volume surrounding the primary star which is ionised by radiation directly from the primary star. The second




Figure 6.4: Helium ionisation schematic of Eta Carinae: He^{0+} (blue), He^{1+} (yellow), He^{2+} (grey). The radius of the plot is ~ 155 au and the dashed lines represent the region of the post-shock primary wind formed in the interface of the colliding winds. This plot is reproduced from GRAVITY Collaboration et al. (2018), which itself was based on the simulations of Clementel et al. (2015a) for apastron.

is the region adjacent to the apex of the colliding wind, which is ionised by radiation from the companion. Here we see how only gas sufficiently close to the shock cone is able to be reached by the companion’s ionising flux. In addition to these two regions, He^{1+} also exists in the post-shock primary wind in the walls of the shock cone cavity, denoted by the dashed lines in Figure 6.4, and throughout most of the pre-shock companion wind. The post-shock companion wind is solely comprised of He^{2+} owing to collisional ionisation caused by the high temperature ($T > 10^6$ K) of the shock-heated gas.

We are primarily concerned with the He^{1+} regions of Figure 6.4, since the observed optical He I lines are indirect consequences of recombination into the highly energetic top levels of these lines. In particular, following Damineli et al. (2008b); Clementel et al. (2015a); GRAVITY Collaboration et al. (2018), we presuppose that the He I 7065 Å absorption is formed principally in the post-shock primary wind.

6.3.3 A geometrical model for the colliding winds

There exist certain combinations of system orientations and colliding wind balances in which the shock cone intersects our line of sight to the primary star. If the He I 7065 Å absorption is formed in the post-shock primary wind, along the walls of the shock cone, then this absorption will trace the geometry of the colliding winds as the shock cone swings across our line of sight. The dynamical information imprinted at the site of the absorption corresponds to the post-shock primary wind velocities projected onto our line of sight. This results in faster velocities the more parallel the shock cone wall is to our line of sight. Qualitatively, for orientations favoured by most investigators, with the shock cone on our side of the system during apastron, this leads to absorption velocities that appear fast as the line of sight first enters the shock cone, shortly after periastron, followed by a progressive slowing as the line of sight traces across to the centre of the shock cone at apastron. The process is then reversed as the line of sight traces out to the other side of the shock cone, the absorption velocities becoming faster, until the line of sight exits the shock cone altogether. This description mirrors the dynamical behaviour of the He I 7065 Å absorption found in Section 6.3 and so we pursue this line of thinking by formulating a geometric model of the colliding winds.

We begin the description of our geometric model by considering the wind momentum balance, as previously described in Section 2.4 by Equations 2.35, 2.36, and 2.37. Along the line of centres of the two stars the wind momentum ratio is defined as

$$\eta_{\text{wm}} = \frac{\dot{M}_2 v_{\infty,2}}{\dot{M}_1 v_{\infty,1}}, \quad (6.1)$$

where $\dot{M}$ is the mass-loss rate, v_{∞} is the terminal velocity of the wind, and the subscripts 1 and 2 correspond to the primary and companion stars, respectively. Using Equation 6.1 the apex of the contact discontinuity occurs at

$$r_1 = \frac{D}{1 + \sqrt{\eta_{\text{wm}}}}, \quad (6.2)$$

where r_1 is the distance from the primary star to the contact discontinuity and D is the total separation of the two stars (Stevens et al. 1992). This expression assumes that the winds collide at their terminal velocities which is valid for most of

the apastron passage of Eta Carinae. The balance of ram pressures when the stars have uneven wind strengths leads to a cone-like shape to the contact discontinuity, with the asymptotic half-opening angle of this shock cone, θ , being the fundamental descriptor. Previous hydrodynamic simulations by Madura et al. (2013) found that for Eta Carinae this θ was consistent with the analytic formula of Canto et al. (1996):

$$\eta_{\text{wm}} = \frac{\tan \theta - \theta}{\tan \theta - \theta + \pi}. \quad (6.3)$$

Given analytical expressions for the colliding wind apex's location and asymptotic opening angle we elect to use a hyperboloid as the geometric surface for the contact discontinuity. A hyperboloid of one sheet, orientated such that the apex and opening are along the x -axis, has the following equation:

$$\frac{x^2}{a^2} = \frac{y^2}{b^2} + \frac{z^2}{c^2} + 1, \quad (6.4)$$

where $a = r_1$ sets the distance from the primary star, situated at the origin, to the apex, $b = a \tan \theta$ sets the asymptotic half-opening angle in the x - y plane, and $c = a \tan \theta$ sets the asymptotic half-opening angle in the x - z plane. In this way we have included the flexibility to alter the wind momentum balance in both the equatorial and polar directions (see the model exploration in Figures 6.5 and 6.7 and the discussion in Section 6.4), although for the main model fitting we set both these values to be the same. A hyperboloid is a natural choice owing to its qualitative agreement with hydrodynamic models, and the ease with which the surface can encode the location and angles of the colliding winds.

To synthesise the absorption velocities of the post-shock primary wind of Eta Carinae we work in the rotated frame of reference, which has the line of centres of the two stars fixed along the x -axis, and the primary star at the origin. We compute the separation of the stars at each epoch of an orbital period and from this set the hyperboloid geometry after inputting the wind momentum ratio. Our line of sight to the primary star is then computed in this rotated frame and we can solve for the intersection, if any, between it and the hyperboloid, named as point P_{int} . The solution is found after some algebra, and we are careful to select only the intersections on our side of the system from the resulting quadratic equation. At P_{int} we compute the angle between the tangent plane to the hyperboloid and our line of sight, ϕ , in order to project the shocked gas velocities onto our line of sight.




Figure 6.5: Exploration of our geometrical model for the colliding winds of Eta Carinae for various parameter combinations. Only one parameter is varied per panel, see the text for details.

The final absorption velocities in the post-shock primary wind, $v_{\text{absor}}(t)$, as a function of time, t , are

$$v_{\text{absor}}(t) = v'_{\text{wind}}(t, P_{\text{int}}) \cos^2 \phi, \quad (6.5)$$

where

$$v'_{\text{wind}}(t, P_{\text{int}}) = v_{\text{wind}}(t, P_{\text{int}}) + v_{\text{kep}}(t) + \gamma. \quad (6.6)$$

In these equations v_{wind} is the velocity of the primary wind just prior to being shocked at point, P_{int} , v_{kep} is the Keplerian velocity projected onto our line of sight, and γ is the systemic velocity along our line of sight. The sum total of Equation 6.6 is the velocity of the pre-shock gas in our frame of reference. Equation 6.5 simply projects the pre-shock velocities, first along the shock cone, and then second, onto our line of sight. The angles of these projections are both ϕ as our line of sight is co-linear with the outflowing pre-shock wind. We also choose to calculate v_{wind} from the standard

β -law as first derived by Castor et al. (1975) for line-driven winds and we set $\beta = 1$. However, we find using the terminal velocity of the wind does not affect our results by more than a few km s^{-1} , as at times near to apastron the intersection point is far from the primary star.

Having described a geometrical model for the colliding winds of Eta Carinae, we now look at how the parameters that go into this model change the predicted absorption velocities. The complete parameter set is displayed in Table 6.1: the orbital period, P , time of periastron, T_0 , eccentricity, e , primary star's semi-amplitude, k_1 , systemic velocity along our line of sight, γ , wind momentum ratio, η_{wm} , primary wind terminal velocity, $v_{\infty,1}$, inclination, i , and argument of periastron, ω . In addition to these tabulated parameters we also require a value for the mass of the primary star, m_1 in order to calculate the size of the semi-major axis via the mass function and Kepler's third law. We set $m_1 = 100 M_{\odot}$ which results in a semi-major axis of 17.8 au, but tests show that this choice is arbitrary for our model. The semi-major axis simply rescales the system, having a negligible impact on the absorption velocities, much like the β -law parameterisation mentioned above. Consequently, we leave this parameter fixed and out of the tabulated list of key parameters in Table 6.1.

In Figure 6.5, we explore the dependence of the predicted absorption velocities on the model parameters. In each panel we vary one parameter and fix all of the others. The fixed set is $e = 0.9$, $k_1 = 66.8 \text{ km s}^{-1}$, $\gamma = 0 \text{ km s}^{-1}$, $v_{\infty,1} = 420 \text{ km s}^{-1}$, $i = 137.5^\circ$, and $\omega = 270^\circ$. We further subdivide the wind momentum ratio into its equatorial and polar directions, $\eta_{\text{wm}}^{\text{equatorial}} = 0.2$ and $\eta_{\text{wm}}^{\text{polar}} = 0.2$. We start with the top-left panel where we examine the role of the eccentricity. Here we find the resulting absorption velocities depend strongly on this parameter. The higher the eccentricity, the sooner our line of sight intersects with the shock cone after periastron, and the longer we trace the velocities up to slower values, as the shock cone sweeps more slowly across our line of sight at apastron. In the top-right panel and the middle-left panels we vary the argument of periastron and inclination, respectively. Here we find the model is very sensitive to changes in our line of sight. The argument of periastron shifts the orbital phase at which the slowest velocities (largest projection angles) occur. The inclination changes the overall range of velocities predicted, and to a lesser extent the timing of the start and end of the intersection. This can be

Model parameters	Value	Reference
Orbital elements		
P (days)	2022.7	Damineli et al. (2008a)
T_0 (JD)	2454850.1	Chapter 5
e	0.90	Chapter 5
k_1 (km s ⁻¹)	66.8	Chapter 5
γ (km s ⁻¹)	-8.1	Smith (2004)
Wind parameters		
η_{wm}	0.2	Pittard & Corcoran (2002)
$v_{\infty,1}$ (km s ⁻¹)	456	This work
Line of sight parameters		
i (°)	137.5	Madura et al. (2012)
ω (°)	271	This work

Table 6.1: Parameters used in the geometrical model for the colliding winds of Eta Carinae. “This work” indicates the parameters found in the best fit model shown in Figure 6.6.

interpreted as lower inclinations tracing a line through the centre of the shock cone, intersecting it at a variety of angles. However, for higher inclinations the intersection angles remain small throughout, and so the projected velocities do not change much over the orbit. In the middle-right panel we see how the velocity of the primary wind only serves to shift all of the absorption velocities to slower or faster values.

In the bottom panels of Figure 6.5 we test how the wind momentum ratio affects the predicted absorption velocities of our model. In the bottom-left panel we vary the equatorial ratio and find the model is largely insensitive to this parameter, except for some minor changes to the timing and velocities at the start and end of the intersection. This is the case because, with an inclination set at 137.5°, we trace through the shock cone high above the equatorial plane where the projected angles depend mainly on the polar ratio, as is seen in the bottom-right panel. In the bottom-right panel we find that as the wind momentum ratio increases, meaning the companion star has a more equally balanced wind, the shock cone opening angle increases and




Figure 6.6: Best fit of our geometrical model for the colliding winds of Eta Carinae to the He I 7065 Å absorption velocities. One orbital period has been duplicated.

the absorption velocities span a larger range. Conversely, for smaller wind momentum ratios our line of sight only just intersects with the cone as it becomes more closed, maintaining faster and more even velocities throughout the orbit. Once the wind momentum ratio becomes sufficiently small there will be no intersection at any time, however this is heavily dependent on the inclination angle as can be seen by comparing their influence on the resulting model in their respective panels. For spherically symmetric outflows we lock the equatorial and polar wind momentum ratios together and the model dependence is the sum of the effects in these two bottom panels.

Next, we fit our model to the He I 7065 Å absorption velocities extracted in Section 6.3.1. We fix the orbital elements to literature values: $P = 2022.7$ d (Damineli et al. 2008a); $T_0 = 2454850.1$ (JD), $e = 0.90$, and $k_1 = 66.8$ km s⁻¹ (Chapter 5); $\gamma = -8.1$ km s⁻¹ (Smith 2004). The inclination and wind momentum balance are severely degenerate which means that it is not possible to optimise both of these parameters simultaneously. Instead we use literature values of $i = 137.5^\circ$, the central value in the range estimated from models of the [Fe III] emission (Madura et al. 2012), and $\eta_{\text{wm}} = 0.2$, the value estimated from the duration of the x-ray lightcurve minimum (Pittard & Corcoran 2002). For the model fitting in this section we assume

a spherically symmetric wind balance. Finally, parameters $v_{\infty,1}$ and ω are left free to be optimised. A summary of this parameterisation is detailed in Table 6.1.

The model is optimised using the Levenberg-Marquardt algorithm (Moré 1978; Jones et al. 2001) to find the least-squares solution. We find the best fit values for our free parameters are $v_{\infty,1} = 456 \text{ km s}^{-1}$ and $\omega = 271^\circ$. We display this result in Figure 6.6. Here we see the model produces excellent fits to the data, showing the correct velocity trends as well as the correct timing of the entry (absorption first detected after periastron) and exit (absorption no longer detected before the next periastron) of our line of sight through the shock cone. Additionally, we render the hyperboloid corresponding to this best-fit solution in the top panel of Figure 6.7.

6.4 Discussion

We have presented a geometrical model for the colliding winds of Eta Carinae. The fit of this model to the He I 7065 Å absorption velocities shows excellent agreement. Moreover, the parameters input into the model are concordant with the other panchromatic observations and modelling of Eta Carinae. This includes direct agreement between the Helium ionisation simulations (Clementel et al. 2015a), VLTI/GRAVITY observations of He I 2.0587 μm , models of the x-ray light curve (Pittard & Corcoran 2002), three-dimensional hydrodynamic simulations (Okazaki et al. 2008), HST/STIS observations of [Fe III] (Madura et al. 2012), Gemini/GMOS observations of the Balmer lines (Chapters 4 and 5), and now also the Global Jet Watch observations of He I 7065 Å. This prevailing view of Eta Carinae’s fundamental parameters builds confidence in our understanding of the system: a highly eccentric colliding-wind binary oriented with the companion on the observer’s side of the system during apastron. In the remainder of this section we consider further details of our model within the constraints of this view of Eta Carinae.

6.4.1 A latitudinally dependent wind

During the model fitting undertaken in Section 6.3.3 we assumed spherical symmetry in the colliding winds, locking the wind momentum ratio to identical values in both the x - y and x - z planes. However, as was shown in the exploration of the model parameters in Figure 6.5, specifically the middle-left and bottom-right panels, the

inclination and wind momentum ratio in the polar direction are severely degenerate. In other words, the projected angles between our line of sight and the shock-cone can be held constant by simultaneously decreasing the inclination of the viewing angle and decreasing the polar wind momentum ratio. Consequently, any deductions made about the inclination or wind momentum ratio rely on an accurate value for the other.

To investigate how this may affect our results, we run some further tests varying the inclination angle within the range estimated by Madura et al. (2012). We fit the model with inclination values fixed at 130° and 145° . As for the wind momentum ratio, we keep the equatorial component consistent with the x-ray light curve observations (Pittard & Corcoran 2002), having a value of $\eta_{\text{wm}}^{\text{equatorial}} = 0.2$. We then allow the polar component, $\eta_{\text{wm}}^{\text{polar}}$, to be free in the fitting routine. For both trial inclination values we are able to generate equally good fits to the data. We find the resulting pairs of values are $i = 130^\circ, \eta_{\text{wm}}^{\text{polar}} = 0.11$ and $i = 145^\circ, \eta_{\text{wm}}^{\text{polar}} = 0.40$. For the case with an inclination of 130° , in order to remain consistent with the other observables, the polar wind momentum ratio must be lower than the equatorial wind momentum ratio, 0.11 versus 0.20. This may imply a stronger polar outflow from the primary star. We visualise this case in the bottom panel of Figure 6.7. Here we see subtle changes in the hyperboloid surface and line of sight, relative to the original best-fit seen in the top panel from Section 6.3.3. For the case with an inclination of 145° , the opposite implications would be true.

There exists a precedent in the literature which may indicate that we should expect a latitudinally dependent wind from Eta Carinae. Theoretical work to include non-radial line-forces and gravity darkening in models of rapidly rotating luminous stars predicts the mass flux may be enhanced in the polar directions (Cranmer & Owocki 1995; Owocki et al. 1996). Further to this, observations of the Balmer lines by Smith et al. (2003) indicate that a polar enhanced wind may be present in the primary star of Eta Carinae. Based on these assumptions and our model fitting, we may expect the inclination to be towards the lower end of the range estimated by Madura et al. (2012). However, for us to draw quantitative conclusions we need to improve the certainty in our measurements of the inclination.




Figure 6.7: Three-dimensional renderings of the hyperboloid surfaces from our geometric models. The models are shown in the rotated frame of reference with the primary star at the origin. The dashed lines show the line of sight at three different phases: 0.1, 0.5, 0.9. Top panel: the best-fit model from Section 6.3.3 for a spherically symmetric wind with $i = 137.5^\circ$ and $\eta_{\text{wm}} = 0.20$. Bottom panel: the best-fit model from Section 6.4.1 for a polar enhanced wind with $i = 130^\circ$, $\eta_{\text{wm}}^{\text{equatorial}} = 0.20$, and $\eta_{\text{wm}}^{\text{polar}} = 0.11$.

6.4.2 Model caveats

As discussed above, the model and results presented include several assumptions and simplifications worth noting. First, is the omission of the Coriolis effect from the model’s geometry. This effect causes the shock cone to trail behind the axis of the two stars; however, it is unlikely to cause noticeable changes to the model at times around apastron as can be seen in hydrodynamic simulations of Eta Carinae (Okazaki et al. 2008; Parkin et al. 2011; Clementel et al. 2015a). In fact, the instabilities and turbulence in the shocked flows are likely to be a larger source of uncertainty at these times. However, for the few data points at the very start and end of the absorption signal there will certainly be a Coriolis effect, as well as modifications to the Helium ionisation structure (Clementel et al. 2015b). This may be manifested in our detection of the second Helium absorption component only between phases $\phi = 0.08$ and $\phi = 0.95$, which shows a noticeable asymmetry about periastron, with a delay in our line of sight re-entering the shock cone.

Next, we have simplified the geometry of the colliding winds to the surface of a hyperboloid. Whilst this formalism is certainly a useful tool, and it recovers the asymptotic opening angles at large distances, the exact shape of the apex may differ in reality. We also assumed that the He I absorption only forms in the post-shock primary wind. This is based on the fact that He^{1+} is most dense in the post-shock primary wind, more so than in the pre-shock primary wind, and orders of magnitude more so than in the companion’s wind which is thought to make no observable contribution to the spectra. This assumption is also validated by how well our model fits the data. But there may still be contributions from both the pre-shock and post-shock absorption in the spectra. The faster absorption component in our multi-Gaussian fits, from Section 6.3.1, is perhaps a product of this pre-shock line formation. Overall, in spite of several simplifications made in our geometrical model we find it explains the observations of Eta Carinae’s He I absorption very well.

6.5 Summary and conclusions

In this chapter we have investigated the unique dynamics of the He I 7065 Å absorption velocities in the colliding-wind binary Eta Carinae. Our work is summarised as

follows:

- (i) We made use of Global Jet Watch observations of Eta Carinae covering the last 1.3 orbital periods (~ 7 years). The unprecedented coverage throughout apastron enabled us to extract clear dynamical information at these orbital phases.
- (ii) We employed a multi-Gaussian fitting algorithm to the He I 7065 Å line profile. We found the number of absorption components appears to vary between one and two depending on the orbital phase. The slower of these two absorption components is only detected between phases $\phi = 0.08$ and $\phi = 0.95$ ($\phi = 0.0$ is periastron), and displays velocities that deviate from the orbital-like motion which is found in the line's emission.
- (iii) To interpret these deviations we conjectured that this absorption component of the He I 7065 Å line is formed in the post-shock primary wind and is only detected when our line of sight intersects with the shock cone. We formulated a geometrical model for the colliding winds in terms of a hyperboloid in which the opening angle and location of its apex are parameterised in terms of the system's wind momentum ratio. The absorption velocities of the post-shock primary wind are computed as the sum of the shocked wind velocities projected onto our line of sight, the orbital motion, and the systemic velocity of the system.
- (iv) We fitted our geometrical model to the He I 7065 Å line absorption velocities finding results that are concordant with the panchromatic observations and simulations of Eta Carinae.
- (v) The model presented in this chapter is an extremely sensitive probe of the exact geometry of the wind momentum balance of binary stars. Given more certainty in the inclination of Eta Carinae with respect to our line of sight, the model could be used to probe the latitudinal dependence of the mass loss quantitatively.

Chapter 7

Mahalanobis ensembles of co-trending light curves for correcting spatially-correlated photometric artefacts

7.1 Introduction

In this chapter we shift our focus from spectroscopic data to photometric data. The brightness of a star can be extracted from a time-series of images to analyse its so-called light curve. These light curves encode valuable time-domain information about their underlying sources. Accurate signals enable the discovery and characterisation of a wide variety of systems from exoplanet transits to variable stars and supernovae. However, these data may be affected by systematics which can hinder analysis of the true astrophysical variability.

Systematics may appear in light-curve datastreams as a consequence of imperfections in instrumentation, algorithms, and ambient observing conditions. From an instrumentation perspective, systematics may be introduced as a result of variations in the point spread function across the frame, CCD efficiency, or non-uniform illumination across the detector. Photometric extraction pipelines may not be perfect and may be hindered by edge cases or blended sources which are difficult to process. For ground-based imaging, atmospheric transparency variations may become the dominant source of systematic errors. The transparency may change as a function of observatory, time, and sky pointing owing to differences in the Rayleigh scattering by air molecules, absorption by water vapour and ozone, and aerosol scattering by




Figure 7.1: Examples of spatially-correlated artefacts from the Global Jet Watch photometric datastream. The left-hand panels correspond to images from the SS433 field and the right-hand panels correspond to images from the GG Carinae field. The image brightness has been truncated below and above the 5th and 95th percentiles to best visualise the large-scale systematics.

particulates such as dust, water droplets/clouds, and pollutants (Hayes & Latham 1975). Additionally, for small-aperture telescopes we can expect the statistical noise budget to be substantial due to the lower signal-to-noise observations obtained by these instruments.

Previously, algorithms have been developed to correct for these systematics. For example, TFA (Kovács et al. 2005) and SysRem (Tamuz et al. 2005) which both aim to remove linearly varying deviations that are built from the entire sample of stars. The next generation of this approach, based on principal component analysis, resulted in the Kepler Presearch Data Conditioning Bayesian maximum a posteriori

(PDC-MAP) pipeline used for Kepler (Stumpe et al. 2012; Smith et al. 2012; Stumpe et al. 2014) and TESS (Jenkins et al. 2016; Feinstein et al. 2019). There have also been algorithms developed that utilise hierarchical tree clustering to identify subsets of stars with similar systematics, such as PDT (Kim et al. 2009), and algorithms that select comparison stars by iterative correlation searches, such as DOHA (Mislis et al. 2017). Specific to K2 (Howell et al. 2014), the extension of the Kepler mission, are algorithms for handling the loss in pointing accuracy after the satellite’s second reaction wheel failed. For example, EVEREST (Luger et al. 2016) which applies pixel level decorrelation and a Gaussian process to disentangle the astrophysical signals. Finally, for some ground-based surveys, such as the Dark Energy Survey (Dark Energy Survey Collaboration et al. 2016) and the Vera Rubin Observatory’s Legacy Survey of Space and Time (LSST; LSST Science Collaboration et al. 2009) (LSST Science Collaboration et al. 2009), the atmospheric systematics can be physically modelled by making use of auxiliary instrumentation to calculate the top of atmosphere flux, known as Forward Global Photometric Calibration (Burke et al. 2017).

In this chapter we introduce the newly commissioned photometric datastream at the Global Jet Watch. Each observatory has dual filter cameras that capture 5.6 by 4.5 degree images of the sky simultaneously to the usual spectroscopic observing. Initial analysis of the images reveals the presence of challenging spatially-correlated artefacts, as shown by a few example images in Figure 7.1, that lead to considerable systematics in the extracted light curves. To correct these data we formulate a novel algorithm, specifically targeted at dealing with spatially-correlated artefacts, and describe our approach in Section 7.2. In Section 7.3 we synthesise four test datasets for validating the performance of this algorithm and compare the results with several of the other detrending algorithms from the literature. In Section 7.4 we provide further details about the Global Jet Watch photometric datastream and apply the algorithm to two example fields. Finally, in Section 7.5 we summarise our findings.

7.2 The Algorithm

For ground-based differential photometry data, such as ours, we propose a two-stage algorithm for correcting spatially-correlated artefacts. The algorithm works independently for different filters and corrects each star in each image separately. For a given

image and filter, the data used to determine these corrections are structured such that for every star, i , we measure the raw counts per second, r_i , the median counts per second across all images for the matching filter, m_i , and the position on the CCD in pixels, $\mathbf{p}_i = (x_i, y_i)^T$.

The first-stage correction follows the standard approach to account for the atmospheric extinction (eg. Burki et al. 1995). The atmospheric extinction corrected counts per second for the i th star, c_i , is

$$c_i = r_i 10^{\alpha X_i / 2.5}, \quad (7.1)$$

where α is the atmospheric extinction coefficient and $X_i = \sec z_i$ is the approximate airmass for a plane-parallel atmosphere at an observing angle, z_i , from the zenith. The second-stage correction involves a correction factor, β_i , which refines the counts per second by

$$f_i = \beta_i c_i, \quad (7.2)$$

where f_i is the final corrected counts per second.

The key to our algorithm lies in how it learns the value of β_i . By comparing each star's counts per second to its median counts per second (both after the first-stage correction) we can find the deviation, $d_i = m_i / c_i$. We assume that the overwhelming majority of stars are not intrinsically variable; and therefore, the deviation as a function of pixel location on the detector, and by extension, the sky, tells us how an individual image is affected by artefacts. If we were to set $\beta_i = d_i$, then we would correct every star back to its median counts per second and have perfectly stable light curves. This would work well for all the non-variable stars, but for transient sources this correction would remove the variability we are interested in. Instead, we use the surrounding stars, local to a given source on the detector, to inform the correction factor. In this way, our method uses the population dominance of the stable stars to inform good correction factors, without removing interesting astrophysical signals, and the spatially-correlated artefacts are accounted for.

The question now is: for each star how do we determine this surrounding ensemble of stars from which the correction factor can be learned? Rather than setting an arbitrary fixed radius, inside of which a star is included in the correction factor calculation, we want an algorithm that learns both which stars to include and how

much to weight their contribution. Hereafter, we refer to the collection of stars that will be used to correct a specific source as the co-trending ensemble of stars. Intuitively, the composition of these ensembles will depend on the nature of the artefacts present in the images. For images with minor systematics spread across a wide number of pixels, the ensembles should be large and utilise as many stars as possible. While for images with major systematics that vary fast across the detector, the ensembles should be small and localised to the few stars that are co-trending with each source.

To achieve this behaviour we define a *Mahalanobis ensemble*. A Mahalanobis ensemble includes all the stars in a given image, where each star is assigned a weight according to its Mahalanobis distance. The Mahalanobis distance,

$$d_{\Sigma}^{\text{mah}}(\mathbf{p}_i, \mathbf{p}_j)^2 = (\mathbf{p}_i - \mathbf{p}_j)^T \Sigma_i^{-1} (\mathbf{p}_i - \mathbf{p}_j), \quad (7.3)$$

is a measure of the distance from a nearby star, j , to a distribution, specified by the covariance matrix, Σ_i , centred at the star being corrected, i . The distribution is chosen to be a bivariate normal distribution where the optimal covariance matrix is learned from a *deviations map* that we construct from the image. The resulting weights of the ensemble are used to compute the correction factor, β_i , by way of a weighted mean of each star's deviations.

To help elucidate the workings of the algorithm, over the subsequent sections we first present a one-dimensional version of the Mahalanobis ensemble. We then describe the full two-dimensional implementation and further details of the optimization process.

7.2.1 Proof of concept I: symmetric Gaussians

When the Mahalanobis distance is calculated using a one-dimensional normal distribution, the algorithm is equivalent to finding the ensemble weights using a symmetric Gaussian centred on the star being corrected, i . A normal distribution is selected so that stars closer to star i have more weight in the calculation of the correction factor as they are more likely to have suffered a similar deviation.

In Figure 7.2 we create an illustrative test case for this symmetric Gaussian version of the algorithm. On the left-hand side we generate $55^2 = 3025$ stars on an evenly




Figure 7.2: Left side: deviation surface, d , of a synthetic test image. The dashed black circles show the one standard deviation contours, σ , of a few example symmetric Gaussian weighting distributions. Right side: the top panel shows the mean squared error (MSE) of the corrected stars for different σ . The bottom three panels each display a test showing how the optimal σ changes based on the density of stars, the size of the deviations, and the signal-to-noise of the stars.

spaced square grid of 2500 pixels, each with $c_i = m_i = 9000$ and corresponding Poisson noise. The deviations of each star from their medians are then manipulated by synthetic systematics varying across the detector. In this test we use the so-called monkey saddle function, defined by $d = x^3 - 3xy^2$, to create a smoothly varying problematic surface. We translate the surface to the centre of the detector and stretch it such that the maximum deviation is 0.1. For all stars in the image we calculate the correction factor for 10 symmetric Gaussians with incrementally growing standard deviations. The dashed lines show how a few of these distributions compare to the detector scale for the central star.

In the top panel on the right-hand side of Figure 7.2 we plot how the quality of the correction depends on the standard deviations of the Gaussians, σ . To assess the quality of the correction applied to all the stars in the test image we use the mean squared error (MSE) between the final corrected counts per second, f_i , and the

medians, m_i . We find that there exists a minimum MSE corresponding to $\sigma = 83$ pixels. For too wide a distribution the MSE becomes larger due to the correction factor up-weighting stars from a region of the image that are affected in a slightly different way by the image systematics. These stars have different deviation values and are no longer co-trending with the actual i th star being corrected. For too narrow a distribution the MSE becomes larger as the number of stars with large weights in the ensemble becomes small. Consequently, the correction factor is adversely affected by the statistical noise present. The result of this test is that there exists an optimal ensemble of co-trending stars that minimises the MSE.

There are several factors that affect the optimal ensemble. First, the density of the stars in an image. In the second-from-top panel on the right-hand side of Figure 7.2 we test how changing the total number of stars in our test image affects the optimal σ . We find that as the density of the stars increases, the optimal σ decreases. This is because there are more sources in the ensemble with large weights, even for Gaussians with small σ values, and so the preference to use stars that are spatially correlated with the target star wins out over the statistical noise fluctuations. Second, in the second-from-bottom panel on the right-hand side of Figure 7.2 we test how changing the maximum size of the injected deviations in our test image affects the optimal σ . We find that as the deviations increase, the optimal σ decreases. This is because for larger systematics the gradients in the deviation map increase, and therefore fewer close-by stars occupy the same regime as the target star. Third, in the bottom panel on the right-hand side of Figure 7.2 we test how changing the signal-to-noise ratio of the stars in our test image affects the optimal σ . We find that as the stars get brighter, and the signal-to-noise ratio increases, the optimal σ decreases. This is because for brighter stars the statistical fluctuations are smaller and so the algorithm only needs to trust fewer stars to find an accurate correction factor.

7.2.2 Proof of concept II: Mahalanobis ensembles

After our investigations with symmetrical Gaussians, we now return to the full two-dimensional algorithm. In this case the Mahalanobis distance is calculated for a bivariate normal distribution. This additional flexibility afforded by the two-dimensional distribution allows us to form ensembles that are more resilient to difficult artefacts.

The mathematics of the algorithm are as follows. The correction factor for star i is defined as

$$\beta_i = \frac{\sum_{j \neq i}^n k(i, j) d_j}{\sum_{j \neq i}^n k(i, j)}, \quad (7.4)$$

where d_j is the deviation of star j , $k(i, j)$ is the weight of star j , and n is the number of stars in the image. Note that we do not include the star's own deviation ($j = i$) in its weighted average. We compute the weight of star j , for corrections to star i , via the kernel function

$$k(i, j) = \exp\left(\frac{-d_{\Sigma}^{\text{mah}}(\mathbf{p}_i, \mathbf{p}_j)^2}{2}\right), \quad (7.5)$$

where d_{Σ}^{mah} is the Mahalanobis distance as described in Equation 7.3. The Mahalanobis distance depends on the distribution specified. In our case, a bivariate normal distribution is determined by the covariance matrix, Σ_i , and is centred on the star being corrected, i , such that the squared Mahalanobis distance can be written as

$$d_{\Sigma}^{\text{mah}}(\mathbf{p}_i, \mathbf{p}_j)^2 = \mathbf{S}_{ij}^T \Sigma_i^{-1} \mathbf{S}_{ij}, \quad (7.6)$$

where $\mathbf{S}_{ij}$ is the separation vector of the stars i and j .

The crux of the algorithm is in how we formulate Σ_i . As we have found in Section 7.2.1, the optimal distribution depends on several quantities. In particular, we make use of the deviation of each star as a function of its positions on the image. We create a deviation map, $d(x, y)$ by interpolating the deviations of each star onto the entire image grid, using a nearest-neighbour approach. This map is then Gaussian smoothed to spread the deviation information by a chosen distance. For all the data in this study we select a smoothing length-scale of 250 pixels after some empirical tests.

We compute the gradients in the deviation map, $\nabla d(x, y)$, to detect how the deviations vary locally at each star's position. The idea is that we can then construct Σ such that its first principal component is perpendicular to the direction of the deviation gradient, and the length of this component is learned from the magnitude of the gradient. As Σ is positive semi-definite we decompose it into

$$\Sigma_i = \mathbf{V}_i \Lambda_i \mathbf{V}_i^T, \quad (7.7)$$

where $\mathbf{V}_i$ is a matrix where the columns are the principal components (eigenvectors) and Λ_i is a diagonal matrix where the values are the lengths of the principal




Figure 7.3: Synthetic test images with various deviation maps injected across panels a, b, c, and d. The cut-out triangle in panel a reveals the underlying density of the stars used in each image. The white dashed lines indicate the 1σ outlines of the learned distributions used in the Mahalanobis distance computation at various locations in each image.

components (eigenvalues) for the bivariate normal distribution. The first principal component is a normalised vector in the direction perpendicular to the gradient at star i , and the second principal component is constructed so that is orthogonal to the first. The eigenvalues that fill the diagonal of $\mathbf{\Lambda}_i$ are

$$\lambda_{1,i} = (a + b \|\nabla d(\mathbf{p}_i)\|)^2, \quad (7.8)$$

$$\lambda_{2,i} = (c - d \|\nabla d(\mathbf{p}_i)\|)^2, \quad (7.9)$$

where a , b , c , and d are free parameters to be learned, which, together, comprise a linear transformation of the amplitude of the deviation gradient at the i th star's position. In Section 7.2.4 we describe how these parameters are found for each image.

Our formulation of the covariance matrix in this way utilises both the gradient's direction and magnitude to learn the optimal ensemble weights across the detector. This construction of the bivariate normal distribution enables it to extend furthest across the image in the direction along which the star's deviations are the smallest. The overall size of the distribution is also controlled by the steepness of the gradient. In other words, the co-trending ensemble of stars assigns the largest weights to the surrounding stars that are located along contours of similar deviation. This leads to correction factors that are more accurate than in the symmetric Gaussian case, especially when the deviations have complex morphologies.

In Figure 7.3, we show the results of the Mahalanobis ensembles algorithm for four illustrative synthetic images. We generate $30^2 = 900$ randomly positioned stars on a square grid of 2500 pixels each with $c_i = m_i = 5000$ and corresponding Poisson noise. The deviations of each star from their medians are then synthetically manipulated by systematics varying across the detector. In each panel we plot the deviations map and overlay the 1σ outlines of the bivariate normal distributions learned at various positions in the images. In Figure 7.3a, there are no deviations present in the image and so the algorithm learns to incorporate stars from all over the image in the correction factor. In Figure 7.3b, we inject a boundary at which the deviations rapidly change. Here we see how the algorithm is able to construct distributions that elongate along the deviation contours, avoiding stars from regions with different deviations, whilst still utilising many stars in the calculation in regions far away from the boundary. In Figure 7.3c, we create a particularly difficult ring-shaped problem. This is a good example of how the orientations of the distributions are learned from the deviation gradients. In Figure 7.3d, we have constructed a more realistic and subtle case of systematics slowly varying across the image. Overall, the panels in Figure 7.3 showcase the flexibility of the ensembles and how, in the limit of no systematics, the algorithm utilises most stars in the image for correction.

7.2.3 Proof of concept III: intensity correlations

The algorithm we have described so far is designed to correct spatially-correlated artefacts in the images. However, the deviations in an image may also be dependent on the absolute intensity of the sources.

Our algorithm, in the way that it is formulated, can be thought of as a kernel method. Kernel methods enable different kernel functions to be easily combined. We extend Equation 7.5, to also account for intensity correlations, and so it becomes

$$k(i, j) = \exp\left(\frac{-d_{\Sigma}^{\text{mah}}(\mathbf{p}_i, \mathbf{p}_j)^2}{2}\right) \times \exp\left(\frac{-d^{\text{int}}(c_i, c_j)^2}{2}\right), \quad (7.10)$$

where d^{int} is a measure of the intensity similarity between stars i and j . We define this measure as

$$d^{\text{int}}(c_i, c_j)^2 = \left(\frac{\log(c_i/c_j)}{l^2}\right)^2, \quad (7.11)$$

where the numerator quantifies the difference in logarithmic counts per second, and l is a free parameter that scales the weights for a given image. In the limit of there being no correlation in intensity, the value of l grows to large values and the weights are only dependent on d_{Σ}^{mah} .

7.2.4 Optimization

To implement the algorithm we rewrite Equation 7.6 as

$$d_{\Sigma}^{\text{mah}^2} = \mathbf{S}^T (\mathbf{V} \mathbf{\Lambda} \mathbf{V}^T)^{-1} \mathbf{S} = \mathbf{S}^T \mathbf{V} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{V}^T \mathbf{S} = \mathbf{W}^T \mathbf{W}, \quad (7.12)$$

where

$$\mathbf{W} = \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{V}^T \mathbf{S}, \quad (7.13)$$

and we have used the fact that $\mathbf{V}^{-1} = \mathbf{V}^T$ as it is orthonormal, and we have dropped the indices for notational convenience. This form enables us to more easily vectorise the computation.

In total the algorithm has five free parameters, $\{a, b, c, d, l\}$, which are optimized for each image separately. The objective function is based on minimising the MSE between the final corrected counts per second, f_i , and the medians, m_i . This loss function is defined as

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (m_i - f_i)^2, \quad (7.14)$$

where n is the total number of stars being corrected. To optimize Equation 7.14 we utilise `pytorch` (Paszke et al. 2019), along with its automatic differentiation and the stochastic gradient-based optimizer, Adam (Kingma & Ba 2015).

This optimisation problem requires that the eigenvalues of $\mathbf{V}$ remain positive. To constrain the optimisation we pass Equations 7.8 and 7.9 through a sigmoid function and rescale the output values such that

$$\lambda_{i,\text{constrained}}^{\frac{1}{2}} = \frac{\sigma_{\max} - \sigma_{\min}}{1 + e^{-\lambda_i^{\frac{1}{2}}}} + \sigma_{\min} \quad (7.15)$$

where we set $\sigma_{\min} = 25$ pixels and $\sigma_{\max} = 2500$ pixels which act as the minimum and maximum standard deviations along each principal component of the bivariate Gaussian. We also find that initially rescaling the gradient magnitudes, such that $0 < \|\nabla d(\mathbf{p}_i)\| < 1$ for all i , helps to improve the optimization.




Figure 7.4: Top panel: optimization tests for various mini-batch sizes in units of number of sources per sky area. Bottom panel: optimization example using a learning rate scheduler to improve convergence. Note that the learning rate is fixed at $\rho_A = 4.0 n_{\text{srcs}} \text{deg}^{-2}$ and the y-axis is logarithmic in this panel.

To speed up the optimization process we implement mini-batches, where we only compute the loss function for a subset of the data within an image for each iteration. This reduces the size of the matrices when computing the co-trending ensembles, as well as meaning the gradients and parameters are updated several times per optimization epoch. In the top panel of Figure 7.4 we test how different sizes of mini-batches affects the optimization. We run the tests based on the density of stars in the image to make the results transferable between different observing modes and fields of view. We find that as the density of sources decreases the loss improves more rapidly. However, once the density becomes too low the optimization becomes unstable (eg. the dark blue line at $\rho_A = 4.0 n_{\text{srcs}} \text{deg}^{-2}$).

To overcome this limitation of small mini-batch sizes we utilise a learning rate scheduler. We start with an initial learning rate of $\eta = 0.3$ and decay the learning rate by a factor of 10 when the loss fails to decrease by at least 0.01% over 3 subsequent epochs. The learning rate is decayed twice and on the third occasion the decay condition is met we consider the optimization converged. In the bottom panel of Figure 7.4 we show an example of how our learning rate scheduler prevents overshooting and the convergence of the algorithm is improved.

7.3 Synthetic benchmarks

In this section we create synthetic datasets and use them to validate the performance of our algorithm (hereafter referred to as the Mahalanobis ensembles algorithm), along with a handful of other algorithms from the literature.

7.3.1 Data synthesis

We synthesise four datasets each with 100 images evenly spaced in time. We use the sources extracted from a 300 second Global Jet Watch exposure of the SS433 field as the template for the brightness and position of the synthetic sources. This template image contains 5918 sources distributed across a 5.6 by 4.5 degree field of view. In each image we imprint these sources onto a 2749 by 2199 pixel grid with Gaussian point spread functions having standard deviations of 0.65 pixels. These sources are kept steady in their intrinsic brightness throughout each dataset. We add 1800 counts of sky background to every pixel, similar to what is found in the Global Jet Watch

images for a 300 second exposure, and apply an atmospheric extinction correction with $\alpha = 0.17$ and with an airmass that is linearly increasing over the 100 images in the interval $1.1 < X < 2.0$. Finally, we add Global Jet Watch dark and bias frames to the images, as well as Poisson noise.

In addition to the 5918 steady sources, we inject 10 eclipsing binaries using the stellar and exoplanet light-curve occultation code **starry** (Luger et al. 2019). These eclipsing binaries have varying stellar luminosity ratios, eclipse timings, and signal-to-noise ratios. These transients will serve as an important test of whether the correction algorithms remove any of the precious astrophysical signals from the light curves.

The four synthetic datasets are then separated and into each we generate various spatially-correlated artefacts. For dataset 1 we make no changes, leaving only Poisson noise, as though we have ideally observed images. For dataset 2 we apply an attenuating layer evenly across the entire image. This attenuation creates deviations that are linearly increasing over the 100 images in the interval $0.0 < d < 0.2$. For dataset 3 we apply systematic deviations to all pixels with x -values greater than a certain value. These x -values are randomly sampled per image from a value within the width of the image. The amplitude of the deviations per image is randomly sampled from an exponential distribution with a rate parameter of 0.1. For dataset 4 we apply smoothly-varying systematic deviations on multiple scales by randomly sampling a two-dimensional Gaussian process. This Gaussian process is characterised by a Rational Quadratic Kernel, which has the form

$$k(\mathbf{X}_p, \mathbf{X}_q) = A \prod_{r=1}^2 \left[1 + \frac{(\mathbf{X}_{pr} - \mathbf{X}_{qr})^2}{2al_r^2} \right]^{-a}, \quad (7.16)$$

where A controls the amplitude, l_r controls the characteristic length scale of the correlations in each input dimension, and a determines the relative weighting of variations on different scales. The values of l_r are randomly sampled from a uniform distribution over the interval 300 to 3000 pixels, in both x and y , and a is set to 10. The amplitude of the deviations per image is sampled in the same way as for dataset 3. Overall these four datasets will challenge the correction algorithms in both the absence and presence of systematics, and with artefacts that cause problems in both time and space across a variety of different length scales.

Algorithm (option)	Dataset 1		Dataset 1	
	Steady	Transients	Steady	Transients
None	1429.10	6.55	1322.73	29.46
α correction	1.77	1.34	135.08	18.92
Constant factor	1.37	1.29	1.31	1.16
SysRem (1 component)	1.69	2.33	1.31	4.75
SysRem (2 components)	1.30	6.50	1.19	5.64
SysRem (3 components)	1.27	8.33	1.16	7.15
PDT	1.77	1.34	1.23	2.92
DOHA	1.02	2.35	1.25	2.74
Mahalanobis ensembles	1.36	1.30	1.30	1.15

Table 7.1: Benchmark scores on synthetic datasets 1 and 2 with spatially-correlated artefacts. We benchmark published algorithms plus the Mahalanobis ensembles algorithm presented in this study. Each number listed is a score which is the mean squared error between the corrected light curves and the true injected light curves for the steady and transient sources separately. Bold numbers indicate the best score for each dataset on the transient light curves.

7.3.2 Synthetic tests

We run our Mahalanobis ensembles algorithm on the four synthetic datasets outlined in Section 7.3.1. In addition to this algorithm we benchmark several algorithms from the literature to assess and compare the effectiveness of our new methodology. We score each algorithm by calculating the mean squared error between the corrected light curves and the true injected light curves, and this calculation is applied to the steady and transient sources separately. These scores are presented in Tables 7.1 and 7.2.

In Tables 7.1 and 7.2 we start by computing the scores for the raw light curves with no correction applied. In row two we compute scores after only the α correction factor for atmospheric extinction has been applied. These are the baseline scores from which the other algorithms should improve. In row three we implement a simplistic algorithm in which the entire image is corrected by a constant factor which is found

Algorithm (option)	Dataset 3		Dataset 4	
	Steady	Transients	Steady	Transients
None	1864.83	14.34	1634.93	15.28
α correction	242.58	16.76	119.21	13.12
Constant factor	94.73	7.11	49.45	7.87
SysRem (1 component)	56.27	5.80	35.88	5.76
SysRem (2 components)	29.01	3.41	18.58	3.98
SysRem (3 components)	20.40	2.98	10.72	2.99
PDT	2.31	1.82	38.74	5.89
DOHA	7.98	1.99	5.23	1.59
Mahalanobis ensembles	14.45	1.53	2.63	1.33

Table 7.2: The same as Table 7.1 but for synthetic datasets 3 and 4.

through minimising the MSE. This algorithm provides a great comparison for our Mahalanobis ensembles algorithm in the limit of no spatially-correlated deviations. In datasets 1 and 2 where the deviations are absent or affect the entire image evenly we find the performance between this algorithm and the Mahalanobis ensembles algorithm are almost identical. However, in datasets 3 and 4, where spatially-correlated systematics are present, the Mahalanobis ensembles algorithm substantially outperforms the constant factor approach.

In the remaining rows of Tables 7.1 and 7.2 we compute scores for some previously published algorithms. We implement the algorithm SysRem (Tamuz et al. 2005) which removes linearly varying deviations that are found in the entire sample of stars. The next generation of this approach resulted in the Kepler Presearch Data Conditioning Bayesian maximum a posteriori (PDC-MAP) pipeline used for Kepler (Stumpe et al. 2012; Smith et al. 2012; Stumpe et al. 2014) and TESS (Jenkins et al. 2016; Feinstein et al. 2019). However, we are not able to include a score for PDC-MAP owing to the code not being open-source. Whilst the SysRem algorithm is not as sophisticated as PDC-MAP it does help to demonstrate some of the capabilities of principal component analysis for correcting light curve data. We implement SysRem




Figure 7.5: Light-curve normalised root-mean-squared deviations (NRMSD) as a function of instrumental source magnitude for each of the four synthetic datasets. The data points represent the raw extracted light curves (black), after atmospheric extinction correction (α correction, purple), and after our Mahalanobis ensembles algorithm has been applied (β correction, orange).

with up to three components and find it produces mixed results across the four synthetic datasets. In datasets 1 and 2, although the scores for the steady stars improve as more components are removed, the scores for the transient sources worsen. This is an indication that the SysRem algorithm is reducing the variance across the entire sample of stars, but at the expense of the transients' variability we want to preserve. In datasets 3 and 4 we find the scores for both the steady and transient sources decrease as the number of components removed increases. After three components have been removed the scores are not better than the Mahalanobis ensembles algorithm.

We implement two further algorithms: PDT (Kim et al. 2009) a hierarchical




Figure 7.6: Light curves for four selected transients from synthetic dataset 4. The left-hand column shows the light curves after the atmospheric extinction correction (α correction), and in the middle column we show the light curves after our Mahalanobis ensembles algorithm has been applied (β correction). The black dashed lines are the injected signals. The right-hand column shows colour-corresponding histograms of the deviations between the corrected data and the injected signals.

tree clustering algorithm to identify subsets of stars with similar deviations, and DOHA (Mislis et al. 2017) an algorithm that selects comparison stars by a two-step correlation search. For both of these algorithms we find they produce good scores across all four synthetic datasets. In particular, the DOHA algorithm is very effective on dataset 4.

Lastly, in the final row of Tables 7.1 and 7.2 we compute the scores for our Mahalanobis ensembles algorithm. In datasets 1 and 2 this algorithm and the constant factor approach have almost identical scores and beat all the other algorithms. This demonstrates that for deviations that affect the whole image our algorithm converges to the optimal approach. In datasets 3 and 4 we find that our Mahalanobis ensembles algorithm performs the best out of all the algorithms we benchmarked. These scores validate the algorithm as being effective at correcting for all sorts of spatially-correlated systematics.

In Figure 7.5 we display the resulting improvements to each light curve for each synthetic dataset after the α and β correction factors have been applied. We plot each light curve’s normalised root-mean-squared deviation (NRMSD) as a function of instrumental source magnitude ($-2.5 \log_{10}(f_i)$). In this figure we see how all the light curves have a signal-to-noise ratio greater than 5 ($\text{NRMSD} < 20\%$) due to the usual detection constraints in the source extraction pipeline. The light curves follow a curved distribution that results from the Poisson noise in the images: brighter sources have light curves with lower variance.

For dataset 1, in the top-left panel of Figure 7.5, as expected the atmospheric extinction correction removes almost all of the systematic variability in the light curves, owing to the images in this dataset having no artefacts. For the other three datasets, in the remaining panels of this figure, we find a clear improvement as each stage of correction is applied. For dataset 2, these sequential improvements are the most obvious with the final data being of similar quality to dataset 1. For dataset 3, we deduce clear striations in the uncorrected and first-stage corrected data points due to the distinct systematic regimes imposed on these data. In the β -corrected data we find there are a few sources that lie above the bulk of the distribution. We identify these sources as often coinciding with the hard boundaries of the injected systematics, making them particularly difficult sources to correct. For dataset 4, comprised of images with the most realistic systematics, we find our algorithm produces particularly good results despite the challenges of this dataset.

In Figure 7.6 we select four transient sources from synthetic dataset 4 and display their light curves after each of the two-stage correction procedure. In the left-hand column we plot the light curves after the atmospheric extinction correction (α correction), and in the middle column we plot the light curves after our Mahalanobis ensembles algorithm has been applied (β correction). We find clear improvements in all of the transient light curves, as is shown by the accompanying histograms in the right-hand column. Importantly, these improvements are not at the expense of any of the true astrophysical variability.

Whilst our results here are primarily concerned with fixing spatially-correlated noise, we ran an additional test to check how our algorithm handles systematics that are correlated by the source intensity. We augmented synthetic dataset 4 to include

Algorithm (option)	Dataset 4 + intensity correlations	
	Steady	Transients
None	364.05	18.19
α correction	29.34	18.81
Constant factor	7.53	6.65
SysRem (1 component)	5.76	6.14
SysRem (2 components)	3.54	3.42
SysRem (3 components)	2.81	3.33
PDT	2.29	2.17
DOHA	3.23	1.97
Mahalanobis ensembles	2.03	1.99

Table 7.3: Benchmark scores on synthetic dataset 4 with intensity-correlated systematics added. We benchmark published algorithms plus the Mahalanobis ensembles algorithm presented in this study. The score is the mean squared error between the corrected light curves and the true injected light curves for the steady and transient sources separately.

correlations that vary linearly with intensity at a similar amplitude as for the spatial dimension. We re-ran all the benchmark algorithms and present the results in Table 7.3. Here, we find good performance from our Mahalanobis ensembles algorithm, but we also find that the PDT and DOHA algorithms both gain ground in this regime. These results indicate that for datasets where the artefacts are more heavily dependent on the source intensity and less dependent on their position, other algorithms, such as PDT, DOHA, and PDC-MAP, may be the more favoured approach.

7.4 The Global Jet Watch MC2 data stream

Details of the Global Jet Watch are described in Section 1.3, however we briefly re-state the important information regarding the photometry cameras here. Each observatory has two cameras that piggyback on the main telescope. These cameras benefit from the pointing and guiding of the telescope but have a separate optical pathway; such instrumentation is thus described as commensal. These many com-




Figure 7.7: A Mollweide projection in equatorial coordinates showing a map of the sources detected in the Global Jet Watch MC2 datastream as of July 2021. Dark blue represents areas of the sky not observed thus far. Also shown are the positions of the sources SS433 (light blue triangle) and GG Carinae (light blue cross) which are the fields analysed in Section 7.4.2.

mensal cameras are hereafter referred to as the MC2 cameras.

7.4.1 The observing campaign

The MC2 cameras capture 5.6 by 4.5 degree images of the sky onto cooled 2749 by 2199 pixel CCDs. There are two filters centered at 656nm (hereafter named the $H\alpha$ filter) and 736nm (hereafter named the 736 filter) that operate in tandem. These cameras take predominately 300 second exposures, but shorter modes are available for particularly bright targets of interest.

The observing campaign for the MC2 cameras runs simultaneously to the spectroscopic observations. The Global Jet Watch spectroscopic observations target time-series observations of several targets, and as a result we obtain time-series images for many patches of the sky. In Figure 7.7 we show the current map of sources detected over the initial 2.5 year observing campaign. In total there are currently $\sim 300,000$ sources with light curves of at least 100 data points.

The photometry pipeline for source extraction, astrometry, and source association is built using the superb open-source tools from `photutils` (Bradley et al. 2020), `astrometry.net` (Lang et al. 2010), and software pipelines for the Large Synoptic

Survey Telescope (LSST, Axelrod et al. 2010), respectively. All data are stored in a relational database which includes flags for data quality and transient discoveries. Once a source has been marked as variable, we thereafter prohibit its deviation data from being included in the calculations of correction factors.

After the photometry pipeline has processed all of the currently available images we examine the raw light curves. We find deviations between the light curves and their median values that occur both intra-image and on a night-by-night basis. We examine the deviation maps within the images and find that these systematics are predominately spatially correlated, with little to no correlation with source intensity.

Given the validation of our Mahalanobis algorithm – detailed in Section 7.2 and tested in Section 7.3 – and the nature of the systematics present in the Global Jet Watch images, we select this algorithm for correcting our data rather than any of the previously published algorithms. In addition to being well-suited to correcting spatially-correlated artefacts, the Mahalanobis algorithm has no requirements for evenly-spaced observations and no requirements for a consistent set of sources either. For these reasons this algorithm is ideal for correcting the MC2 data given the agile observing modes utilised by the Global Jet Watch.

7.4.2 Global Jet Watch tests

In this section we elucidate the results of applying our Mahalanobis algorithm to the Global Jet Watch MC2 datastream. In particular, we analyse two fields that well-represent the broader Global Jet Watch MC2 dataset: the SS433 field and GG Carinae field. These fields are named after the central spectroscopic target of interest which the MC2 cameras capture, as well as the surrounding sources. We select ~ 200 days of images for each field for showcasing the corrected light curves.

In Figure 7.8, we display the resulting improvements to each light curve, for each of the two example Global Jet Watch fields, after the α and β correction factors have been applied. We plot each light curve’s NRMSD as a function of source magnitude. The left-hand panels correspond to data from the SS433 field and the right-hand panels correspond to data from the GG Carinae field. The GG Carinae field is more densely populated by sources relative to the SS433 field having ~ 30000 and ~ 5000 sources respectively, due to the proximity of the former to the Galactic plane. The

Global Jet Watch observing campaign of the GG Carinae field operates at a number of shorter exposure lengths in this region. Accordingly, there are brighter sources present, that remain unsaturated, and extend to slightly lower magnitudes. The panels in the top row correspond to data from the $H\alpha$ filter and the panels in the bottom row correspond to data from the 736 filter.

Across all the panels in Figure 7.8 we see a broad similarity in structure with the synthetic data presented in Figure 7.5. The two-stage correction improves the quality of the light curves shifting the distributions down, to lower variance, and slightly left, to brighter values. The brightening of the sources results from the α correction, as the counts per second of the sources are corrected to their above-atmosphere levels, and the β correction, which largely accounts for the variable and spatially-correlated attenuation. The final corrected light curves show the characteristic distribution associated with the Poisson noise underlying the data, with the brighter sources having least variance.

The main difference between the synthetic test data and these real fields is the extent and variety of the noise present in the raw light curves of the real fields. These data, especially for the GG Carinae field, are afflicted by systematics that imprint substantial deviations in the light curves. Additionally, we find that the sky background has a large number of counts in many of the images, which may limit the quality of the corrected light curves even with optimal correction factors. These results are apparent in the right-hand panels of Figure 7.8 where the fully-corrected data (yellow points) maintain a noticeable spread to larger variances. Overall, our Mahalanobis algorithm yields corrected light curves that show a decrease in their variance, and for the vast majority of data the corrected light curves are now viable data products for analysis.

In Figures 7.9 and 7.10 we select four transient sources from each of the 200-day snapshots of the Global Jet Watch SS433 and GG Carinae fields, respectively. We display the transient light curves after each of the two-stage correction procedure. In the left-hand column we plot the light curves after the atmospheric extinction correction (α correction), and in the right-hand column we plot the light curves after our Mahalanobis ensembles algorithm has been applied (β correction). In general, the systematics remaining after the α correction act to attenuate the signal to lower




Figure 7.8: Light curve normalised root-mean-squared deviations (NRMSD) as a function of source magnitude for each of the two example Global Jet Watch fields (SS433 and GG Carinae) in two filters (H α and 736). The data points represent the raw extracted light curves (black), after atmospheric extinction correction (α correction, purple), and after our Mahalanobis ensembles algorithm has been applied (β correction, orange).




Figure 7.9: Light curves for four selected transients from the example SS433 Global Jet Watch field. The left-hand column shows the light curves after the atmospheric extinction correction (α correction), and in the right-hand column we show the light curves after our Mahalanobis ensembles algorithm has been applied (β correction). The black dashed lines are splines to merely aid visual comparison of the light curves before and after the second-stage correction.

values. We see a clear improvement in these systematics after the β correction. The light curves become far more cohesive showing clear trends through time. A few outliers do remain in these data, and subsequent inspection of the original images shows severe artefacts are present. For these data it may be the case that they must simply be pruned owing to the instantaneous difficulties with the original images.

Starting from the top row in Figure 7.9, these transient sources are identified as the following objects¹. First, Global Jet Watch MC2 source id = 29933 is identified as IRAS 19198+0224 (alternatively Gaia DR2 4288895419164344960) situated at RA = 19h 22m 22.396s and DEC = +02° 30′ 14.566″, with a G-band magnitude of 9.9162 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.0 and an instrumental magni-

¹We cross-matched our Global Jet Watch data using the SIMBAD astronomical database (Wenger et al. 2000).




Figure 7.10: The same as Figure 7.9, but for the GG Carinae field.

tude of -5.40. We find variability in this source at the $\sim 43\%$ level over timescales of ~ 200 days. Second, Global Jet Watch MC2 source id = 12130 is identified as IRAS 19167+0500 (alternatively Gaia DR2 4293633523999388928) situated at RA = 19h 19m 14.515s and DEC = $+05^\circ 06' 05.434''$, with a G-band magnitude of 9.0205 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.6 and an instrumental magnitude of -6.74. We find variability in this source at the $\sim 25\%$ level over timescales of ~ 100 days. Third, Global Jet Watch MC2 source id = 8966 is identified as IRAS 19101+0237 (alternatively Gaia DR2 4268167219828665856) situated at RA = 19h 12m 37.862s and DEC = $+02^\circ 43' 00.856''$, with a G-band magnitude of 9.9706 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.4 and an instrumental magnitude of -5.86. We find variability in this source at the $\sim 42\%$ level over timescales of ~ 85 days. Fourth, Global Jet Watch MC2 source id = 29697 is identified as IRAS 19184+0142 (alternatively Gaia DR2 4264815770927400576) a known variable star of Mira Cet type (Alfonso-Garz3n et al. 2012; Heinze et al. 2018). This star is situated at RA = 19h 20m 57.005s and

DEC = $+01^{\circ} 48' 28.732''$, with a G-band magnitude of 10.4741 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.2 and an instrumental magnitude of -5.50. We find variability in this source at the $\sim 134\%$ level over timescales longer than the 200-day snapshot used to generate the figures.

Moving on to Figure 7.10, these transient sources are identified as the following objects. First, Global Jet Watch MC2 source id = 628 is identified as TYC 8958-1828-1 (alternatively Gaia DR2 5337991415170945664) situated at RA = 10h 57m 44.181s and DEC = $-60^{\circ} 58' 37.079''$, with a G-band magnitude of 9.9135 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.2 and an instrumental magnitude of -6.39. We find variability in this source at the $\sim 13\%$ level over timescales of ~ 100 days. Second, Global Jet Watch MC2 source id = 3158 is identified as IRAS 10505-5903 (alternatively Gaia DR2 5338439569235719680) situated at RA = 10h 52m 35.828s and DEC = $-59^{\circ} 19' 38.9174''$, with a G-band magnitude of 9.5835 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 21.3 and an instrumental magnitude of -6.58. We find variability in this source at the $\sim 9\%$ level over timescales of ~ 100 days. Third, Global Jet Watch MC2 source id = 3663 is identified as V* V669 Car (alternatively Gaia DR2 5338420602658479744) situated at RA = 10h 53m 23.761s and DEC = $-59^{\circ} 35' 12.336''$, with a G-band magnitude of 11.0978 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 17.8 and an instrumental magnitude of -4.87. We find variability in this source at the $\sim 106\%$ level over timescales of $\gtrsim 150$ days. Fourth, Global Jet Watch MC2 source id = 35186 is identified as SV* HV 1230 (alternatively Gaia DR2 5338203891444566528) situated at RA = 10h 52m 58.760s and DEC = $-60^{\circ} 40' 11.886''$, with a G-band magnitude of 10.7726 (Gaia Collaboration 2018). In the Global Jet Watch data this corresponds to a median signal-to-noise ratio of 18.5 and an instrumental magnitude of -5.11. We find variability in this source at the $\sim 42\%$ level over timescales of ~ 70 days.

7.5 Conclusions

In this chapter we set out to create a new algorithm for correcting spatially-correlated photometric artefacts, motivated by the systematics present in the new Global Jet Watch MC2 datastream. Our work is summarised as follows:

- (i) We introduced a novel two-stage algorithm for correcting systematics in images of photometric datastreams. The first-stage correction accounts for atmospheric extinction and the second-stage targets spatially-correlated artefacts remaining in the images.
- (ii) We formulated the second-stage correction per image per star by defining a Mahalanobis ensemble: an ensemble of stars that are weighted according to how similarly they are co-trending with the star being corrected. These weights were computed via the Mahalanobis distance where the distribution was chosen to be a bivariate Gaussian which itself was learned from the artefact gradients present in the images.
- (iii) We synthesised four test datasets for validating the performance of our algorithm. Each synthetic dataset was injected with steady and transient sources, as well as various spatially-correlated systematics. We found that our algorithm recovered the true light curves best when compared with several other detrending algorithms from the literature. Importantly, these improvements were not at the expense of any of the true astrophysical variability in the transient sources.
- (iv) We detailed the new Global Jet Watch MC2 datastream and applied our Mahalanobis algorithm to two example fields over a 200-day period. We found our algorithm substantially improved the quality of the light curves despite challenging systematics being present. The Global Jet Watch MC2 datastream has a database currently containing $\sim 300,000$ light curves ready for analysis, and observations are scheduled to continue over the next few years.

Chapter 8

Epilogue

In this thesis we have studied some of the most luminous and massive emission-line stars in various binary systems. We have developed new and innovative methods for characterising the fundamental parameters of these systems and, applying these to several case studies, we have presented improved inferences for the orbits and outflows of several famous LBVs and WR-stars. Here, we summarise our findings from each of the previous chapters and describe possible avenues for future research.

In Chapter 4, we began by considering binary systems which are composed of stars driving powerful stellar winds. For a star with dense outflows, the effective photospheric radius is increased, precluding direct observation of the Keplerian motion of the star itself. Instead, the observables are broad lines emitted over large radii in the stellar wind (Beals 1929). In fact, the winds show stratification in the line emission regions dependent on each line’s excitation energy (Bowen 1928; Kuhl 1973; Willis 1982) and opacity (Hillier 1987).

To analyse the effect of extended and stratified line emission on the extracted radial velocities, and hence on the inferred orbital solutions, we investigated the dynamics of 10 Balmer lines from the LBV Eta Carinae. Eta Carinae was selected because, characteristic of LBVs, it has an enormous mass-loss rate, $8.5 \times 10^{-4} \text{ M}_{\odot} \text{ yr}^{-1}$, and a relatively slow wind, 420 km s^{-1} (Hillier et al. 2001; Groh et al. 2012a; Clementel et al. 2015a), resulting in the longest wind flow times to the emission-line regions of almost any emission-line star known today. As a consequence, the observational effects will be at their most pronounced in this system. As an instructive exercise, we fitted each of the Balmer line’s radial velocities with a purely Keplerian model. As expected, this revealed strong, systematic discrepancies among the inferred orbital

solutions. The trends showed that the higher energy Balmer lines have associated radial velocity curves that appear phase shifted later in time, as well as being less eccentric and lower in amplitude. The monotonic trends were also asymptotic: the higher energy lines appear to converge towards a single set of parameter values.

To overcome these discrepancies, we formulated a novel semi-analytical model which encapsulates both the star’s orbital motion and the propagation of the wind. The model encodes the integrated velocity field of the out-flowing gas in terms of a convolution of past motion due to the finite flow speed of the wind, named convolutional Keplerian motion. For Eta Carinae, this model enabled the 10 Balmer lines to be fit concurrently with a single set of orbital parameters: time of periastron $T_0 = 2454848$ (JD), eccentricity $e = 0.91$, semi-amplitude $k = 69 \text{ km s}^{-1}$ and longitude of periastron $\omega = 241^\circ$.

After this successful application of the model to Eta Carinae, we investigated how the model may translate to a larger population of more typical emission-line star systems. We produced forward models for the Wolf-Rayet star in RMC 140b, having a mass-loss rate of $1.26 \times 10^{-5} \text{ M}_\odot \text{ yr}^{-1}$ and terminal wind velocity of 1300 km s^{-1} (Shenar et al. 2019). For this WR-star, we demonstrated that for commonly used lines in radial velocity analysis, such as He II and N IV/V, deviations are still expected between the underlying Keplerian orbit and the predicted radial velocities. We also investigated the B[e] supergiant, GG Carinae, with a mass-loss rate of $2.2 \times 10^{-6} \text{ M}_\odot \text{ yr}^{-1}$ and terminal wind velocity of 265 km s^{-1} (Porter et al. 2021a). Analysis of 6 optical emission-lines found the lines apparently displayed significant deviations from Keplerian motion. We applied our model simultaneously to all 6 of these lines and estimated improved orbital parameters of $P = 31.01$ days, $T_0 = 2452069.4$ (JD), $e = 0.50$, $k = 48.6 \text{ km s}^{-1}$, and $\omega = 339.9^\circ$.

More generally, we found that the magnitude of the deviations between radial velocity curves and the true Keplerian motion depends on the interplay between the line-formation kernels and the orbit. The more extended the line emission regions relative to the orbital time-scales, the larger the deviations. The work from this chapter indicates that corrective modelling, such as the model presented in this thesis, is necessary in order to identify a consistent set of orbital parameters, independent of the emission line used. This is important for future high accuracy work.

In Chapter 5, we continued to build on the work from the previous chapter. Whilst the model for convolutional Keplerian motion made substantial strides towards unifying the discrepant radial velocities of various spectral lines, there are often residual systematics present in the data. Even for emission-line stars in which the radial velocities of specifically chosen lines are expected to show purely Keplerian motion and well-represent only the stars’ orbital motion, the complex and extended circumstellar environments may lead to systematics. This is problematic when orbital parameter probability distributions are estimated assuming radial velocity data are solely comprised of an orbital model plus Gaussian white noise, leading to overconfident and inaccurate orbital solutions, with implications for the inferred dynamical masses and hence evolutionary models.

To tackle these problems we want to create a sea change in the approach taken when analysing these systems. We conceived of a data analysis work flow in which the investigator first tests their models on a benchmark dataset, before fitting their models to their observations. In this way, the investigator can ensure their model is robust and performant, as well as providing a quantitative measure by which the best modelling techniques can take precedence in the field. To this end, we synthesised benchmark datasets, each with a different and challenging noise formulation, for testing the performance of different algorithms. We have made these datasets freely available with the aim of making model validation an easy and standardised practice in this field.

We described a novel application of Gaussian processes to model the radial velocity systematics of emission-line binaries, named marginalised $\mathcal{GP}$. This algorithm is comprised of a mean model plus a Gaussian process, to model the orbital motion plus any systematics, respectively. The hyperparameters of the Gaussian process are marginalised over to best capture the predictive uncertainty in the orbital parameters. We benchmarked this algorithm, along with current standardised algorithms, on the synthetic benchmark datasets and find our marginalised $\mathcal{GP}$ algorithm performs significantly better than the standard algorithms when the data includes systematics.

We applied the marginalised $\mathcal{GP}$ algorithm to four prototypical emission-line binaries. We revisited Eta Carinae and GG Carinae from the chapter before, as well as analysing WR 140 and WR 133. For Eta Carinae, GG Carinae, and WR 140 we found

obvious systematics in the radial velocity data, and as a result the orbital parameter distributions predicted by our marginalised $\mathcal{GP}$ algorithm are more uncertain than the literature distributions. For WR 133 we did not discern any systematics and our predicted distributions are consistent with those in the literature. Overall, given the validation of the marginalised $\mathcal{GP}$ algorithm on the benchmark datasets, we expect the estimated parameter distributions for Eta Carinae, GG Carinae, and WR 140 to be more accurate than those previously calculated. As a direct consequence, the dynamical masses inferred for these systems, and other similar systems, may need to be updated.

In Chapter 6, we took an in-depth look at the outflows of Eta Carinae. Eta Carinae has a companion star which also has a strong wind, resulting in an extremely luminous and energetic colliding-wind binary. The combination of its orbit and orientation, with respect to our line of sight, enables direct investigation of the conditions and geometry of the colliding winds. We analysed He I line profiles from the Global Jet Watch observatories covering the last 1.3 orbital periods (~ 7 years) of Eta Carinae. The unprecedented coverage throughout apastron reveals the distinct dynamics of the emitting versus absorbing components: the emission follows orbital velocities whilst one of the absorption components is detected only around apastron and exhibits velocities that deviate substantially from the orbital motion.

To interpret these deviations, cognisant of the known Helium ionisation maps, we conjectured that this absorption component of the He I line is formed in the post-shock primary wind, and is only detected when our line of sight intersects with the shock cone. We formulated a geometrical model for the colliding winds in terms of a hyperboloid in which the opening angle and location of its apex are parameterised in terms of the ratio of the wind momentum of the primary star to that of the secondary. We fitted this geometrical model to the absorption velocities, finding results that are concordant with the panchromatic observations and simulations of Eta Carinae. The model presented in this chapter is an extremely sensitive probe of the exact geometry of the wind momentum balance of binary stars in general. Given more certainty in Eta Carinae’s inclination, we demonstrate how our model could be used to probe the latitudinal dependence of the mass loss.

In Chapter 7, we switched data types from spectroscopy to photometry. In this chapter we described the newly commissioned commensal imaging datastream at the Global Jet Watch’s observatories. These cameras are designed to capture light curves of hundreds of thousands of stars, however we found many of the light curves suffer from predominately spatially-correlated artefacts in the raw images. Motivated by these issues, and wanting to optimise the utility of the data, we formulated a novel algorithm that learns optimal ensembles of co-trending stars for computing correction factors per image per star, called Mahalanobis ensembles. Each Mahalanobis ensemble is comprised of the surrounding stars weighted according to their Mahalanobis distances: the distance from a comparison star to a bivariate normal distribution centred at the target star, which is itself learned from the artefact gradients present in the images. We synthesised test datasets for validating the performance of this algorithm, before applying it to the Global Jet Watch data. We found our algorithm substantially improves the quality of the light curves despite the challenging systematics present.

8.1 Future directions

The physical and statistical models presented in Chapters 4 and 5 provide a framework for making inferences of the orbits of emission-line stars from a variety of spectral lines. These algorithms have been applied to observational datasets of the LBV Eta Carinae, the B[e] supergiant GG Carinae, and the WR-stars WR 140 and WR 133. In order to strengthen these results, simply expanding the number of case studies would provide further information on the success of these models. For testing convolutional Keplerian motion it is key that any new observations and radial velocity data span a range of spectral line energies, and are extracted from systems with a range of orbital periods, as these two scales control the magnitude of the expected deviations from pure Keplerian motion. As for the marginalised $\mathcal{GP}$ algorithm, moving forward it is important that the uncertainties in the inferred orbital solutions, and hence dynamical masses are propagated through to any subsequent analysis of stellar evolution. Changes in these masses, or increases in uncertainty, may lead to different conclusions that can be made about a system’s evolution. Future studies should be cognisant of

this when using results from studies that have swept any systematic effects under the carpet.

From a theoretical perspective, the model presented in Chapter 4 makes several simplifications in its formulation, as discussed in Section 4.4. It would be interesting to compare the semi-analytical formulation with the output of a three-dimensional hydrodynamic simulation. In particular, a hydrodynamic simulation would be able to incorporate the effects due the wind flow time, self-consistently with the location of the star. The influence of a companion on the primary wind could also be quantified by investigating different colliding-wind cavities and radiative inhibition or braking regimes. The entire simulation could be post-processed with a radiative transfer code to understand how convolutional Keplerian motion may manifest itself in the line profile, thereby making it easier to identify and extract the underlying Keplerian velocities.

In Chapter 5, we built benchmark datasets of radial velocities for emission-line stars. The algorithm presented in this chapter represents a new baseline against which investigators should confront their models. These benchmark datasets are freely available for such testing. In the future we encourage the continued development of new and ingenious algorithms to improve upon our scores on the synthetic data, and thereby help the field make increasingly better estimates of the orbits and dynamical masses of emission-line binaries. Possible performance enhancements may be found through utilising different samplers, Gaussian process kernels, mean models, parameter combinations, and other algorithms entirely. In addition, accompanying radial velocity data from companion stars less affected by systematics, and well-sampled astrometry is paramount in driving down parameter uncertainties when the emission-line star is heavily effected by systematics.

In our models in Chapters 4 and 6, we assume, for the most part, that the winds are spherically symmetric. However, we know from observations (eg. Smith et al. 2003) and theory (eg. Cranmer & Owocki 1995; Owocki et al. 1996) that we may expect latitudinally-dependent winds from massive rotating stars. Specific to our geometric modelling in Chapter 6, in the future it may be possible to make quantitative insights into the latitudinal dependence of the winds of Eta Carinae. But, in order to do this the inclination must be constrained with greater certainty. This may involve

improving on the work of Madura et al. (2012) and Teodoro et al. (2016) to beat down the uncertainties in their models, or utilising observations of Eta Carinae’s surrounding polar nebula to infer the inclination (Smith 2006). However, this requires assumptions about the alignment of the orbital plane with respect to the nebula. More generally, further research is needed into quantifying the latitudinal dependence of stellar winds through high-quality observations and models.

8.2 Open-source software

All of the work presented in this thesis stands on the shoulders of a flourishing open-source community of developers, both in and outside of astrophysics. We are grateful to the immense work that has been dedicated to creating tools that enable research to advance at a faster pace than would otherwise be possible. We acknowledge the use of the following software: **CMFGEN** (Hillier & Miller 1998; Hillier & Lanz 2001), **SciPy** (Jones et al. 2001), **NumPy** (Van Der Walt et al. 2011), **Pandas** (McKinney 2010), **emcee** (Foreman-Mackey et al. 2013), **george** (Ambikasaran et al. 2015), **barycorrpy** (Kanodia & Wright 2018), **pytorch** (Paszke et al. 2019), **astropy** (The Astropy Collaboration et al. 2013; 2018), **photutils** (Bradley et al. 2020), **starry** (Luger et al. 2019), and **Matplotlib** (Hunter 2007). Software developed for this thesis is open-source and freely available at <https://github.com/DavoGrant/>. We hope the algorithms and datasets presented will lead to the next generation of ideas, and incrementally we, as a community, will make better and better inferences of the orbits and outflows of emission-line stars, and advance our understanding of massive stars as a whole.

Bibliography

- Abbott D. C., 1982, *The Astrophysical Journal*, 259, 282
- Abbott D. C., Lucy L. B., 1985, *The Astrophysical Journal*, 288, 679
- Aigrain S., Parviainen H., Pope B. J., 2016, *Monthly Notices of the Royal Astronomical Society*, 459, 2408
- Aitken R. G., 1964, *The Binary Stars*. Dover Publication, New York
- Aksulu M. D., Wijers R. A., Van Eerten H. J., Van Der Horst A. J., 2020, *Monthly Notices of the Royal Astronomical Society*, 497, 4672
- Alfonso-Garzón J., Domingo A., Mas-Hesse J. M., Giménez A., 2012, *Astronomy and Astrophysics*, 548, 79
- Allen D., Hillier D., 1993, *Publications of the Astronomical Society of Australia*, 10, 338
- Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O’Neil M., 2015, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38, 252
- Axelrod T., Kantor J., Lupton R. H., Pierfederici F., 2010, in *Software and Cyberinfrastructure for Astronomy*. International Society for Optics and Photonics, p. 774015
- Bailer-Jones C. A. L., Rybizki J., Fouesneau M., Mantelet G., Andrae R., 2018, *The Astronomical Journal*, 156, 58
- Beals C. S., 1929, *Monthly Notices of the Royal Astronomical Society*, 90, 202
- Belczynski K., Kalogera V., Bulik T., 2002, *The Astrophysical Journal*, 572, 407

- Bell A. R., 1978a, Monthly Notices of the Royal Astronomical Society, 182, 147
- Bell A. R., 1978b, Monthly Notices of the Royal Astronomical Society, 182, 443
- Bjorkman J. E., Cassinelli J. P., 1993, The Astrophysical Journal, 409, 429
- Blaauw A., 1961, Bulletin of the Astronomical Institutes of the Netherlands, 15, 265
- Blundell K. M., Bowler M. G., Schmidtbreick L., 2007, Astronomy & Astrophysics, 474, 903
- Blundell K. M., Bowler M. G., Schmidtbreick L., 2008, The Astrophysical Journal, 678, L47
- Blundell K. M., Schmidtbreick L., Trushkin S., 2011, Monthly Notices of the Royal Astronomical Society, 417, 2401
- Bowen I. S., 1928, The Astrophysical Journal, 67
- Bradley L., et al., 2020, Zenodo, 1.0.0
- Branch M. A., Coleman T. F., Li Y., 1999, SIAM Journal on Scientific Computing, 21, 1
- Brewer B. J., Stello D., 2009, Monthly Notices of the Royal Astronomical Society, 395, 2226
- Burke D. L., et al., 2017, The Astronomical Journal, 155, 41
- Burki G., Rufener F., Burnet M., Richard C., Blecha A., Bratschi P., 1995, Astronomy & Astrophysics Supplement Series, 112, 383
- Canto J., Raga A. C., Wilkin F. P., 1996, The Astrophysical Journal, 469, 729
- Castor J. I., 1970, Monthly Notices of the Royal Astronomical Society, 149, 111
- Castor J. I., 1974, Monthly Notices of the Royal Astronomical Society, 169, 279
- Castor J., Abbot D., Klein R., 1975, The Astrophysical Journal, 195, 157
- Cherepashchuk A. M., 1976, Soviet Astronomy Letters, 2, 138

- Chlebowski T., Garmany C. D., 1991, *The Astrophysical Journal*, 368, 241
- Clementel N., Madura T. I., Kruip C. J. H., Paardekooper J.-P., Gull T. R., 2015a, *Monthly Notices of the Royal Astronomical Society*, 447, 2445
- Clementel N., Madura T. I., Kruip C. J., Paardekooper J. P., 2015b, *Monthly Notices of the Royal Astronomical Society*, 450, 1388
- Conti P. S., 1976, *Ministere de l'Education National*, 9, 193
- Corcoran M. F., Ishibashi K., Swank J. H., Petre R., 2001, *The Astrophysical Journal*, 547, 1034
- Cranmer S. R., Owocki S. P., 1995, *The Astrophysical Journal*, 440, 308
- Crowther P. A., 2007, *Annual Review of Astronomy and Astrophysics*, 45, 177
- Currie D. G., et al., 1996, *The Astronomical Journal*, 112, 1115
- Damineli A., Conti P. S., Lopesb D. F., Van Den Heuvel E. P. J., 1997, *New Astronomy*, 2, 107
- Damineli A., et al., 2008a, *Monthly Notices of the Royal Astronomical Society*, 384, 1649
- Damineli A., et al., 2008b, *Monthly Notices of the Royal Astronomical Society*, 386, 2330
- Damineli A., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 963
- Dark Energy Survey Collaboration et al., 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 1270
- Davidson K., 1987, *The Astrophysical Journal*, 317, 760
- Davidson K., 1997, *New Astronomy*, 2, 387
- Davidson K., Humphreys R. M., 1997, *Annual Review of Astronomy and Astrophysics*, 35, 1

- Davidson K., Ishibashi K., Gull T. R., Humphreys R. M., Smith N., 2000, The Astrophysical Journal, 530, 107
- Davidson K., Mehner A., Humphreys R. M., Martin J. C., Ishibashi K., 2015, Astrophysical Journal Letters, 801, L15
- De Koter A., Lamers H. J., Schmutz W., 1996, Astronomy and Astrophysics, 306, 501
- Dorland B. N., Currie D. G., Hajian A. R., 2004, The Astronomical Journal, 127, 1052
- Dupuy T. J., Liu M. C., 2017, The Astrophysical Journal Supplement Series, 231, 15
- Ebbets D., 1982, The Astrophysical Journal Supplement Series, 48, 399
- Ekström S., et al., 2012, Astronomy and Astrophysics, 537
- Evans C. J., Taylor W. D., 2011, Astronomy and Astrophysics, 530
- Fahed R., et al., 2011, Monthly Notices of the Royal Astronomical Society, 418, 2
- Feinstein A. D., et al., 2019, Publications of the Astronomical Society of the Pacific, 131
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306
- GRAVITY Collaboration et al., 2018, Astronomy & Astrophysics, 618, A125
- Gaia Collaboration 2018, VizieR Online Data Catalog, p. I/345
- Gaia Collaboration et al., 2016, Astronomy and Astrophysics, 595, 14
- Gaia Collaboration Brown A. G. A., Vallenari A., Prusti T., de Bruijne J. H. J., Babusiaux C., Bailer-Jones C. A. L., 2018, J. Guiraud 16 , R. Gutiérrez-Sánchez, 60, 61
- Gaia Collaboration et al., 2021, Astronomy and Astrophysics, 649, A1
- Gayley K. G., 2000, The Astrophysical Journal, 529, 1019

- Gayley K. G., 2009, *Astrophysical Journal*, 703, 89
- Gayley K. G., Owocki S. P., Cranmer S. R., 1997, *The Astrophysical Journal*, 475, 786
- Gibson N. P., Aigrain S., Roberts S., Evans T. M., Osborne M., Pont F., 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 2683
- Gosset E., Hutsemekers D., Surdej J., Swings J., 1985, *Astronomy & Astrophysics*, 153, 71
- Grant D., Blundell K., 2021, *Monthly Notices of the Royal Astronomical Society*, 509, 367
- Grant D., Blundell K., Matthews J., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 17
- Griffiths D. J., 2007, *Introduction to Electrodynamics*. Pearson Education, Dorling Kindersley
- Groh J., Damineli A., 2004, *Information Bulletin on Variable Stars*, 5492
- Groh J. H., Hillier D. J., Madura T. I., Weigelt G., 2012a, *Monthly Notices of the Royal Astronomical Society*, 423, 1623
- Groh J. H., Madura T. I., Hillier D. J., Kruip C. J., Weigelt G., 2012b, *Astrophysical Journal Letters*, 759, 2
- Hamaguchi K., et al., 2014, *Astrophysical Journal*, 784, 125
- Hamaguchi K., et al., 2018, *Nature Astronomy*, 2, 731
- Hayes D. S., Latham D. W., 1975, *The Astrophysical Journal*, 197, 593
- Heinze A. N., et al., 2018, *The Astronomical Journal*, 156, 241
- Hernandez C. A., Sahade J., Lopez L., Thackeray A. D., 1981, *Publications of the Astronomical Society of the Pacific*, 93, 747
- Hillier D. J., 1987, *The Astrophysical Journal Supplement Series*, 9, 965

- Hillier D. J., 2011, in Proceedings of the International Astronomical Union. pp 229–234
- Hillier D. J., Lanz T., 2001, ASP Conference Series, 247, 343
- Hillier D. J., Miller D. L., 1998, Astrophysical Journal, 496, 407
- Hillier D. J., Davidson K., Ishibashi K., Gull T., 2001, The Astrophysical Journal, 553, 837
- Hillier D. J., et al., 2006, The Astrophysical Journal, 642, 1098
- Hills J. G., 1983, The Astrophysical Journal, 267, 322
- Hirai R., Mandel I., 2021, Publications of the Astronomical Society of Australia
- Hirai R., Podsiadlowski P., Owocki S. P., Schneider F. R. N., Smith N., 2021, Monthly Notices of the Royal Astronomical Society, 503, 4276
- Howell S. B., et al., 2014, Publications of the Astronomical Society of the Pacific, 126, 398
- Humphreys R. M., Davidson K., 1994, Publications of the Astronomical Society of the Pacific, 106, 1025
- Humphreys R. M., Smith N., Davidson K., 1999, Publications of the Astronomical Society of the Pacific, 111, 1124
- Humphreys R. M., Davidson K., Koppelman M., 2008, Astronomical Journal, 135, 1249
- Hunter J. D., 2007, Computing in Science and Engineering, 9, 99
- Jenkins J. M., et al., 2016, in Software and Cyberinfrastructure for Astronomy IV.
- Jones E., Oliphant T., Peterson P., 2001, SciPy: Open source scientific tools for Python
- Kanodia S., Wright J., 2018, Research Notes of the AAS, 2, 4

- Kaper L., Henrichs H. F., Nichols J. S., 1996, *Astronomy and Astrophysics Supplement Series*, 116, 257
- Kashi A., Soker N., 2016, *The Astrophysical Journal*, 825, 105
- Kim D. W., Protopapas P., Alcock C., Byun Y. I., Bianco F. B., 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 558
- Kingma D. P., Ba J. L., 2015, in *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*.
- Kogure T., Leung K., 2010, *The astrophysics of emission-line stars*. Springer Science & Business Media.
- Kolmogorov A. N., 1933, *Giornale dell’Istituto Italiano degli Attuari*, 4, 83
- Konacki M., 2005, *The Astrophysical Journal*, 626, 431
- Konacki M., Muterspaugh M. W., Kulkarni S. R., Hellminiak K. G., 2010, *Astrophysical Journal*, 719, 1293
- Kovács G., Bakos G., Noyes R. W., 2005, *A trend filtering algorithm for wide-field variability surveys*
- Kruip C., 2011, PhD thesis, Univ. Lieden
- Kudritzki R.-p., Puls J., 2000, *Annual Reviews of Astronomy & Astrophysics*, 38, 613
- Kuhi L. V., 1973, *The Astrophysical Journal*, 180, 783
- LSST Science Collaboration et al., 2009, Technical report, *The LSST Science Book*. arXiv:0912.0201
- Lang D., Hogg D. W., Mierle K., Blanton M., Roweis S., 2010, *Astronomical Journal*, 139, 1782
- Lattimer A. S., Cranmer S. R., 2021, *The Astrophysical Journal*, 910, 48
- Leitherer C., 1988, *The Astrophysical Journal*, 326, 356

- Lucy L. B., 1970, *The Astrophysical Journal*, 159, 879
- Luger R., Agol E., Kruse E., Barnes R., Becker A., Foreman-Mackey D., Deming D., 2016, *The Astronomical Journal*, 152, 100
- Luger R., Agol E., Foreman-Mackey D., Fleming D. P., Lustig-Yaeger J., Deitrick R., 2019, *The Astronomical Journal*, 157, 64
- Luo D., McCray R., Mac Low M.-M., 1990, *The Astrophysical Journal*, 362, 267
- MacGregor K. B., Hartmann L., Raymond J. C., 1979, *The Astrophysical Journal*, 231, 514
- Madura T. I., Gull T. R., Owocki S. P., Groh J. H., Okazaki A. T., Russell C. M., 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 2064
- Madura T. I., et al., 2013, *Monthly Notices of the Royal Astronomical Society*, 436, 3820
- Maeder A., Meynet G., 2010, *New Astronomy Reviews*, 54, 32
- Marchenko S. V., Moffat A. F. J., Koenigsberger G., 1994, *The Astrophysical Journal*, 422, 810
- Marchenko S. V., et al., 2003, *The Astrophysical Journal*, 596, 1295
- Marchiano P., Brandi E., Muratore M. F., Quiroga C., Ferrer O. E., García L. G., 2012, *Astronomy and Astrophysics*, 540, 91
- Martínez Fiorenzano A. F., Gratton R. G., Desidera S., Cosentino R., Endl M., 2005, *Astronomy & Astrophysics*, 442, 775
- Mason B. D., Hartkopf W. I., Gies D. R., Henry T. J., Helsel J. W., 2009, *Astronomical Journal*, 137, 3358
- McGregor P. J., Hyland A. R., Hillier D. J., 1988, *The Astrophysical Journal*, 324, 1071
- McKinney W., 2010, in *Proceedings of the 9th Python in Science Conference*. pp 51–56

- Mehner A., Davidson K., Ferland G. J., Humphreys R. M., 2010, *Astrophysical Journal*, 710, 729
- Mehner A., Davidson K., Martin J. C., Humphreys R. M., Ishibashi K., Ferland G. J., 2011, *Astrophysical Journal*, 740, 80
- Mehner A., Davidson K., Humphreys R. M., Ishibashi K., Martin J. C., Ruiz M. T., Walter F. M., 2012, *Astrophysical Journal*, 751, 73
- Mislis D., Pyrzas S., Alsubai K. A., Tsvetanov Z. I., Vilchez N. P., 2017, *Monthly Notices of the Royal Astronomical Society*, 465, 3759
- Moe M., Di Stefano R., 2017, *The Astrophysical Journal Supplement Series*, 230, 15
- Moffat A. F. J., Niemela V. S., Phillips M. M., Chu Y.-H., Seggewiss W., 1987, *The Astrophysical Journal*, 312, 612
- Mokiem M. R., et al., 2007, *Astronomy and Astrophysics*, 473, 603
- Monnier J. D., et al., 2011, *Astrophysical Journal Letters*, 742, 1
- Moré J. J., 1978, in , *Numerical analysis*. Springer, Berlin, Heidelberg, pp 105–116
- Morris P. W., Gull T. R., Hillier D. J., Barlow M. J., Royer P., Nielsen K., Black J., Swinyard B., 2017, *The Astrophysical Journal*, 842, 79
- Murray C. D., Dermott S. F., 2000, *Solar System Dynamics*. Cambridge University Press
- Najarro F., 2001, in *P Cygni 2000: 400 Years of Progress*. p. 133
- Newton I., 1687, London
- Nielsen K. E., Corcoran M. F., Gull T. R., Hillier D. J., Hamaguchi K., Ivarsson S., Lindler D. J., 2007, *The Astrophysical Journal*, 660, 669
- Okazaki A. T., Owocki S. P., Russell C. M., Corcoran M. F., 2008, *Monthly Notices of the Royal Astronomical Society: Letters*, 388, 39
- Oskinova L. M., 2016, *Advances in Space Research*, 58, 739

- Oskinova L. M., Feldmeier A., Hamann W. R., 2006, *Monthly Notices of the Royal Astronomical Society*, 372, 313
- Owocki S. P., Castor J. I., Rybicki G. B., 1988, *The Astrophysical Journal*, 335, 914
- Owocki S. P., Cranmer S. R., Gayley K. G., 1996, *The Astrophysical Journal*, 472, L115
- Owocki S. P., Gayley K. G., Shaviv N. J., 2004, *The Astrophysical Journal*, 616, 525
- Panagia N., Felli M., 1975, *Astronomy and Astrophysics*, 39, 1
- Parkin E. R., Pittard J. M., Corcoran M. F., Hamaguchi K., Stevens I. R., 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 1758
- Parkin E. R., Pittard J. M., Corcoran M. F., Hamaguchi K., 2011, *Astrophysical Journal*, 726, 105
- Paszke A., et al., 2019, in *Advances in Neural Information Processing Systems*.
- Pauldrach A., Puls J., Kudritzki R. P., 1986, *Astronomy and astrophysics* (Berlin. Print), 164, 86
- Pietrzyński G., et al., 2013, *Nature*, 495, 76
- Pittard J. M., Corcoran M. F., 2002, *Astronomy & Astrophysics*, 383, 636
- Pittard J. M., Stevens I. R., Corcoran M. F., Ishibashi K., 1998, *Monthly Notices of the Royal Astronomical Society*, 299, L5
- Podsiadlowski P., Joss P. C., Hsu J. J. L., 1992, *The Astrophysical Journal*, 391, 246
- Pollock A., 1987, *Astronomy and astrophysics*, 171, 135
- Porter A., Grant D., Blundell K., Lee S., 2021a, *Monthly Notices of the Royal Astronomical Society*, 501, 5554
- Porter A., Blundell K., Podsiadlowski P., Lee S., 2021b, *Monthly Notices of the Royal Astronomical Society*, 503, 4802

- Porter A., Blundell K., Lee S., 2021c, *Monthly Notices of the Royal Astronomical Society*, 509, 1720
- Prilutskii O. F., Usov V. V., 1976, *Soviet Astronomy*, 20, 2
- Puls J., Vink J. S., Najarro F., 2008, *Astronomy and Astrophysics Review*, 16, 209
- Quiñonero-Candela J., Rasmussen C. E., Sinz F., Bousquet O., Schölkopf B., 2005, Technical report, LNAI 3944 - Evaluating Predictive Uncertainty Challenge. MLCW
- Rasmussen C. E., Williams C. K. I., 2006, *Gaussian processes for machine learning*
- Richardson N. D., Gies D. R., Henry T. J., Fernández-Lajés E., Okazaki A. T., 2010, *Astronomical Journal*, 139, 1534
- Richardson N. D., Gies D. R., Gull T. R., Moffat A. F., St-Jean L., 2015, *Astronomical Journal*, 150
- Richardson N. D., et al., 2021, *The Astrophysical Journal*, 908, L3
- Riener M., Kainulainen J., Henshaw J. D., Orkisz J. H., Murray C. E., Beuther H., 2019, *Astronomy & Astrophysics*, 628, A78
- Rivet J. P., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 218
- Ruiz M. T., Melnick J., Ortiz P., 1984, *The Astrophysical Journal*, 285, L19
- Russell C. M. P., 2013, PhD thesis, University of Delaware
- Russell C. M., Corcoran M. F., Hamaguchi K., Madura T. I., Owocki S. P., John Hillier D., 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 2275
- Saar S. H., Butler R. P., Marcy G. W., 1998, *The Astrophysical Journal*, 498, L153
- Sana H., et al., 2012, *Science*, 337, 444
- Sana H., et al., 2014, *Astrophysical Journal, Supplement Series*, 215, 15

- Sander A. A., Hamann W. R., Todt H., Hainich R., Shenar T., 2017, *Astronomy and Astrophysics*, 603, 1
- Savitzky A., Golay M., 1964, *Analytical Chemistry*, 36, 1627
- Schwarz G., 1978, *The Annals of Statistics*, 6, 461
- Scott D. W., 2015, *Multivariate density estimation: Theory, practice, and visualization: Second edition.* wiley
- Shenar T., et al., 2019, *Astronomy & Astrophysics*, 627, 151
- Shylaja B. S., 1986, *Journal of Astrophysics and Astronomy*, 7, 171
- Shylaja B. S., 1987, *Journal of Astrophysics and Astronomy*, 8, 183
- Silverman B. W., 2018, *Density estimation: For statistics and data analysis.* Routledge
- Simpson F., Lalchand V., Rasmussen C. E., 2020, arXiv preprint
- Smirnov N. V., 1939, *Bull. Math. Univ. Moscou*, 2, 3
- Smith N., 2004, *Monthly Notices of the Royal Astronomical Society*, 351, 1
- Smith N., 2006, *The Astrophysical Journal*, 644, 1151
- Smith N., 2014, *Annual Review of Astronomy and Astrophysics*, 52, 487
- Smith N., Frew D. J., 2011, *Monthly Notices of the Royal Astronomical Society*, 415, 2009
- Smith N., Davidson K., Gull T. R., Ishibashi K., Hillier D. J., 2003, *The Astrophysical Journal*, 586, 432
- Smith J. C., et al., 2012, *Publications of the Astronomical Society of the Pacific*, 124, 1000
- Smith N., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 1466
- Steiner J. E., Damini A., 2004, *The Astrophysical Journal*, 612, 133

- Stevens I. R., 1993, *The Astrophysical Journal*, 404, 281
- Stevens I. R., Pollock A. M. T., 1994, *Monthly Notices of the Royal Astronomical Society*, 269, 226
- Stevens I. R., Blondin J. M., Pollock A. M. T., 1992, *The Astrophysical Journal*, 386, 265
- Stickland D. J., Bromage G. E., Budding E., Burton W. M., Willis A. J., 1984, *Astronomy & Astrophysics*, 134, 37
- Stumpe M. C., et al., 2012, *Publications of the Astronomical Society of the Pacific*, 124, 985
- Stumpe M. C., Smith J. C., Catanzarite J. H., Van Cleve J. E., Jenkins J. M., Twicken J. D., Girouard F. R., 2014, *Publications of the Astronomical Society of the Pacific*, 126, 100
- Szostek A., Dubus G., Mcswain M. V., 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 3521
- Tamuz O., Mazeh T., Zucker S., 2005, *Monthly Notices of the Royal Astronomical Society*, 356, 1466
- Teodoro M., et al., 2012, *Astrophysical Journal*, 746, 73
- Teodoro M., et al., 2016, *The Astrophysical Journal*, 819, 47
- The Astropy Collaboration et al., 2013, *Astronomy and Astrophysics*, 558, 33
- The Astropy Collaboration et al., 2018, *The Astronomical Journal*, 156, 123
- Thomas J. D., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 5221
- Tonry J., Davis M., 1979, *The Astronomical Journal*, 84, 1511
- Underhill A. B., Hill G. M., 1994, *The Astrophysical Journal*, 432, 770

- Van Der Walt S., Colbert S. C., Varoquaux G., 2011, *Computing in Science and Engineering*, 13, 22
- Van Dyk S. D., Matheson T., 2012, *Astrophysical Journal*, 746, 179
- Verner E. M., Gull T. R., Bruhweiler F., Johansson S., Ishibashi K., Davidson K., 2002, *The Astrophysical Journal*, 581, 1154
- Verner E., Bruhweiler F., Gull T., 2005, *The Astrophysical Journal*, 624, 973
- Vink J. S., De Koter A., Lamers H. J., 1999, *Astronomy and Astrophysics*, 350, 181
- Vink J. S., De Koter A., Lamers H. J., 2000, *Astronomy and Astrophysics*, 362, 295
- Vink J., de Koter A., Lamers H., 2001, *Astronomy & Astrophysics*, 369, 574
- Walborn N. R., 1973, *The Astrophysical Journal*, 179, 517
- Walborn N. R., Liller M. H., 1977, *The Astrophysical Journal*, 211, 181
- Walder R., 1995, in Hucht K. A. v. d., Williams P., eds, *WR stars: binaries, colliding winds, evolution.. IAU Symposium No. 163*, pp 420–424
- Weigelt G., Ebersberger J., 1986, *Astronomy & Astrophysics*, 163, L5
- Wenger M., et al., 2000, *Astronomy and Astrophysics Supplement Series*, 143, 9
- White R. L., Chen W., 1994, *Astrophysics and Space Science*, 221, 295
- Williams P. M., 2008, in *Revista Mexicana de Astronomia y Astrofisica: Serie de Conferencias*. pp 71–76
- Williams P. M., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 503, 643
- Willis A. J., 1982, *Monthly Notices of the Royal Astronomical Society*, 198, 897
- Wright A. E., Barlow M. J., 1975, *Monthly Notices of the Royal Astronomical Society*, 170, 41

Wu Y.-L., Smith N., Close L. M., Males J. R., Morzinski K. M., 2017, The Astrophysical Journal Letters, 841, L7

de Vaucouleurs G., Eggen O. J., 1952, Publications of the Astronomical Society of the Pacific, 64, 185