Doctoral Thesis

The search and study of extreme neutron stars

Author:
James Stringer

Supervisor:
Prof. Rene Breton

A thesis submitted to the University of Manchester
for the degree of Doctor of Philosophy

in the

Department of Physics and Astronomy, School of Natural Sciences
Faculty of Science and Engineering

2021

Contents

List of Figures

List of Tables

List of Abbreviations

AP	**A**ccretion-**P**owered
ASAS-SN	**A**ll-**S**ky **A**utomated **S**urvey for **S**uper**N**ovae
BIC	**B**ayesian **I**nformation **C**riterion
CCD	**C**harge **C**oupled **D**evice
CSS	**C**atalina **S**ky **S**urvey
CV	**C**ataclysmic **V**ariable
DM	**D**istance **M**odulus OR **D**ispersion **M**easure
FAP	**F**alse **A**larm **P**robability
FWHM	**F**ull **W**idth at **H**alf **M**aximum
GEV	**G**eneralised **E**xtreme **V**alue
GLS	**G**eneralised **L**omb-**S**cargle
GOTO	**G**ravitational-wave **O**ptical **T**ransient **O**bservatory
HMXB	**H**igh **M**ass **X**-ray **B**inary
IBS	**I**ntra-**B**inary **S**hock
LMXB	**L**ow **M**ass **X**-ray **B**inary
LSST	**L**egacy **S**urvey of **S**pace and **T**ime
MCMC	**M**arkov **C**hain **M**onte **C**arlo
MHAOV	**M**ulti-**H**armonic **A**nalysis **O**f **V**ariance
ML	**M**achine **L**earning
MSP	**M**illi**S**econd **P**ulsar
NTT	**N**ew **T**echnology **T**elescope
PDF	**P**robability **D**ensity **F**unction
PTF	**P**alomar **T**ransient **F**actory
PWN	**P**ulsar **W**ind **N**ebula
RL	**R**oche **L**obe
RP	**R**otation-**P**owered
RV	**R**adial **V**elocity
SDSS	**S**loan **D**igital **S**ky **S**urvey
SNR	**S**ignal-to-**N**oise **R**atio
SSB	**S**olar **S**ystem **B**arycentric

tMSP	**t**ransitional **M**illi**S**econd **P**ulsar
UT	**U**nit **T**elescope
ZTF	**Z**wicky **T**ransient **F**acility

THE UNIVERSITY OF MANCHESTER

Abstract

Faculty of Science and Engineering
Department of Physics and Astronomy, School of Natural Sciences

Doctor of Philosophy

The search and study of extreme neutron stars

by James Stringer

Millisecond pulsars – highly magnetised neutron stars rotating with millisecond periods – are usually observed in binary systems as they are 'spun up' to such fast periods by the accretion of matter from their companion. The 'spider' class of binary contains a millisecond pulsar and a low mass, degenerate companion in a tight orbit which is irradiated by the pulsar wind, resulting in distinct observable properties across the whole electromagnetic spectrum. The first part of this research presents new optical photometry of two systems, PSR J1023+0038 and PSR J1227–4853, which display transitions between the aforementioned accretion phase and the spider phase. We find that through modelling these data we gain insights into the nature of this transition. The second concerns the search for new spider systems with the creation of a classification pipeline for the new Gravitational-wave Optical Transient Observatory telescope project, employing periodicity search and machine learning classification techniques. We discuss the application of this pipeline to commissioning data from this fledgling project, and how the shortfalls we encountered can be overcome.

Declaration of Authorship

I, James Stringer, declare that this thesis titled, "The search and study of extreme neutron stars" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Copyright Statement

Acknowledgements

This thesis would not have been possible without the advice, encouragement, and patience of a great number of people, many more than I can acknowledge here. First and foremost, an enormous thank you to my supervisor, Prof. Rene Breton, for his continued and committed support and for helping me to remain enthusiastic about my research during tough times. I am incredibly grateful also for the team spirit and assistance of Drs. Colin Clark, Mark Kennedy, Daniel Mata Sanchez, and Guillaume Voisin, throughout my PhD. I am immensely glad to have had the pleasure of working with them. I also wish to extend a general thank you to the pulsar and time domain group, and indeed the JBCA as a whole, for putting up with the problem child that is myself, but mostly for the countless times I've come for advice, help, and company. Thank you to Josh, Max, Luke, Nialh, Dan, Emma, and the countless other friends I've made along the way. Last but not least, to Sara for standing by me through all the highs and lows.

I acknowledge support of the European Research Council, under the European Union's Horizon 2020 research and innovation program (grant agreement No. 715051; Spiders). This work has made use of data from the European Space Agency (ESA) mission *Gaia* (`https://www.cosmos.esa.int/gaia`), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, `https://www.cosmos.esa.int/web/gaia/dpac/consortium`). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

Chapter 1

Introduction

1.1 Overview of this work

The study of spider pulsars is a relatively new field, even compared to pulsar astronomy; it is only within the last decade that it has evolved from the study of a few outlier systems to the diverse field it is today. As such spiders, consisting of a millisecond pulsar (MSP) and a low-mass ($< 1\mathrm{M}_{\odot}$) companion star in a tight orbit ($P_{\mathrm{orb}} \sim 0.1 - 1$ day), are still a small population (< 100 known) with a plethora of open questions about their evolution, formation, and the physical processes responsible for their characteristic observable features. The tidally-locked companion in these systems is strongly irradiated by the pulsar wind, resulting in heating of the inner face and ablation of stellar plasma (Fruchter et al., 1988). Spider systems host some of the most massive and fastest spinning MSPs (Linares, 2020), and can display unique properties; a key focus of this research is the transitional MSP (tMSP). tMSP systems are observed to transition between two distinct binary classes on very short timescales of a few weeks to months, displaying the complete accretion timescale during each phase of the transition. As such they are excellent laboratories for accretion physics, and can provide insights on the formation of spider systems. This work aims to tackle two key questions; how can we best search for new systems, and how and why do these tMSPs transition between their two states?

First, we introduce the field of pulsar astronomy and overview the state of research of spider pulsars. Pulsar astronomy is an extremely diverse field and as such we keep the discussion of, for example, the radio timing of pulsars, to a minimum. Instead, we focus on millisecond pulsars in binary systems and their optical properties. Following this, we overview some of the key statistical and computational tools used, such as the `Icarus` binary light curve synthesis code.

The next chapter deals with three important facets of optical astronomy. First, we present the instrumentation used in this research, including the optical telescopes and photometry instruments, as well as an overview of the data reduction and calibration process. Then, we discuss some of the observation work performed during the research which otherwise does not belong in its own chapter. This includes the presentation of selected optical light curves. Lastly, we introduce some of the light curve simulation techniques used. These simulations form the backbone of chapter 4, however in this chapter we also discuss some preliminary results that were obtained.

The following two chapters constitute the bulk of the research, and attempt to answer the two key questions. The first, chapter 3, covers our analysis of optical photometry of two transitional millisecond pulsars. In this chapter we discuss the modelling of the light curves of the sources using the `Icarus` code and how the results can be best interpreted. These sources presented significant challenges in modelling due to the highly asymmetric nature of their light curves, something which is unexpected with standard heating models.

The second extends the simulation work of chapter 2 in the context of a new optical survey telescope, GOTO. We created a pipeline to identify periodic sources and classify them by star type, with the aim of identifying the optical counterparts of new spider systems. In this chapter we introduce the periodicity search pipeline and the machine learning classifier, and discuss its performance with both simulated and real photometry. This work reveals some significant shortfalls of the GOTO photometry in its current state, as revealed by the poor performance of the periodicity search and classifier.

In the final chapter we discuss the consequences of the research in these two chapters: what lessons can be drawn from our modelling and attempts at light curve classification? In the first case, we obtained excellent photometry but our modelling did not easily reveal concrete results. Conversely, our periodicity search pipeline performed well on simulated data, however our results with real GOTO data revealed significant problems. we discuss the causes of these problems, and what lessons we can learn from this research to improve results in the future. For instance, what can we learn from other survey telescopes that can be applied to GOTO, or what has our extensive modelling taught us about the intricacies of the `Icarus` model?

1.2 Introduction to pulsars

1.2.1 Discovery, population, and properties

Emerging in the mid 1960s with the discovery of the first radio pulsar (Hewish et al., 1968), pulsar astronomy (along with neutron star physics) is one of the most active research areas in astrophysics today. While it is now known that pulsars emit radiation along the axis of their co-rotating magnetic fields, this early discovery of regularly spaced radio pulses of period 1.337 s, was initially attributed to a variety of terrestrial sources as well as a pulsating compact object. Subsequent parallax measurements indicated that the source was extrasolar, and Gold (1968) confirmed that the radio pulsations were indeed those emitted by a rotating neutron star with spin period P. Notably, this discovery was in fact preceded by theories from Pacini (1967) and Shklovsky (1967) predicting the emission of radio-frequency radiation from rotating neutron stars, and the emission of X-rays from accreting neutron stars respectively. The physics of this emission is covered in section 1.2.2. Later, the discovery of the Crab and Vela pulsars in 1968 confirmed the theory that neutron stars were born in supernovae (SNe).

In the 50 years since this first discovery, there have been over 3000 new pulsars identified (as of November 2020) with more discoveries every week (Manchester, 2014). A typical pulse profile is illustrated in figure 1.1. This large number of sources gives a great deal of insight into the common properties of the pulsar population, including the distribution of neutron star masses, radii, spin periods, and ages, as well as canonical values for these properties. As well as this, the theoretical backbone to pulsar astronomy has also matured, such that there are now several accepted theories for many of these observed phenomena. These will be discussed in section 1.2.2.

Neutron stars are composed of matter at densities beyond the nuclear density, meaning that standard theories of stellar structure are not sufficient (Lattimer & Prakash, 2004). This makes it hard to easily determine neutron star masses and radii. The reverse is also true; the composition of neutron star interiors is challenging to determine from astrophysical observables (Lattimer & Prakash, 2007). It is possible that gravitational wave measurements of neutron star binary mergers may shed some light on the structure (Hinderer et al., 2010). However, from observations and timing of binary pulsar systems (Stairs, 2004), and theoretical predictions from neutron star formation (Lattimer & Prakash, 2001), the neutron star canonical mass is accepted to be $M_{\rm can} = 1.4\ \mathrm{M}_\odot$. The observational studies of Stairs (2004) confirm this, though varying formation routes can result in different masses. However, accretion and pulsar recycling, discussed in section 1.3, often lead to more

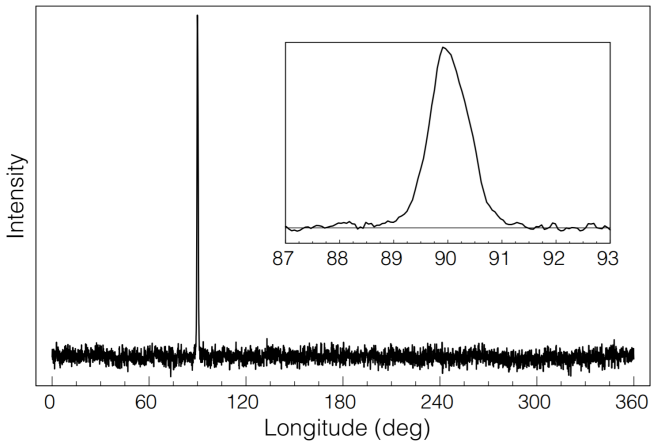

FIGURE 1.1: Pulse profile of PSR J2144–3933, a pulsar with an unusually long 8.5 second period. The main plot shows the radio pulse with a reference of one orbit, while the inset shows a close-up of the pulse. The shape of the pulse is defined by the shape of the radio-frequency emission region on the pulsar. Image credit: Young et al. (1999)

massive pulsars (Lattimer & Prakash, 2001; van Kerkwijk et al., 2011). Similarly, the canonical neutron star radius is taken to be $R_{\rm can} \sim 10-12$ km. Special relativity provides a lower radius limit by requiring the sound speed be less than the speed of light (Glendenning, 1992), and centrifugal force provides an upper radius limit by requiring the neutron star remain gravitationally bound. Lattimer & Prakash (2001) summarise that this canonical radius is agreed on by the majority of models. It is worth noting that this is only 3 times the Schwarzchild radius of a black hole of comparable mass, quantifying the magnitude of the densities and gravitational fields of these objects (Lorimer & Kramer, 2004).

As well as high densities, a characteristic common to all pulsars is an incredibly high magnetic field strength with surface field densities on the order of $B \sim 10^{12} - 10^{14}$ G, suggested by X-ray measurements of electron cyclotron emission lines (Truemper et al., 1978; Wheaton et al., 1979). These fields arise as a result of the conservation of magnetic flux when the progenitor star collapses through the compression of field lines as it rapidly decreases in radius. The magnetic field is dipolar and offset from the rotation axis of the pulsar, and as is described in section 1.2.2, this is responsible for the radio emission. The commonly assumed model is the so-called lighthouse model, illustrated in figure 1.2, where the magnetosphere co-rotates with the pulsar.

The fastest spinning pulsar to date, PSR J1748–2446ad (Hessels et al., 2006), has

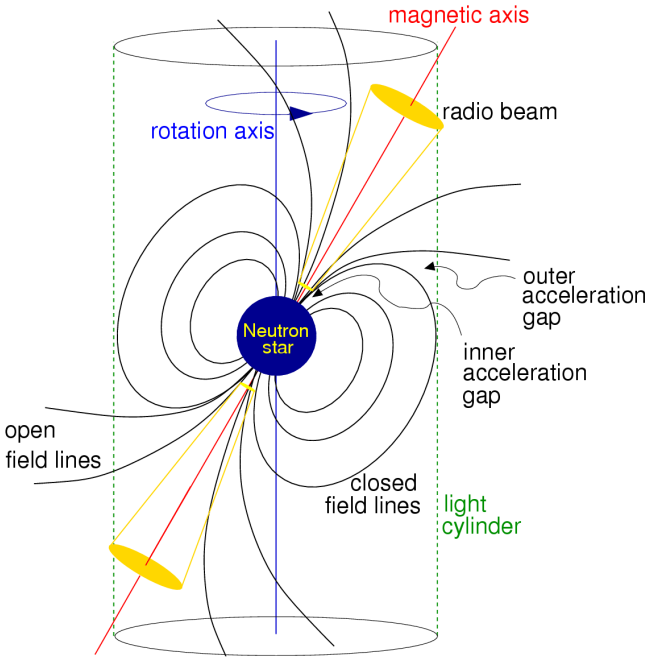

FIGURE 1.2: Illustration of a pulsar and its magnetosphere. The co-rotating magnetic field lines are only closed within the so-called light cylinder, of radius $R_{lc} = \frac{c}{2\pi}P$, which represents the radius from the pulsar centre at which magnetic field lines must rotate at the speed of light. The radio emission region is slightly above the poles of the pulsar. Image credit: Lorimer & Kramer (2004)

a spin frequency of 716 Hz, while the slowest spinning pulsar, PSR J0250+5854 (Tan et al., 2018), has a period of 23.5 s. It is important to note that the fast spinning pulsar is part of a class known as 'millisecond pulsars' (MSPs), which are discussed in detail from section 1.3 onwards. The canonical spin periods are usually taken as $P \sim 0.1 - 1$ s for 'standard' pulsars, and $P \lesssim 20$ ms for millisecond pulsars. These periods represent two ends of a diverse spectrum of pulsars, though almost every one has a common trait of extremely stable spin periods over time. Measurements of the rate of rotational slowdown, that is, the rate of increase of spin period, $\dot{P} = \mathrm{d}P/\mathrm{d}t$, show that the period increases at a rate of around $\dot{P} \sim 10^{-20}$ s s^{-1} for millisecond pulsars, to a more modest $\dot{P} \sim 10^{-15}$ s s^{-1} for 'standard' pulsars (Lorimer & Kramer, 2004). The pulsar spin down rate can be equated to the rate of loss of rotational kinetic energy,

$$\dot{E} \equiv \frac{\mathrm{d}E_{rot}}{\mathrm{d}t} = -\frac{\mathrm{d}I\Omega^2/2}{\mathrm{d}t} = -I\Omega\dot{\Omega} = 4\pi^2\frac{I\dot{P}}{P^3}, \tag{1.1}$$

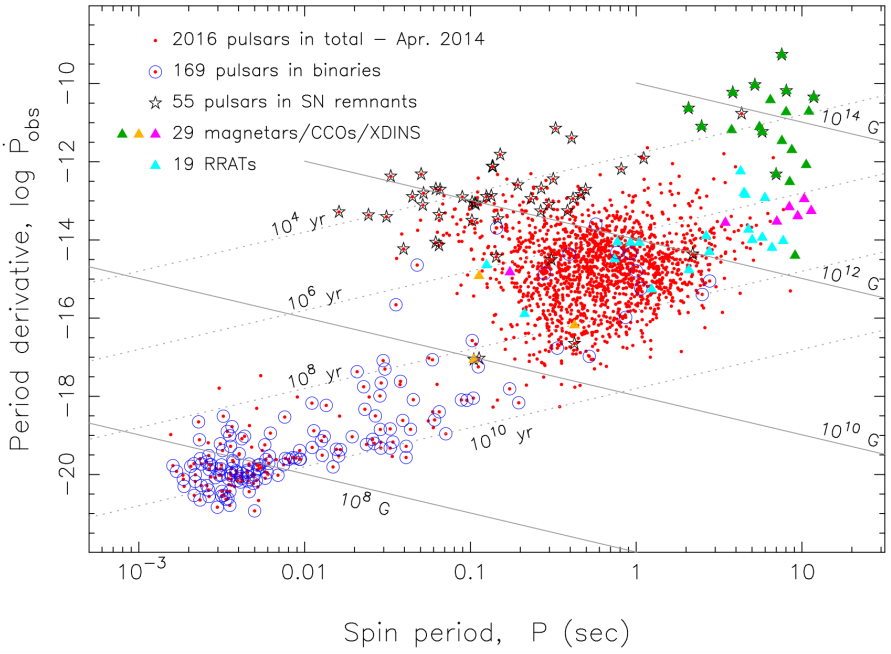

FIGURE 1.3: A $P - \dot{P}$ diagram, showing the distribution of known pulsars in 2014. Note the diagonal lines denoting the pulsar ages and magnetic field strengths, both of which can be expressed in terms of P and $\dot{P}$. The population of isolated and binary pulsars are discussed in the text. There are a wide range of other pulsar 'species', not covered in this literature review. Image credit: Tauris (2015), with data from Manchester (2014)

where $\Omega = 2\pi/P$, and I is the moment of inertia of the pulsar, approximated by

$$I = \frac{2}{5}MR^2 \tag{1.2}$$

for a solid sphere; the value derived from the canonical mass and radius is $I = 10^{45}$ g cm^2. $\dot{E}$ is known as the spin-down luminosity[1] which represents the total pulsar power output. These spin-down rates, compared to the periods of the pulsars, are so small and so stable that many pulsars are as accurate timekeepers as the best atomic clocks on long timescales (Hobbs et al., 2010; Verbiest et al., 2008). Note that in absolute terms $\dot{E}$ is still significant; a canonical MSP with period $P = 3$ ms and $\dot{P} = 10^{-20}$ s s^{-1} will have a spin-down luminosity of order $L \sim 10^{29}$ erg s^{-1}.

The population of pulsars itself is best described by means of a so-called $P - \dot{P}$ diagram, as shown in figure 1.3. From this diagram it can be seen that the 'zoo' of pulsars splits into several distinct groups, each with distinct properties. 'Standard'

[1]The spin down luminosity is usually measured in erg s^{-1}, where 1 erg = 10^{-7} J

radio pulsars are represented by single red dots, while millisecond pulsars are typically found in binary systems, indicated by blue circles in the bottom-left of the diagram: these include the spider pulsars. These pulsars are generally significantly older than the main pulsar population, with ages around 10^9 yr, while 'young' pulsars in SN remnants are typically only a few tens of thousands of years old. This is due to the fact that MSPs have been 'recycled'; spun up by accretion of mass from their companion stars (Bhattacharya & van den Heuvel, 1991). The canonical pulsar falls somewhere between these two extremes, with a median age of around 10^7 yr. Notably, it can also be seen that very few, if any, pulsars appear below the 10^{10} year age line. Here the 'death line', which represents the theoretical age at which a pulsar can no longer emit radio waves (Chen & Ruderman, 1993), is found, though note that it is not equivalent to this age line. The majority of this literature review will focus on the millisecond pulsars, and more specifically those in binary systems.

1.2.2 Theoretical background

Energy loss, rotation and magnetism

Timing observations indeed observe that pulsar periods increase over time, indicating that the pulsar is losing energy. As is detailed in section 1.2.2, only a small fraction of the power output manifests as radio emission. Instead, the majority of energy is lost to magnetic dipole radiation and the so-called pulsar wind in the form of high energy radiation. It is known from classical electrodynamics (Jackson, 1962; Lorimer & Kramer, 2004; Sturrock, 1971) that a magnetic dipole of moment $|\mathbf{m}|$ rotating with angular frequency Ω will radiate an electromagnetic wave with frequency equal to Ω. The radiation power of this wave will be

$$\dot{E}_{\rm dipole} = \frac{2}{3c^3}|\mathbf{m}|^2\Omega^4 \sin^2\alpha, \tag{1.3}$$

where α is the offset angle between the magnetic moment axis and the rotation axis. If it is assumed that the majority of power lost is indeed in the form of magnetic dipole radiation, equation 1.3 can be equated to equation 1.1 such that

$$\dot{\Omega} = -\left(\frac{2|\mathbf{m}|^2 \sin^2\alpha}{3Ic^2}\right)\Omega^3 \tag{1.4}$$

describes the evolution of spin frequency. Note that $\dot{\Omega}$ is necessarily negative unless an external effect spins up the pulsar. This can be expressed as a general power law,

using the non-angular frequency $\nu = 1/P$,

$$\dot{\nu} = -K\nu^n, \tag{1.5}$$

where n is the so-called braking index and K is a constant. By differentiating this equation it is possible to express n purely in terms of the spin frequency and its derivatives,

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}. \tag{1.6}$$

From equation 1.4 we can see that $n = 3$ for systems where energy loss is caused entirely by magnetic dipole radiation, that is, rotating neutron stars in a vacuum. However, as the pulsar magnetosphere is filled with plasma extracted from the neutron star, the observed braking index will be different. From Kaspi & Helfand (2002), a range of braking indices from $n = 1.4$ to $n = 2.9$ are observed, indicating that the assumptions made in this treatment are inaccurate. That is, there are significant contributions to energy loss other than magnetic dipole radiation. A second contribution, known as the Goldreich-Julian term, is often included to account for the effects of plasma currents in the pulsar magnetosphere (Goldreich & Julian, 1969). However, this early treatment is insufficient and there has been little success in analytically describing the structure and action of the pulsar magnetosphere since. For example, work from Mestel (1972) and Mestel & Pryce (1992) on an oblique rotator model produced results that differed very little from that of Goldreich & Julian (1969). More recently, numerical simulations have offered promising results. (see, e.g., Contopoulos et al. (1999); Chen & Beloborodov (2014); Kalapotharakos & Contopoulos (2009)).

The electromagnetic field strength, B, can also be estimated by considering the spin-down of the pulsar, following the treatment in Tauris (2015), by writing the radiation power as

$$\dot{E}_{\text{dipole}} = \frac{2}{3c^2}|\ddot{\mathbf{m}}|^2, \tag{1.7}$$

where

$$|\ddot{\mathbf{m}}|^2 = BR^3\Omega^2 \sin\alpha, \tag{1.8}$$

and equating this to equation 1.1 such that

$$B_{\text{dipole}} = \sqrt{\frac{3c^3 I}{8\pi^2 R^6} P\dot{P}} \text{ Gauss} \sim 3.2 \times 10^{19}\sqrt{P\dot{P}} \text{ Gauss} \tag{1.9}$$

for $\alpha = 90°$, $R = 10$ km, and $I = 10^{45}$ g cm^2. As shown in figure 1.2, there is some angle from the equator at which field lines cannot close. The regions of open field

lines at the magnetic poles are known as the polar caps, and are important when considering the emission of the pulsar.

Emission

Early models, such as those based on observations of pulse polarisation and spectra by Radhakrishnan & Cooke (1969) and Komesaroff (1970), are still useful for describing several properties of pulsar emission. These models describe a radio beam originating from the narrow cone of open field lines at the magnetic poles, as in figure 1.2. The open field lines in this region cause plasma to flow from the neutron star surface into the magnetosphere, causing the emission of photons parallel to the field lines. As such, the beam width of the photon emission is proportional to the width of the emission region, some height above surface at the polar cap, with beams emitted at a lower height narrowest. It must be noted that the *observed* pulse width depends on other factors, for example the inclination of the pulsar magnetic poles to the observer. Even with this caveat, Cordes (1978) showed that higher radio frequencies are emitted at lower emission heights as the observed pulse widths are narrower. However, observations of millisecond pulsars suggest that the emissions heights are roughly independent of frequency. Intuitively, this is as expected as for millisecond pulsars the light cylinder (and thus magnetosphere and emission region) are significantly smaller (Kramer et al., 1999).

While the observed properties of the radio pulses are well studied, the emission process is still not completely understood. The radio pulses are highly polarised, coherent and are effectively homogeneous across several orders of magnitude of spin period and magnetic field strength. As well as this, they are observed over 3 orders of magnitude of frequency, such that the emission process must also be broad band. Early models considered antenna mechanisms, such as emission by bunches of plasma moving along magnetic field lines (Komesaroff, 1970; Sturrock, 1971; Sutherland, 1975). However, these processes were inefficient (Lesch et al., 1998) and could not justify how the bunches maintained their shape over time. Later models proposed relativistic plasma emission (Melrose, 1992), where the radiation stems from instabilities in the plasma, which form so-called Langmuir turbulence waves which are in turn, by a non-linear process, converted into radiation (Asseo, 1993). Other proposed emission mechanisms include maser emission (Rowe, 1995) and coherent cyclotron emission (Machabeli & Usov, 1979). As for the high energy emission, observed as non-thermal radiation in the optical, X-ray, and γ-ray frequencies, there is debate even as to where the emission originates. Harding (1986) and Harding & Muslimov (2002) suggest the emission originates from the polar caps, via inverse

Compton scattering, while Cheng et al. (1986) and Romani (1996) suggest emission from the outer acceleration gaps, involving particle acceleration and high-energy radiation emission at large radii.

Propagation and observational quantities

For the radiation emitted by pulsars to be observed on Earth, it must first pass through the interstellar medium (ISM); a cold, ionised plasma composed mostly of hydrogen. As it propagates through the interstellar medium, the signal will be affected primarily by four effects: frequency dispersion, Faraday rotation, scintillation, and scattering. The first two can be described using a simple, homogeneous model of the interstellar medium, while the latter two require the consideration of turbulence and will not be discussed in this review.

Frequency dispersion The frequency-dependent refractive index of the interstellar medium, $\mu(f)$, has the form

$$\mu = \sqrt{1 - \left(\frac{f_p}{f}\right)^2}, \tag{1.10}$$

where f_p is the plasma frequency, $f_p = \sqrt{e^2 n_e / \pi m_e}$, with n_e the electron number density, and m_e the electron mass. It can be seen that an electromagnetic wave will only propagate if $f > f_p$, and that necessarily $\mu < 1$ such that the group velocity $v_g = c\mu < c$. As the refractive index is frequency dependent, so too is the group velocity. As such, the arrival times of a signal emitted some distance d away along line of sight l will be dependent on frequency. Therefore we can define a time delay, t, compared to some undispersed signal, given by

$$t = \int_0^d \frac{\mathrm{d}l}{v_g} - \frac{d}{c}. \tag{1.11}$$

By assuming that $f \gg f_p$, it can be written that

$$t = \frac{\mathcal{D}}{f^2} \int_0^d n_e \mathrm{d}l \equiv \frac{\mathcal{D}}{f^2} \times \mathrm{DM}, \tag{1.12}$$

where $\mathcal{D}$ is the so-called dispersion constant, $\mathcal{D} \equiv \dfrac{e^2}{2\pi m_e c}$, and DM is the dispersion measure,

$$\mathrm{DM} = \int_0^d n_e \mathrm{d}l. \tag{1.13}$$

The dispersion measure is measurable by comparing the pulse arrival time at different frequencies, and can be used to determine the distance to the pulsar if a model of the Galactic electron density is known (Taylor & Cordes, 1993).

1.2.3 Timing

In this section we follow Lorimer & Kramer (2004) for a very brief introduction to pulsar timing. In order to accurately time the pulse periods of pulsars, it is important to be able to correctly define the pulse time of arrival (TOA). The pulse time as measured on earth, the topocentric arrival time (t_{top}), needs to be transferred to the reference frame of the rotating pulsar. First, the topocentric arrival time measured at the observatory must be transferred to the Solar System centre-of-mass, the Solar System Barycentre (SSB). This corrects for motion of solar system bodies and gravitational effects too, and is obtained with the expression

$$t_{\rm SSB} = t_{\rm top} + t_{\rm cor} - \frac{\Delta D}{f^2} + \Delta R_\odot + \Delta S_\odot + \Delta E_\odot. \tag{1.14}$$

The first term, $t_{\rm cor}$, represents clock corrections, ΔD is a correction for the dispersion measure to remove frequency dispersion from the data, and the final three terms correct for various time delays. The first of these, $\Delta R_\odot$, is the Römer delay correction for the classical light travel time between the telescope and the SSB. The second, $\Delta S_\odot$, is the Shapiro delay to correct for gravitational time delays caused by masses in the solar system. Ideally, this should sum over all the bodies in the solar system, though in practice the sun dominates in almost all cases. Lastly, $\Delta E_\odot$ is the Einstein delay, correcting for gravitational redshift and time dilation due to the Earth's motion. In most cases the motion of the pulsar relative to the SSB must be taken into account, which means that a similar series of corrections must be applied. For pulsars in binary systems, the corrections are further complicated by the role of the pulsar companion, which can introduce its own relativistic delays.

1.2.4 Pulsar wind

First proposed by Pacini (1967), the pulsar wind is produced as ultrarelativistic particles, accelerated along magnetic field lines of the pulsar magnetosphere, break free of confinement at the light-cylinder radius. These particles carry away magnetic flux from the pulsar and typically have energies from 10-100 MeV. The pulsar wind is a significant source of energy loss for the pulsar, as introduced in section 1.2.2. Due to the necessarily subluminal speeds of the plasma the magnetic field lines outside of light cylinder 'fold' backwards relative to those within, causing a differential rotation

that results in further energy loss due to so-called magnetic braking. This energy in turn drives the plasma wind (Spitkovsky & Arons, 2002). Kirk et al. (2009) notes that this energy loss is independent of the alignment of the magnetic field relative to the rotation axis, such that an aligned rotator would still experience this effect. However, pulsars are oblique rotators which means another effect is seen, known as a striped wind. The striped wind is a spatially periodic structure in the expelled plasma traced out by the rotating pulsar, observed as so-called current sheets of higher temperature and density (Kirk et al., 2009).

Additionally, the difference in pressure between the pulsar wind and the surrounding material results in a termination shock. The termination shock sources the high-energy particles which comprise pulsar wind nebulae (PWNe), such as the Crab nebula (e.g., Hester et al. (2002)) Illustrated in figure 1.4a is a false-colour image of the X-ray emission from the Crab nebula. Figure 1.4b shows a diagram of the PWNe interior; these nebulae are sources of radiation from high-energy X- and γ-rays to radio. Of importance to spider systems is that the pulsar wind is thought to cause the ablation of the companion, and thus also cause the radio eclipses observed (Fruchter et al., 1988). The magnetisation parameter, σ, the fraction of power carried in the magnetic field, greatly varies: it is as small as $\sigma \sim 0.003$ in the Crab PWN and as large as $\sigma \sim 1$ in the Vela PWN (Helfand et al., 2001).

1.3 Binary pulsars

1.3.1 Introduction

While the vast majority of known pulsars are isolated, roughly 8% (Manchester, 2014) exist in binary systems. The first binary pulsar system was discovered by Hulse & Taylor (1975) from detections of a periodically varying pulse period. Characteristically, pulsars in binary systems have short spin periods, small period derivatives, small magnetic fields, and large ages compared to isolated pulsars (Tauris, 2015). Pulsars in binary systems provide excellent laboratories for many areas of physics, a notable example being gravitational physics with the Hulse-Taylor binary (a neutron star and pulsar) providing one of the best tests of general relativity (Hulse & Taylor, 1975; Weisberg et al., 2010). While this particular pulsar system is non-eclipsing (and non-interacting outside of gravitational interaction), discoveries of eclipsing binary systems with non-compact pulsar companions (e.g. Fruchter et al. (1988)) have provided new opportunities to study behaviour like accretion, ablation, and

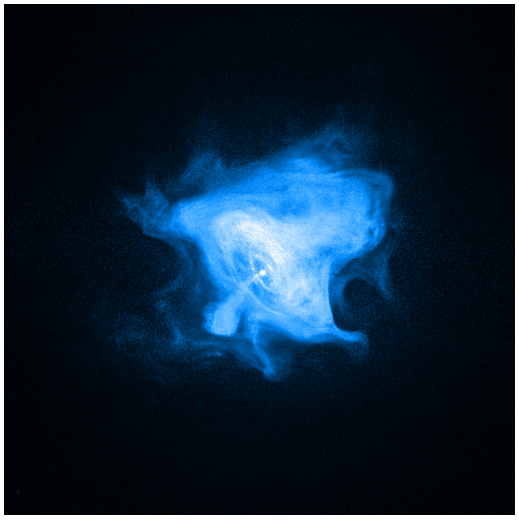

(A) Image of the Crab nebula produced using CHANDRA Observatory X-ray data. The spiralling sheet structure can be seen, as well as the shock-heated region near the centre and a jet in the lower left. Image credit: Seward et al. (2006) via NASA APOD

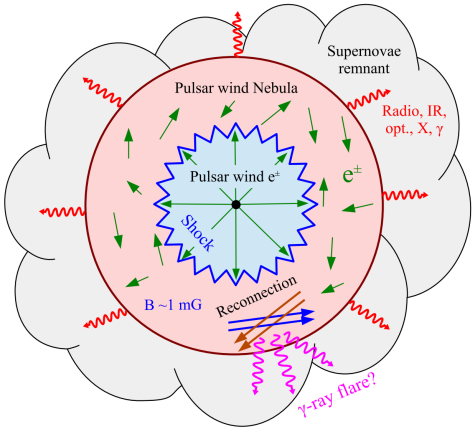

(B) Illustration of a PWN shock. The inner shock represents the edge of the expanding pulsar wind and is the point where electron-positron pairs are produced, which emit synchrotron radiation. Other regions of emission are shown, such as the proposed emission of γ-rays from magnetic reconnection. Image credit: Cerutti et al. (2014)

FIGURE 1.4: Left: X-ray image of the Crab pulsar wind nebula. Right: Illustration of a pulsar wind nebula.

optical effects during ingress and egress of the eclipse. Likewise, the double pulsar system provides further tests of general relativity, and can be used to probe the pulsar magnetosphere (Breton et al., 2008)

1.3.2 Timing in binaries

Keplerian delays

Binary pulsars show periodic variability in observed pulse period (i.e. the pulse arrival time), caused by both Keplerian and relativistic delays (Lorimer, 2008). For binary systems with lower orbital velocities, outside the relativistic regime, the Keplerian description will suffice. As 'spider' binaries fit this criterion, the relativistic delays will not be discussed in this work.

The mass function

The mass function, $f(M_1, M_2)$, of a binary system can be used to relate the mass of the pulsar and companion, based only on Keplerian parameters. Here we define the pulsar as the primary object, such that $M_1 \equiv M_{\rm psr}$. For a circular orbit, the

mass function can be derived from Kepler's third law,

$$GM_{\rm tot} = a^3 \frac{4\pi^2}{P_{\rm orb}^2}, \tag{1.15}$$

where G is the gravitational constant and $P_{\rm orb}$ is the orbital period, and a is the semi-major axis, defined as the sum of the distance of the two bodies from the centre of mass (CoM): $a = a_1 + a_2$. Similarly, $M_{\rm tot}$ is the total mass of the system, defined by $M_{\rm tot} = M_1 + M_2$. The general CoM equation defines the vector position, $\mathbf{R}$, of the centre of mass of number of bodies $i = 1, .., N$. For bodies of mass m_i and position $\mathbf{r}_i$, this is defined as

$$\mathbf{R} = \frac{1}{M_{\rm tot}} \sum_{i=1}^{N} m_i \mathbf{r}_i. \tag{1.16}$$

Conventionally, we assume that the CoM position is the origin in our reference frame; $\mathbf{R} \equiv \mathbf{0}$. This definition then allows us to relate the masses and separations of our binary system as $M_1 a_1 = M_2 a_2$, which can be substituted into equation 1.15 to give

$$\frac{M_2^3}{(M_1 + M_2)^2} = \frac{4\pi^2 a_1^3}{G P_{\rm orb}^2}. \tag{1.17}$$

By defining the projected radial velocity of the primary as $K_1 = v_1 \sin i$, where i is the orbital inclination and the velocity is given by $v_1 = \frac{2\pi a_1}{P_{\rm orb}}$, we arrive at the binary mass function:

$$f \equiv \frac{M_2^3 \sin^3 i}{(M_1 + M_2)^2} = \frac{P_{\rm orb} K_1^3}{2\pi G}. \tag{1.18}$$

The orbital inclination is the angle between the orbital angular momentum vector and the vector along our line of sight. By convention, an orbital inclination of $i = 90°$ corresponds to a completely edge-on system.

The mass function is frequently defined in terms of the mass ratio, $q = \frac{M_1}{M_2}$ (that is, the pulsar mass divided by the companion mass), and M_1 as

$$f \equiv \frac{M_1 \sin^3 i}{\left(1 + \frac{1}{q}\right)^2} = \frac{P_{\rm orb} K_2^3}{2\pi G}, \tag{1.19}$$

where we have used the relationship $q = \frac{M_1}{M_2} = \frac{a_2}{a_1} = \frac{K_2}{K_1}$. It must be noted that since the mass function depends of the inclination of the orbit, i, that is in almost all cases unknown, often we may only obtain an upper limit on the mass function (with $\sin\ i\ =\ 1$).

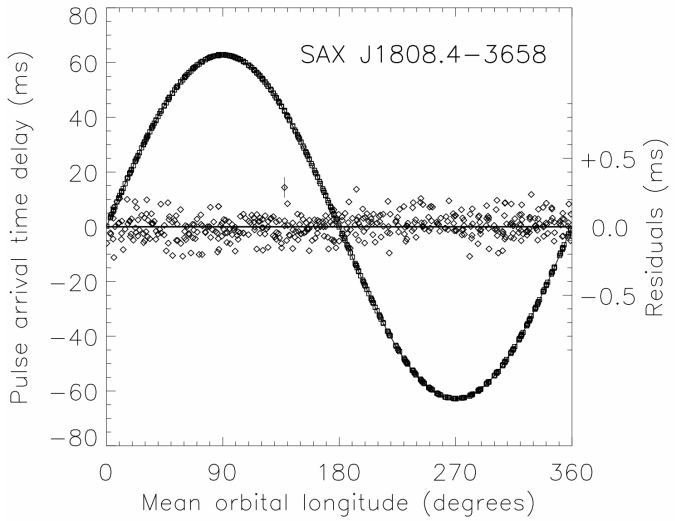

FIGURE 1.5: Timing delays over the two-hour orbit of binary pulsar SAX J1808.4-3658. The squares are the measured arrival time delays, and an orbital model curve is fitted. From this model, the mass function and other orbital parameters can be extracted. The residuals are plotted as diamonds. Image credit: Chakrabarty & Morgan (1998).

1.3.3 Formation

Recycling

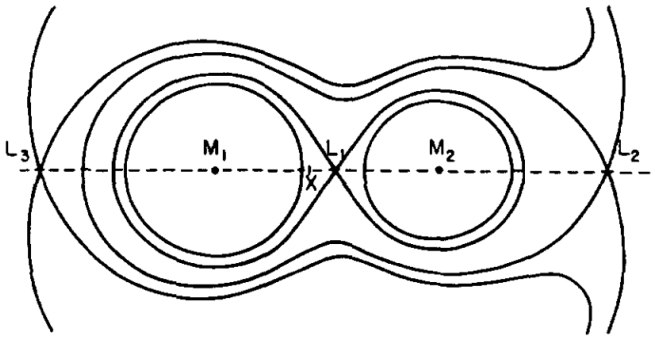

FIGURE 1.6: Illustration of equipotential surfaces and Lagrange points for two masses $M_1 > M_2$. The Roche lobes of the two masses are the two pear-shaped regions in a figure-8 configuration. If matter overflows the Roche lobe of M_1, it will transfer to mass M_2 via the L_1 Lagrange point. Image credit: Bhattacharya & van den Heuvel (1991).

It is generally accepted that binary pulsar systems start their lives as binary stellar systems (Tauris, 2015; Tauris & van den Heuvel, 2006), with differences in the formation route taken accounting for the different types of binary pulsar system. At the original orbital separation, little interaction is expected between the

two components and their evolution can be considered independently. In this situation, the more massive star typically evolves faster and reaches the end of its life by collapsing into a neutron star in a supernova. The binary system will be maintained if the companion is not disturbed in the supernova. The neutron star inherits a spin period from the progenitor during collapse, and so rotates as a pulsar at birth. Over time, the pulsar's spin period will increase as it approaches the death line as it loses angular momentum. However, if the pulsar enters a so-called recycling phase, where mass is transferred onto the pulsar from the companion, the pulsar will spin up as angular momentum is transferred from the companion (Alpar et al., 1982; Bhattacharya & van den Heuvel, 1991). In most cases this spins the pulsar up to periods significantly shorter than its initial spin period, into the region of millisecond periods (Bhattacharya & van den Heuvel, 1991). It is now important to discuss the conditions under which the transfer of mass in binary systems typically occurs: Roche lobe overflow. The Roche lobe (RL) of a star, illustrated in figure 1.6 represents an equipotential surface within which material is gravitationally bound to that star. If the companion expands (due to, for example, entering the red giant phase), it will overflow its Roche lobe and material will be free to accrete onto the pulsar. This process, occurring during the so-called X-ray binary (XRB) phase where the accretion of matter onto the pulsar occurs via a X-ray bright accretion disc, also tends to reduce the eccentricity of the binary orbit through tidal forces (Verbunt & Phinney, 1995) and reduce the magnetic field of the pulsar through dissipation of magnetic flux (Bhattacharya, 2002).

Angular momentum transfer

The loss and exchange of angular momentum, both orbital and spin, is important in determining the evolution of the binary system as it determines the orbital period and spin frequencies of the stars. The loss and transfer can shrink the orbital separation and reduce the size of the Roche lobes, causing overflow of material, or widen the orbit and stabilise mass transfer (Tauris & Savonije, 2001). There are several sources of loss, such as magnetic braking due to co-rotation with the pulsar wind (Verbunt & Zwaan, 1981), spin-orbit coupling (Witte & Savonije, 1999), gravitational waves, and mass loss from the system (Soberman et al., 1997). With magnetic braking, particles carried along the co-rotating magnetic field lines of the pulsar will eventually escape the system and carry with them angular momentum. However, the stellar wind of the companion can also carry away angular momentum; if the two stars are rotating synchronously with the orbital motion then these winds carry away angular momentum directly from the orbit (Rappaport et al., 1983).

Note that this process can also occur in LMXB systems without a pulsar wind as the magnetic braking from the companion star alone is sufficient to shrink the orbit. In this case, the tidal forces responsible for the synchronous rotation of the companion with the orbit produce a change in the orbital period as the spin period of the companion decreases. This process is related to the spin-orbit coupling described in Witte & Savonije (1999). Gravitational waves can also carry away significant amounts of angular momentum, though only in systems with very small orbital separations. Bhattacharya & van den Heuvel (1991) and Tauris & van den Heuvel (2006) provide thorough descriptions of these processes.

X-Ray binaries and their evolution

Following Tauris & van den Heuvel (2006), we will briefly detail two sub-classes of X-ray binary. High-mass X-ray binaries (HMXBs), shown in figure 1.7a, are characterised by a generally luminous, giant companion of mass $> 10M_{\odot}$ which almost fills its Roche lobe, with orbital periods of $\sim 1 - 10$ days. Low-mass X-ray binaries (LMXBs), shown in figure 1.7b, are characterised by a comparatively dim companion stars of mass $< 1M_{\odot}$, with typical orbital periods of a few hours to a few days. These companion stars, while low mass, are not dwarf stars. In fact they are generally slightly evolved from the main sequence such that they are able to fill their Roche lobes.

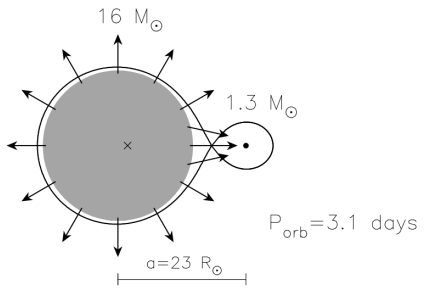

(A) Illustration of a HMXB, showing accretion via the strong stellar wind of the giant companion. Note that at later stages in the accretion phase Roche lobe overflow can occur.

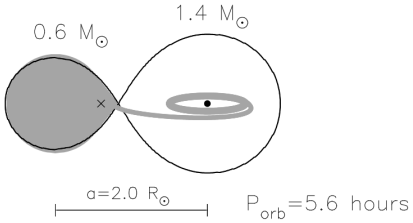

(B) Illustration of a LMXB, showing accretion via Roche lobe overflow through the L_1 Lagrange point. The accretion disk around the neutron star companion can be seen.

FIGURE 1.7: Illustrations of low- and high-mass X-ray binaries. Image credit: Tauris & van den Heuvel (2006)

Accretion in HMXBs initially occurs via interception of a strong stellar wind from the giant companion, though near the end of the giant's life it may engulf the compact object in a common envelope evolution. If this occurs large amounts

of orbital angular momentum will be lost and the orbital separation will decrease, constricting the Roche lobes of the system. Mass will instead transfer via Roche lobe overflow, resulting in sufficient spin-up to millisecond pulsar-periods (Tauris & van den Heuvel, 2006). This accretion process is often relatively short, lasting on the order of 10^5 yr. Subsequent supernovae of the giant star can leave a compact remnant in the binary if it is not disturbed, with this being the likely source of double-pulsar binaries.

The optical companion in LMXBs does not have a strong stellar wind, but instead has filled its smaller Roche lobe so overflow can occur. If this happens, the companion will overflow its Roche lobe and lose material through the L_1 Lagrange point towards the neutron star. As such, the accretion usually occurs via an X-ray bright accretion disk around the neutron star. LMXBs have typically long accretion timescales, on the order of 10^8 yr, as this accretion process is slow and stable. Tauris & Savonije (1999) suggest that this results in the these systems hosting the fastest-spinning, most massive MSPs as significant angular momentum transfer is allowed to occur.

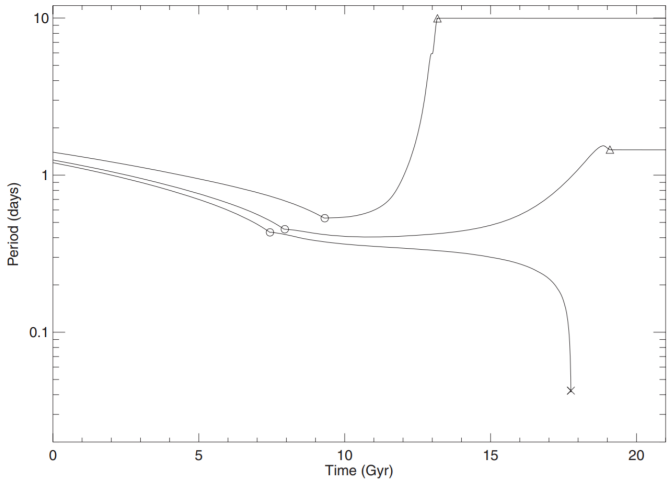

FIGURE 1.8: Three simulated orbital evolution trajectories with similar initial periods, illustrating the bifurcation process. From top to bottom, are two diverging systems and one converging system. The 'o' symbol indicates the start of Roche lobe overflow, the '△' symbol the end of the overflow, and the 'x' indicates the end of the simulation. Image credit: Ma & Li (2009)

The orbital evolution of LMXBs depends sensitively on the orbital period of the binary prior to mass transfer. Ma & Li (2009) describe a bifurcation process with two outcomes: a converging system with decreasing orbital periods, and diverging systems with increasing orbital periods and separations. An illustration of this

bifurcation can be seen in figure 1.8. The converging system is the most relevant to this discussion, as it forms a compact binary with a low-mass companion, rather than the typically large-separation white dwarf binaries formed from the diverging system. These systems have periods of < 1 day, and very low companion masses, $M_c < 1\mathrm{M}_\odot$. A significant factor which contributes to this bifurcation is the nature of the mass transfer (King, 1988). If the mass transfer is conservative, that is, there is no net change of total angular momentum in the system (such as under the assumption that 100% of the mass lost from the companion is accreted by the primary), then the binary separation will expand. This is because mass concentrates near the binary centre of mass (i.e. within the neutron star) and so the companion must orbit at a larger distance to conserve angular momentum. This of course has the effect of expanding the companion's Roche lobe and therefore halting mass transfer. Note that this occurs in general where there is conservative mass transfer from the low-mass object to the high-mass object.

The masses of companions in LMXBs are such that further expansion is unlikely as we do not expect them to evolve onto the giant branch with such small masses, though there may be a less significant increase in radius if the star evolves off the main sequence. Instead, we may consider the case where angular momentum is lost from the system; only a fraction of overflowing mass accretes onto the neutron star and the rest is carried away by winds. In this case, the Roche lobe of the companion does not significantly expand and mass transfer can continue.

However, to arrive at our desired binary MSP, it is not enough for the neutron star to be spun up to millisecond periods; the radio pulsar must also switch on and the accretion disc must disappear. As the LMXB system continues accreting mass from the companion, along this trajectory, the orbital period decreases due to angular momentum loss through magnetic braking until the companion star becomes fully convective. This occurs when its mass is reduced to around $0.2 - 0.3$ $\mathrm{M}_\odot$ or a period of a few hours (Pylyser & Savonije, 1989), resulting in the weakening of the companion star magnetic field and so the halting of magnetic braking. When this occurs the only remaining effect on the orbit of the binary is that of mass transfer, which expands the orbital separation and Roche lobe, leading to a definitive end to the mass transfer (King, 1988; Rappaport et al., 1983). At this point it is assumed that there has been sufficient mass transfer to spin up the neutron star to millisecond periods (Tauris et al., 2012) and the radio pulsar turns on. This occurs once the magnetosphere expands past the light cylinder radius as the accretion disc recedes. The resulting pulsar wind irradiates the companion star and starts to evaporate it (van den Heuvel & van Paradijs, 1988), causing further mass loss. At this stage, the companion may expand and fill its Roche lobe as it bloats due to adiabatic expansion.

However, mass transfer will not occur as any material that flows through the L_1 point is ejected from the system due to the radio ejection mechanism (Kluzniak et al., 1988), where the pressure of the pulsar wind prevents any material reaching the pulsar. Evaporation of the companion can widen the orbit, however (Chen et al., 2013).

1.4 Spider pulsars

1.4.1 Classification and formation

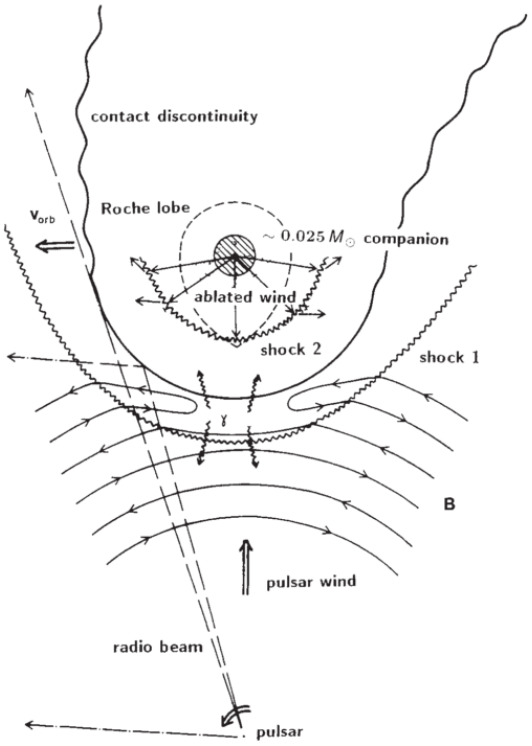

FIGURE 1.9: Illustration of the effect of the pulsar wind on the irradiated companion. The radio eclipsing by the ablated material can be seen, as well as the intrabinary shocks that occur as a result of the pressure difference between the pulsar wind and the ablated material outflowing from the companion. The emission of γ-rays in this region can also be seen. Image credit: Podsiadlowski (1991)

The first 'spider' system, detected by Fruchter et al. (1988), is a canonical example of the so-called 'black widow' pulsars, characterised by very low companion mass ($0.01\mathrm{M}_\odot < M_c < 0.05\ \mathrm{M}_\odot$). This eclipsing binary displays the characteristic millisecond pulsar period (1.6 ms), a nearly-circular orbit with a sub-day period (9.17 hours), periodic radio eclipses, and the ablation of the companion by

the pulsar wind. Later discoveries (e.g., Lyne et al. (1990)) showed this to indeed be part of a separate class of compact, eclipsing pulsar binaries with small companion masses. However, more recent discoveries of more massive irradiated companions, $0.1\mathrm{M}_{\odot} < M_c < 0.5\mathrm{M}_{\odot}$, by, for example, Ferraro et al. (2001) and Thorstensen & Armstrong (2005a) suggest a second class of spider pulsars, typically known as 'redbacks'. Redbacks are characterised by intermediate mass companions, and marginally longer orbital periods than black widows.

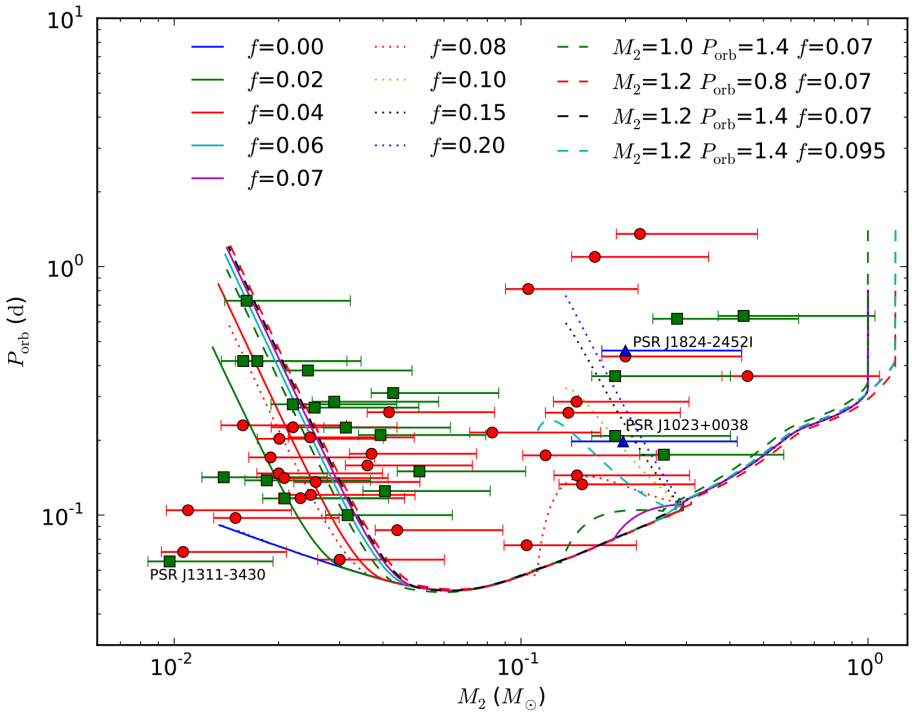

FIGURE 1.10: Simulated evolution tracks of binary systems with varying evaporation efficiency factors, f, showing the formation of two distinct populations. The left-hand cluster are black widows and the right-hand cluster are redbacks. Image credit: Chen et al. (2013)

From the compact eclipsing binaries arrived at in the previous section, Chen et al. (2013) suggest that the evolutionary routes of redbacks and black widows are separate, determined by the efficiency with which the pulsar evaporates the companion. By considering the angular momentum loss by gravitational waves, magnetic braking, ejection of material into a circum-binary disc, as well as mass loss by the evaporation of the companion, Chen et al. (2013) model the evolution of these compact binaries. The evaporation mass loss is dependent on an efficiency factor, f, whose value is dependent on the geometry of the system. For these compact binaries of orbital periods $\sim$ few hours, there is a delicate balance between the gravitational wave radiation which shrinks the orbit and mass loss due to evaporation, which widens it.

For larger values of f, indicating more efficient evaporation, the companion does not re-fill its Roche lobe and for smaller efficiencies it is able to. As the radio ejection mechanism is still active, any overflowing material is ejected from the system; as such where the companion remains detached from its Roche lobe, the orbit shrinks less quickly. This results in the seemingly unintuitive result shown in figure 1.10 that more efficient evaporation leads to the more massive redback systems, while less efficient evaporation leads to black widows. Conversely, Benvenuto et al. (2014) suggests that black widows evolve from redbacks with short periods. This is still an open question, as there is no conclusive evidence to determine if redbacks and black widows are separate populations or descendants.

This evolution of spider pulsars from LMXBs is supported by discoveries of so-called transitional millisecond pulsars (tMSPs) by, for example, Roy et al. (2015b) and Deller et al. (2015), where a compact binary pulsar is observed to transition between an accretion-powered LMXB-like state and a rotation-powered redback-like state. This class of system is discussed in detail in section 1.5

1.4.2 Physical properties

Irradiation and ablation

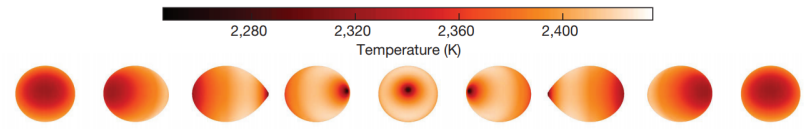

(A) Temperature distribution of the distorted, irradiated companion. Note the teardrop shape, and that the hottest part is not closest to the pulsar. The phase is matched to figure 1.11b

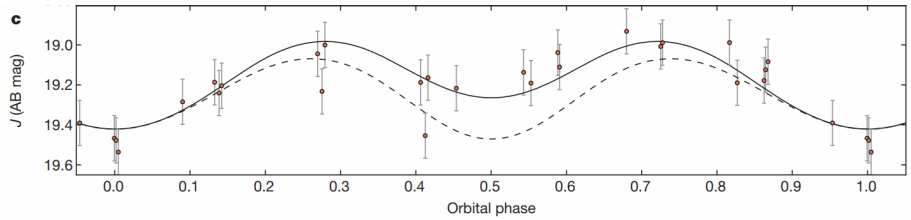

(B) Resulting optical light curve, with the fitted model shown as a solid line, showing both ellipsoidal variation and the effect of irradiation, such that the maxima occur when the companion star is 'side-on' to the observer. The two modulation components (fundamental and first harmonic) can be seen, resulting in the double-peaked light curve. The dashed line shows the model with no irradiation component, isolating the effect of ellispoidal modulation.

FIGURE 1.11: Optical light curve and illustration of the irradiated companion star in SDSS J143317.78+101123.3. Image credit: Hernández Santisteban et al. (2016)

The companion stars in black widow and redback systems are highly irradiated by the pulsar wind, resulting in a group of phenomena common to all spider systems. The two primary effects are the heating of the inner face of the tidally locked companion, and the ablation of the companion's convective envelope via the stripping of surface material (van den Heuvel & van Paradijs, 1988; Phinney et al., 1988; Kluzniak et al., 1988). This ablated ionised plasma is responsible for the characteristic eclipses of the radio pulsar observed in spiders. Podsiadlowski (1991) suggests that the irradiation causes the companion to bloat; outward radiation pressure is increased as the companion is heated, while the inward gravitational pressure is reduced due to the stripping of surface material. As such it is likely the companion will fill its Roche lobe, and lose more mass. As the orbital separation is small the shape of the companion's Roche lobe is such that, if filled, the star it will be substantially tidally distorted into a teardrop shape. It is expected that the efficiency of irradiation is 10-30% (Breton et al., 2012), including direct irradiation from TeV electrons and positrons from the high-energy pulsar wind, and indirectly from γ-rays emitted by a heated intrabinary shock (Chen et al., 2013). The heating and gravitational distortion effects can be seen in figure 1.11a and the resulting optical light curve can be seen in figure 1.11b.

We note two phenomena that affect the observed luminosity of the companion at different points, in addition to the heating. First, limb darkening (e.g., Milne, 1921) is usually an effect where the 'edge' of a star appears dimmer than the centre. This effect manifests in bloated companions as the emission temperature of a star tends to decrease with distance from the centre, such that the distorted shape of the companion means we observe non-uniform temperatures. Second, gravity darkening (e.g., Lara & Rieutord (2011)) is where the brightness of a point on the star is proportional to the radius of the star at that point, such that the 'tip' of the companion closest to the pulsar is dimmer than the rest as it has a lower surface gravity and therefore a lower temperature.

Eclipses

Eclipses of the pulsar's radio beam are observed in phase with the orbital period at the companion's inferior conjunction, when the companion is between the observer and the pulsar). They are often longer than expected from obscuration of the companion alone and as such Fruchter et al. (1990) ascribes them to the 'tail' of ablated, ionised plasma surrounding the system, resulting in radio delays that are longer than could be caused by eclipsing by the companion star alone. Lower radio frequencies are more easily eclipsed (Archibald et al., 2009); for PSR J1023+0038 these eclipses obscure pulses for up to 60% of the orbit at frequencies below 400 MHz,

while at 1.4 GHz this drops to 25%. This ablated material also causes frequency-dependent dispersion, most noticeable in the ingress and egress of the eclipse (Polzin et al., 2020) as shown in figure 1.12.

Many attempts have been made to explain the mechanisms of eclipse (see Thompson et al. (1994) for a review), but the most likely mechanism is thought to be cyclotron-synchrotron absorption (Polzin et al., 2019, 2020). This method relies on the magnetised pulsar wind surrounding the companion, or indeed a companion magnetosphere, to absorb radio waves at the cyclotron frequency of the magnetic field and higher harmonics. This successfully explains the frequency dependence, as the resulting optical depth of the magnetosphere is lower for higher frequencies. Bhattacharyya et al. (2013) and Stappers et al. (2001) report results that support this mechanism of radio eclipses.

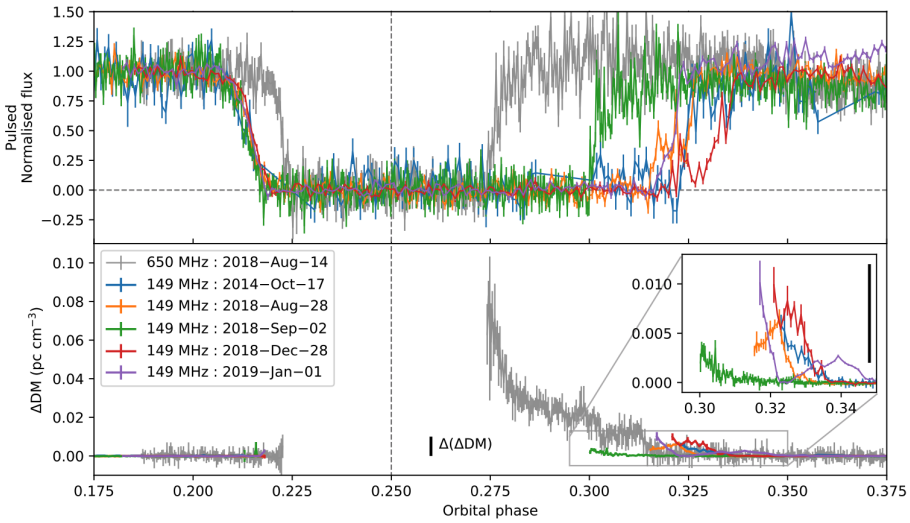

FIGURE 1.12: The radio eclipse of PSR J1816+4510. *Top:* pulsed flux densities from continuum images, normalised such that the flux density is unity out of eclipse. The frequency-dependent eclipse width can be clearly seen, with the higher frequencies eclipsed for shorter duration. *Bottom:* deviation from mean out-of-eclpise DM, from the same observations as the top panel. This panel shows how the DM drastically increases as the radio beam passes through the ablated material. Image credit: Polzin et al. (2020)

Optical properties

As can be seen in figure 1.11b, optical light curves of black widow and redback systems show quasi-sinusoidal luminosity variation with additional harmonics. The fundamental variation at the orbital period is caused by the heating of the day side of the companion by the pulsar wind. This results in a greater mean temperature on the day side, and so a greater mean luminosity. In addition to this, the tidally

distorted teardrop shape of the companion projects a large area in the sky when viewed side on (corresponding to phase 0.25 in figure 1.11); this results in the so-called ellipsoidal modulation at half the orbital period. However, the gravity and limb darkening effects discussed in section 1.4.2 also contribute to this effect.

To treat the irradiation effects on the companion star, Breton et al. (2013) introduce the so-called irradiation temperature, $T_{\rm irr}$, defined as

$$T_{\rm day}^4 = T_{\rm night}^4 + T_{\rm irr}^4, \tag{1.20}$$

where $T_{\rm day}$ and $T_{\rm night}$ are the day and night temperatures respectively. This relation assumes that all incident flux on the companion day side is thermalised and re-radiated within this side of the star. With this, the irradiation temperature can be related to the pulsar spin-down luminosity, L, as

$$\epsilon_{\rm irr} L = 4\pi a^2 \sigma_{\rm SB} T_{\rm irr}^4, \tag{1.21}$$

where a is the orbital separation, $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant, and $\epsilon_{\rm irr}$ is the effective fraction of pulsar spin-down luminosity that is absorbed by the companion (Breton et al., 2013). This term accounts for the fraction of L that is incident on the companion, as well as the albedo of the companion which is defined as the fraction of incident irradiation luminosity that heats the companion. Typically, the albedo is around 0.3, such that 30% of the pulsar's radiation does not heat the companion and is reflected (Baglio et al., 2016).

Optical photometry of these light curves can be used to constrain various orbital and stellar parameter by fitting a numerical model. These parameters can include the day and night side temperatures of the companion, its Roche lobe filling factor, as well as the inclination or mass function of the orbit. However, several of these parameters are degenerate and so additional information is needed. This degeneracy is typically broken by constraining the temperature distribution of the companion through the use of multi-band photometry. This technique is used in the modelling of two sources in chapter 3. Optical spectroscopy (Kaplan et al. (2013), for example) is another tool that can be used not only to determine the radial velocity semi-amplitude of the companion, but also infer the composition of the companion stars which can give clues as to the evolutionary origin of these systems. As such, a combination of the two (and the dynamical parameters obtained using radio timing) is essential in fully determining the behaviour of a system.

Magnetic fields and binary properties

Due to the small orbital separation of spider systems, the magnetic field of the companion becomes important as it interacts with both the pulsar magnetic field and that of the pulsar wind. As introduced in section 1.2.4, interactions between plasma and magnetic fields drive many processes that result in observable phenomena, such as synchrotron emission in the intrabinary shock, or the radio eclipses by the companion magnetosphere (Khechinashvili et al., 2000). It is thought that magnetic fields can also have an impact on the orbital period, via gravitational quadrupole coupling. Applegate & Shaham (1994) and Applegate (1992) describe an observable orbital period modulation with amplitude $\Delta P/P \sim 10^{-5}$ on timescales greater than 10^3 times the orbital period, and explain it via a coupling of the quadrupole moments of the stars. The quadrupole moments are dependent on the oblateness and internal structure of the stars so can offer insight into the magnetic structure of pulsars. These orbital period variations are observed in PSR B1957+20, for example Arzoumanian et al. (1994). Aside from these orbital period variations, spiders have very circular orbits due to tidal circularisation during the LMXB phase; this results in eccentricities, when detected, as small as 10^{-5} (Voisin et al., 2020).

1.5 Transitional millisecond pulsars

1.5.1 Accretion and the transition

The discovery of the radio pulsar PSR J1023+0038 (Archibald et al., 2009) at the position of a target previously classified as a cataclysmic variable (Bond et al., 2002) was the first identification of a new type of system: the transitional millisecond pulsar (tMSP) (Stappers et al., 2014a). Since then, two more systems, PSR J1227-4853 and PSR J1824-2452I, have been discovered (Roy et al., 2015b; Deller et al., 2015). tMSPs are observed to transition between two distinct states; an accretion powered state like LMXBs and a rotation-powered MSP-like state. In the former, the presence of an accretion disc results in an increase in luminosity by $\sim 10\times$ over the latter, which behave mostly indistinguishably from redbacks. These systems appear to straddle the divide between radio MSPs and LMXBs, at the point where the gravitational pull of the neutron star and the outward pressure of the pulsar wind on outflowing matter from the companion are balanced. As such, the study of tMSPs is important in uncovering the evolutionary history of spiders and LMXBs: they may be a missing link between these two populations (Archibald et al., 2009). In this section we cover much of the same material presented in Papitto & de Martino (2020), a very thorough review of the current body of research surrounding tMSPs.

Accretion discs in these binaries form as plasma outflowing from the companion, usually through the L1 point, falls toward the pulsar (Pringle & Rees, 1972). Note that this plasma has non-zero angular momentum, which is conserved in the following process. This plasma increases in density and temperature as it spirals towards the pulsar, at which point it begins to form a disc within the light cylinder of the pulsar. This disc differentially rotates due to Keplerian dynamics and the resulting friction heats it up further. The closed field lines within the light cylinder intersect with the disc and are bent by this differential rotation, causing an increase in magnetic stress. As such the disc is truncated at the radius where this stress dominates the viscous stress of the disc; this radius is known as the accretion radius, $R_{\rm acc}$. From here, there are two outcomes. If the accretion radius is less than the radius at which the disc material co-rotates with the neutron star, $R_{\rm co}$, then the inflowing matter accretes onto the neutron star surface. Conversely, if $R_{\rm acc} > R_{\rm co}$ then the material has insufficient angular momentum to accrete and is instead reflected away: this is the propellor mechanism (Illarionov & Sunyaev, 1975). If the accretion rate (i.e. the rate of mass loss from the companion) decreases, $R_{\rm acc}$ increases. First past $R_{\rm co}$, at which point accretion from the disc stops, and eventually past the light cylinder radius at which point mass transfer is assumed to have stopped; at this point any mass overflowing from the companion is ejected from the system. As such, it is the accretion radius and in turn the mass transfer rate that drives the transition; if the companion ceases to overflow its Roche lobe then the accretion will stop (Kluzniak et al., 1988). In the case of tMSPs and spiders, this occurs due to the irradiation of the companion by the high energy pulsar wind, however the exact mechanism is not known. Similarly, if the companion again overflows its Roche lobe, the pressure of in-falling matter can again form an accretion disc.

1.5.2 Known tMSPs

At the time of writing there are three tMSPs which have been observed to transition between the two states. These are PSR J1023+0023 (J1023), PSR J1227+4853 (J1227), and PSR J1824-2452I (J1824), which is in a globular cluster. The first two of these are the subject of chapter 3, and are discussed in further detail there.

J1023 was first identified as a variable radio source (Bond et al., 2002) with a flickering optical counterpart. It was initially classified as a polar cataclysmic variable (CV), a magnetised white dwarf in an accreting binary (Szkody et al., 2003). However, just a few years later its optical counterpart became dimmer and periodic sinusoidal variability was found in place of the previous flickering (Woudt et al., 2004). It remained in this new state and in 2007 a radio MSP was finally discovered, with a period of 1.69 ms (Archibald et al., 2009). As well as this, the

strong emission lines (He I, He II, and Balmer lines) reported in (Szkody et al., 2003) which indicated the presence of an accretion disc were replaced with a spectrum showing only the absorption lines characteristic of a mid-G type star (Thorstensen & Armstrong, 2005b). These absorption lines showed variable Doppler shift on an orbital period of 0.198 days, matching the period of the optical light curve reported in Woudt et al. (2004). From this, we can conclude that J1023 had an accretion disc until $\sim$ 2001, at which point it disappeared to reveal a redback system. More recently in June 2013, the strong emission lines reappeared and the radio pulsar disappeared (Stappers et al., 2014a), indicating that the system has returned to the LMXB state. During this transition, the gamma ray, X-ray, and ultraviolet flux observed increased as the accretion disc formed.

Similarly to J1023, J1227 was first classified as a CV based on the detection of a hard X-ray source, strong emission lines, and flickering in the optical light curve (Sazonov & Revnivtsev, 2004; Masetti et al., 2006; Pretorius, 2009). de Martino et al. (2010a) noted the coincidence of the optical and X-ray source with a yet-unassociated FERMI-LAT gamma source, suggesting J1227 was a LMXB with an unusually low luminosity of $L_{\mathrm{X-ray}} \sim L_{\gamma} \sim 10^{33}$ erg s^{-1}. These properties are consistent with those expected from the LMXB-state of a tMSP, and indeed as with J1023 the disappearance of optical emission lines and a significant ($\sim$ 10 times) dimming of X-ray, optical, and radio emission between 2012 November 14 and 2012 December 21 (Bassa et al., 2014a) indicated that a transition had occurred. The radio MSP was detected by Roy et al. (2015b) with a spin period of 1.69 ms, coincidentally an almost identical spin period to J1023.

Lastly, the tMSP J1824 was detected during an accretion outburst in March 2013 and quickly identified as an accreting MSP via the detection of pulsed X-rays at a period of 3.9 ms (Papitto et al., 2013). These pulsations were detected over the duration of the accretion event, which lasted for approximately one month. This source was coincident with a previously-detected radio MSP, and further radio pulsations were detected just two weeks after the last X-ray pulse detection.

1.5.3 tMSP characteristics

As summarised in Britt et al. (2017), observations of the three confirmed tMSP systems have revealed several shared characteristics, though it is important to note that due to the small sample size these could be coincidental. In the AP state, pulsed X-rays have been detected (Jaodand et al., 2016), as well as the bimodal flickering and flaring typically associated with the presence of an accretion disc (Patruno et al., 2014; Linares, 2014). The three confirmed tMSPs also share the same relationship between the luminosities of their radio and X-ray emission, $L_{\mathrm{Radio}} \propto L_{\mathrm{X}}^{0.7}$, also shared

with black hole binaries nearing quiescence (Deller et al., 2015). The X-ray light curves of tMSPs in the disc state are observed to rapidly switch between two states; a 'low' state of luminosity $L_{\rm low} \sim 10^{33}$ erg s^{-1}, and a 'high' state roughly an order of magnitude brighter (Patruno et al., 2014). As well as this, short X-ray flares are observed, presumably originating from the accretion disc as well, which can be several times brighter than the high state (Linares, 2014).

The optical emission in the RP state is indistinguishable from that of redback systems as described in the previous section, while the AP state emission exhibits additional flickering and flaring (Kennedy et al., 2018). Strong emission lines are seen in the AP state (Archibald et al., 2009; Bassa et al., 2014b) which fully disappear in the RP state, suggesting they originate from the accretion disc. In the AP state, tMSPs exhibit a flat radio spectrum suggesting self-absorbed synchrotron emission, while in the RP state the radio emission is pulsed with a spectrum characteristic of synchrotron emission with a steep power law (Archibald et al., 2009; Patruno et al., 2014), typical of rotation-powered MSPs.

1.6 Statistical, numerical, and machine learning methods

With the advent of 'big data' in astronomy, the growing complexity of models, and the increasing availability of very precise data, cutting-edge computational techniques are now an essential part of the astronomers toolkit. In this work, for example, we use machine learning to classify hundreds of thousands of individual light curves from all-sky surveys, and use recently-developed numerical methods to fully explore the parameter space of highly degenerate models. Here we introduce the numerical sampling method used in chapter 3, as well as its mathematical foundations in Bayesian inference. We also introduce the `Icarus` binary light curve synthesis code and how it is used to model the light curves of spider systems. Lastly, we introduce the machine learning classifier used in chapter 4, as well as some of the key metrics used for assessing performance.

1.6.1 Bayesian inference and sampling methods

The equation describing conditional probability,

$$P(A|B) = \frac{P(A \cap B)}{P(B)}, \tag{1.22}$$

which states that the conditional probability of observing event A given event B is equal to the joint probability of events A and B divided by the probability of event B, can be used to derive Bayes' theorem,

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \tag{1.23}$$

Bayes' theorem is a ubiquitous equation in statistics, which allows the conditional probability of an event A (given event B) to be calculated based on some prior knowledge of the events. Bayesian inference is a common application of Bayes' theorem, where it is used to obtain a more accurate model (hypothesis) based on some newly obtained data (evidence), as expressed by the equation

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \tag{1.24}$$

Here, $P(H|E)$ is the posterior probability, $P(E|H)$ is the likelihood function, $P(H)$ is the prior probability, and $P(E)$ is the marginal likelihood. The posterior probability described the probability of some model hypothesis, given the evidence (data) that we have fit it to, while the prior probability describes the probability of this model hypothesis before the data is observed; this usually includes information we have from other sources, such as literature measurements. The likelihood is a measure of how well the model fits the data; in our case it is generally derived from the reduced χ^2 statistic for example. Note that this is acceptable only when the data points are measured independently of each other and the uncertainties follow a Gaussian distribution.

Given some model hypothesis, we would want to calculate the posterior probabilities of the model parameters, that is, the distribution of parameter values which correspond to the model which best fits the data. Conversely, given several competing models, we would want to calculate the model evidence to determine which model is most appropriate given our data. In practice, both of these methods require the solution of multi-dimensional integrals which do not have analytical solutions or even cannot be numerically solved on practical timescales. This is especially true for multi-dimensional problems with many parameters, such as the ones in this work.

Instead, numerical tools such as MCMC (Markov Chain Monte Carlo) sampling or nested sampling are used. MCMC uses Monte Carlo methods to sample a large number of points over the parameter space, and calculate the posterior probability distribution at these points. MCMC methods are generally robust and reasonably efficient for multidimensional problems, however they can struggle where the posterior distributions are multimodal, as is the case in this work. Nested sampling samples

the posterior distributions and also estimates the marginal likelihood, defined as Z, by computing the integral

$$Z = \int \pi(\theta) L(\theta) \mathrm{d}\theta. \tag{1.25}$$

Here, θ are the model parameters, $\pi(\theta)$ is the prior distribution, and $L(\theta)$ is the likelihood function. In nested sampling, the multi-dimensional parameter space is reduced to a single dimension by defining the so-called cumulative prior mass, $X(\lambda)$, as

$$X(\lambda) = \int_{L(\theta)>\lambda} \pi(\theta) \mathrm{d}\theta, \tag{1.26}$$

such that as λ increases, $X(\lambda)$ decreases from 1 to 0. $X(\lambda)$ is the prior mass covering all prior values greater than $L(X)$. This allows the integral in equation 1.25 to be written as a one-dimensional integral:

$$Z = \int_0^1 L(X) \mathrm{d}X, \tag{1.27}$$

which is significantly less computationally expensive to solve. This method also allows us to directly compare the Bayesian evidence (often called the marginalised likelihood; a measure of the likelihood of the model over the prior probability space) of each choice of model for a particular problem. We use a so-called multi-nested sampler (Feroz et al., 2009) implemented in Python as `pymultinest` (Buchner, J. et al., 2014), which is well-suited to handling multi-modal posterior distributions.

1.6.2 Icarus

Here we briefly introduce the Icarus binary light curve synthesis code (Breton et al., 2012), which is used with the above numerical techniques to explore the parameter space. The details of the Icarus model are discussed in more detail in section 3.4.1 in chapter 3. For a given set of orbital and physical parameters, such as the temperature, mass ratio, and companion radial velocity[2], the Icarus model synthesises a companion star. The key model parameters are listed in table 1.1. We aim to determine the posterior distributions of the parameters using the multi-nested sampler.

The star is synthesised by constructed as tessellation over the expected star surface, given the binary parameters and the Roche lobe filling factor. Each surface element has an associated temperature, such that the temperature distribution can

[2]Because `Icarus` models the companion star it is defined as the primary within the `Icarus` code, such that the companion radial velocity `Icarus` parameter is labelled as K_1. To avoid confusion we will use the standard convention in this thesis, with the pulsar as the primary and the companion velocity labelled as K_2.

TABLE 1.1: Fit parameters for the `Icarus` model. Parameters with dimensionless quantities are indicated by a '-' in the central column.

Parameter	Unit	Notes
Inclination, i	deg	Inclination angle of the orbital plane relative to line of sight. $i = 90$ deg is edge-on.
Filling factor, f	-	The Roche lobe filling factor, defined as the ratio of the Roche lobe radius to the star radius in the direction of the L_1 point.
Distance modulus, DM	-	Defined as $\mathrm{DM} = 5\log_{10}(d) - 5$, where d is the system distance in parsecs.
Mass ratio, q	-	Ratio of the pulsar mass to the companion mass, $q = \frac{M_\mathrm{P}}{M_\mathrm{C}}$.
Companion radial velocity, K_2	km s^{-1}	The radial velocity is usually informed by spectroscopy and can be used to derive the mass ratio in combination with K_1.
Orbital period, P_orb	s	Almost always a fixed parameter, obtained from radio timing.
Semi-major axis, a	lt s	a is also taken as a fixed parameter, obtained from radio timing.
Base temperature, T_0	K	The night side temperature of the companion star.
Irradiation temperature, T_irr	K	The irradiation temperature of the companion star, defined in equation 1.20.
Co-rotation parameter, Ω	-	The ratio of the companion spin frequency to the orbital frequency. In almost all cases this is fixed at $\Omega = 1$.
Gravity darkening, T_grav	-	For the redbacks with convective envelopes, we use $T_\mathrm{grav} = 0.25$ (Lucy, 1967).
V-band extinction, A_V	-	This parameter encodes the effects of interstellar dust, including the interstellar reddening. This parameter is converted to the correct extinction for each band when initialising `Icarus`.

be modelled. The synthesised star is then irradiated using a direct heating model, however in chapter 3 we consider two extensions to this model to account for asymmetric heating. In the directed heating model, introduced in Breton et al. (2012), the pulsar wind is treated as a source of energy that increases the temperature of surface elements on the inner face of the companion, without disrupting thermal

equilibrium. Atmosphere grids are used to determine the corresponding specific intensity for a given temperature for each photometric band. After this, the specific intensity can then be evaluated for each surface element over a range of orbital phases to construct a light curve, taking into account additional parameters like the interstellar reddening and system distance. Note that this may cause issues with modelling as the spectrum of an irradiated surface may not be the same as a stellar atmosphere with an equivalent effective temperature. The effect of this is that individual bands may be shifted up or down in brightness. The band calibration offsets configurable in Icarus help to offset this, however. These band calibration offsets allow each band to individually be shifted up or down in flux, up to some defined limit. This way, the Icarus model can fit for the shape of the light curve while any uncertainties in the calibration in each band are absorbed by the offsets. Note that if the magnitudes are incorrect by more than the calibration offsets, the model will be affected. For example, if the band calibration is incorrect such that the star is too blue, the temperature of the model star will be too hot. This can have knock on effects on the other parameters, due to the degeneracy of the parameter space, such that a hotter star could lead to a more edge-on system, or one with a lower pulsar mass and less full Roche Lobe.

1.6.3 Machine learning

The Random Forest classifier

The Random Forest classifier (Tin Kam Ho, 1998; Breiman, 2001) is a machine learning method based on the concept of a decision tree; a set of nested rules to split a dataset into defined categories. In essence, each tree consists of a series of linear decision boundaries, where the data is split into subsamples at each layer based on the decision boundary of the previous until either a defined maximum depth is reached or there are too few subsamples to split. A maximum depth is almost always used in order to avoid overfitting due to variance; if there are only 2 or 3 subsamples at a layer then these could be fitting noise rather than real trends in the data. Conversely, limiting the trees to a depth that is too shallow results in underfitting due to bias. Decision trees are scalable, robust to unnormalised or qualitative data, and also automatically choose the most relevant features from a set of rules. However, they are often poor predictors and generally overfit the data, resulting in significant variance.

Random Forest classifiers are, as the name implies, an ensemble method constructed of many decision trees which succeeds in overcoming the issues of the lone

decision tree. There are several methods used to achieve this, though the most common (and the method implemented in this work) is bootstrap aggregation (Breiman, 2001). For a forest of N decision trees and some datatset $X = x_1, ..., x_k$ with labels $Y = y_1, ..., y_k$, a N random samples are selected with replacement and fit to the trees. For each $n = 1, ..., N$, tree T_n is trained on m samples from X, Y. Then, predictions on testing samples are obtained from the mean prediction (or majority rule in some cases) of all the trees. This method decreases the variance of the classifier while keeping the bias constant, since the trees are uncorrelated due to the bootstrap aggregation process. A further step is used in the Random Forest algorithm; a subset of the classification features is also randomly sampled for each tree such that dominant features do not cause any correlation between the trees.

In this work we use the `scitkit-learn` `python` implementation of the Random Forest algorithm; table 1.2 shows the key hyperparameters used, as well as a brief description of each hyperparameter.

Table 1.2: Hyperparameters of the Random Forest classifier used in this research.

Hyperparameter	Value	Description
`n_estimators`	700	The number of decision trees used.
`max_depth`	12	The maximum depth of each tree; large values lead to overfitting
`min_samples_split`	10	Minimum number of samples at a tree node required for a split
`max_features`	$\log_2$	The number of features considered when locating the best split.
`bootstrap`	`True`	Perform bootstrap aggregating.

Performance metrics

Table 1.3: Simple model confusion matrix indicating the definitions of various metrics. $A*$ and $B*$ indicate two predicted classes, and A and B represent two true classes.

	$A*$	$B*$
A	True positive (TP)	False negative (FN)
B	False positive (FP)	True Negative (TN)

The most common method used to assess classifier performance is the confusion matrix, with each row representing the true classes and each column the predicted classes as illustrated in table 1.3. For this simple binary classification, we define four terms from this matrix, however these are easily extended to the multi-class problems

in this research. We can see that the ideal confusion matrix is diagonal, such that all samples are either true positives (TP) or true negatives (TN). Classifiers with poor performance produce confusion matrices with significant off-diagonal terms, such that there are many false positives (FP) and false negatives (FN) From these terms we define several metrics which describe the performance of the classifier. The true positive rate (TPR), also known as the recall or sensitivity, is defined as the ratio of true positives to all positives,

$$\mathrm{TPR} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}, \tag{1.28}$$

and the true negative rate (TNR) or specificity is defined as the number of true negatives over the total number of predicted negatives,

$$\mathrm{TNR} = \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FP}}. \tag{1.29}$$

The false positive rate (or fall-out) and the false negative rate are defined as $1-\mathrm{TNR}$ and $1-\mathrm{TPR}$ respectively. Two further useful metrics are the positive predictive value (PPV, or precision) and the negative predictive value (NPV), which are defined as

$$\mathrm{PPV} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}, \tag{1.30}$$

and

$$\mathrm{NPV} = \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FN}}, \tag{1.31}$$

respectively. As above, we can also define the false discovery rate (FDR) and the false omission rate (FOR) as $1 - \mathrm{PPV}$ and $1 - \mathrm{NPV}$ respectively. Lastly, the overall accuracy is defined as

$$\mathrm{Accuracy} = \frac{\mathrm{TP} + \mathrm{TN}}{\sum \mathrm{P} + \sum \mathrm{N}} = \frac{\mathrm{TP} + \mathrm{TN}}{\mathrm{TP} + \mathrm{FP} + \mathrm{TN} + \mathrm{FN}}. \tag{1.32}$$

For multi-class classifiers, these metrics are defined the same way on a per-class basis, though the terms are extracted differently from the confusion matrix. For an N-class classifier, the confusion matrix will be an $N \times N$ matrix, C_{ij} where $i, j \in 1, ..., N$. The true positives are the diagonal terms, such that the true positive of class n will be C_{nn}, where $n \in 1, ..., N$. The false positives of class n are the terms in column n excluding C_{nn}, and similarly the false negatives are the terms in row n, excluding C_{nn}. The True negatives are the remaining elements. The above metrics can then be obtained using matrix algebra.

It should be noted that these metrics on their own are often an incomplete

description of the performance of the classifier, and an assessment alongside the confusion matrix is often required. For example, if there is a large imbalance in the populations of predicted classes (for instance, if one class has zero predictions), some of these metrics become misleading.

Chapter 2

Instrumentation, observations, and simulations

2.1 Introduction

In this chapter, we cover three key facets of optical astronomy. First, we give a technical overview of the primary telescopes and instruments that have been used for this research. The bulk of this work uses photometry data from either the ULTRACAM instrument and the New Technology Telescope (NTT), or the Gravitational-wave Optical Transient Observatory (GOTO) facility. Photometry from the former provides the basis for chapter 3, while the latter is discussed extensively in chapter 4.

Following this, we discuss the data reduction procedure for ULTRACAM photometry, including calibration of the light curves. Due to the high signal-to-noise ratio (SNR) of the ULTRACAM light curves, this is an exceptionally important step in the process. As discussed in chapter 3, small errors in calibration can have a substantial influence on the modelling. For instance, if one band is poorly calibrated with respect to the others, the colour information will be incorrect which can result in poor estimation of the system temperature. When modelling with `Icarus`, this can have dire consequences as the colour information is essential for reducing the degeneracy of the model.

This section also includes a discussion of the means by which uncertainties are estimated for the observations, and an overview of a scheduling code. This scheduling code was first used during an ESO P100 (Autumn/Winter 2017) observing run undertaken in October 2017. Lastly, selected targets and calibrated light curves from this run are presented in this section.

The final section in this chapter discusses the light curve simulations used during this research. we introduce the procedure by which the light curves are simulated;

first by estimating the signal-to-nose ratio (SNR) of the observation of a source of a given magnitude, then building the light curve based on a toy model. The SNR estimate depends on a number of factors, some intrinsic to the telescope and instrument and others dependent on the observing conditions. Some of the extrinsic factors can be controlled for with scheduling, such as the moon phase. The initial SNR estimate is useful in determining the quality of observations that can be expected of a spider candidate with a certain telescope configuration; for instance, will observations of a redback with mean R-band magnitude $R = 22$ give an acceptable SNR if observed with ULTRACAM on the NTT in dark conditions? If so, can bright moon conditions be used instead? These considerations are useful when applying for telescope time.

Following this, we discuss the method used for building simulated light curves and two example use cases. The first, for the high-SNR light curves obtained for dedicated photometers like ULTRACAM, is used to estimate how accurately parameters like the pulsar mass or inclination can be determined through modelling. Rather than a full treatment with `Icarus`, we instead attempt to recover the ellipsoidal and irradiation component amplitudes of the light curve, with uncertainties, which are then used to estimate the precision to which selected derived parameters can be determined. This work was used frequently to provide technical evidence for several telescope proposals. The second use case, for the irregularly sampled light curves of telescopes like GOTO, is a foundation for the work discussed in chapter 4. Here, we discuss how simulated light curves are used to determine how well the periods of variable sources can be recovered, and how it depends on factors like the number of observations or the SNR. The work presented in this section was also adapted to stand as technical evidence in a telescope proposal.

2.2 Instrumentation and data analysis

2.2.1 Instrumentation

ULTRACAM

In this section we follow Dhillon et al. (2007) to detail 'ULTRACAM', a three-band optical photometer with very high time resolution that is frequently used in optical studies of spider pulsars. As well as this, it is extensively used for studies of XRBs, gamma ray bursts (GRBs), cataclysmic variables (CVs), and expolanets, to name a few. The instrument is capable of frame rates of up to 500 Hz, and even at this frequency it can provide simultaneous images in each of the three bands. It achieves this using so-called frame transfer CCDs ('charge coupled device'; the dominant imaging

chip in optical astronomy, with fast read-out and low noise), which allow for almost instantaneous read-out and negligible dead time between exposures. A separate CCD is used for each band, with the light split into red, blue, and green channels using two dichroic beam-splitters. Before reaching the CCDs, these channels each pass through a filter. Most commonly used are the filters of the SDSS (Sloan Digital Sky Survey) photometric system (Fukugita et al., 1996) with colour bands u', g', r', i', and z', which span the wavelength range ~ 3000 Å (around atmospheric cutoff) to ~ 11000 Å (around the sensitivity limit of the CCDs). The transmission curves for these filters are shown in figure 2.1. These filters are used because they have consistently high transmission, approaching 100% at longer wavelengths, and very little overlap between the bands. In 2013, these filters were replaced with the so-called Super SDSS filters, which have even higher transmission rates and less overlap.

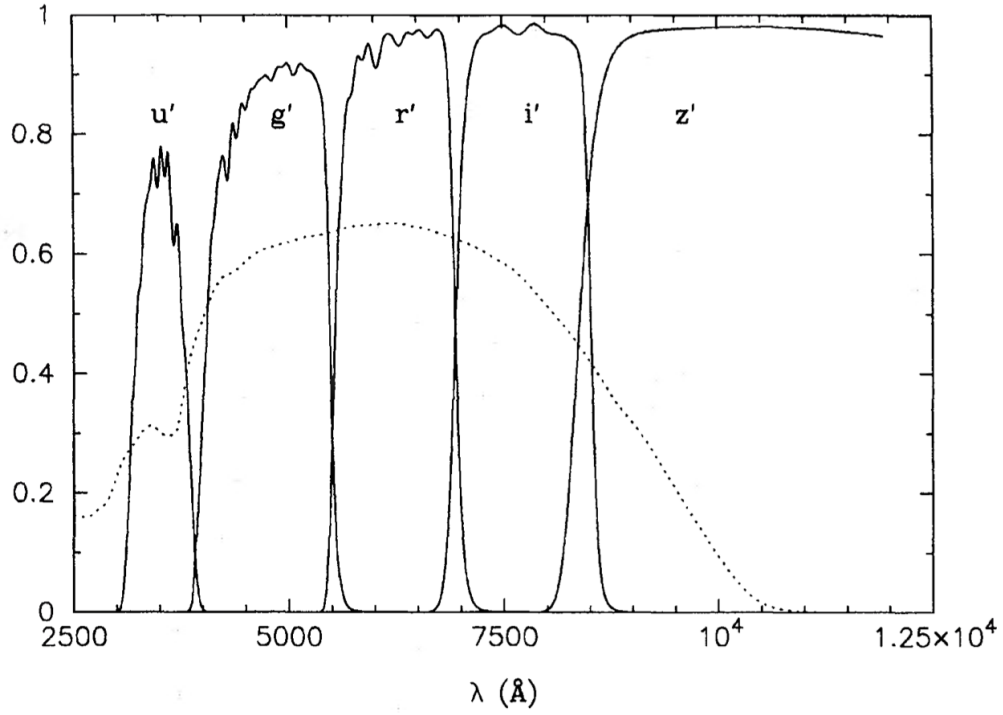

FIGURE 2.1: Transmission curves for the SDSS photometric system u'g'r'i'z' filters. The additional dotted line shows the quantum efficiency of the CCD detector used in the definition of the SDSS photometric system. Image credit: Fukugita et al. (1996).

ULTRACAM has a CCD resolution of 1024×1024 pixels, such that the resolution is 0.3 arcsecs per pixel at a field of view of 5 arcmins. This excellent instrument resolution means that image quality is in all cases seeing-limited for ground-based telescopes. In fact, the two halves of each CCD chip are read out by separate sets of electronics, known as read-out channels), with separate bias and gain. Frame rates

between 500 Hz and 0.05 Hz have been achieved, with an absolute timing accuracy of ~ 1 ms. This range of exposure times means that sources with magnitudes 8 to 26 have been observed: this ability to observe very faint stars is essential in optical studies of spiders as black widows often have optical minima of less than magnitude 25.

The CCDs are cooled to a temperature of 233 K, such that they deliver a dark current of roughly 0.05 e^- pixel^{-1}s^{-1} which is sufficiently small such that even in pessimistic conditions this noise level is negligible. Observations under a new moon, at a dark site, in u' band, with a small aperture telescope give sky noise of around 0.3 e^- pixel^{-1}s^{-1}, for comparison. As such, the noise when using ULTRACAM is sky-dominated for exposures of around 10 seconds or longer, i.e. all cases in this work. However, there are times when a bright source in the frame will cause significant background noise due to diffraction. In these cases a focal plane mask can be used to obscure these sources to further maximise signal-to-noise.

HiPERCAM

Dhillon et al. (2016) describes a new instrument, HiPERCAM: a five-band, high-speed optical photometer. Again based on ULTRACAM, HiPERCAM is capable of imaging five optical bands (SDSS u', g', r', i', z') simultaneously at frame rates of up to 1.6 kHz and exposures as long as 1800 s. Technical improvements have been made in almost every area; for example, the dark current is almost negligible, up to 60 times smaller than with ULTRACAM. First light took place 2017. While photometry from this instrument was not used in this research, the light curve simulation code discussed in this chapter was applied to the configuration of this instrument

Telescopes

Here we briefly introduce a few of the key telescopes our team commonly employs for photometric studies of spiders. These telescopes must have a few key features in order to be useful for this research, such as sufficient size to obtain the required sensitivity, or compatibility with the required instruments like ULTRACAM. The New Technology Telescope (NTT) is a 3.58 m, Richey-Chretien telescope on an Alt-Az mount, located 2375 m above sea level at La Silla observatory in Chile. As such, the observing conditions are excellent all year round, though strong winds can sometimes be a problem. The telescope uses active optics to drastically improve image quality, with a resolution of 0.15 arcsec across its 30 arcmin field-of-view.

ULTRACAM is often mounted on the NTT. The Thai National Telescope (TNT) is a 2.4 m Richey-Chretien telescope on an Alt-Az mount, located at the Thai National Observatory in Thailand. It has a field of view of 16 arcmin. Lastly, the Gran Telescopio Canarias (GTC) is a 10.4 m Richey-Chretien telescope on an Alt-Az mount, located at the Roque de los Muchachos Observatory on the La Palma island (Spain), 2267 m above sea level. Lastly, the William Herschel Telescope (WHT) is a 4.20 m Richey-Chretien telescope also located at Roque de los Muchachos Observatory in the Canary Islands. Both the WHT and GTC boast a median seeing of 0.7 arcsec, and clear skies for around 75% of the year. Over the course of this research we have proposed observations on all four of these telescopes, although the sources discussed in this research were all observed with ULTRACAM + NTT.

In addition to these large area telescopes, we also use data from the GOTO (Gravitational-Wave Optical Transient Observer) project. The aim of this project is to identify optical counterparts to gravitational wave events, such as those detected by LIGO. It accomplishes this using a set of 8 small, identical optical telescopes on a single mount in order to obtain a wide field-of-view of up to 40 deg^2. Note that this the current configuration as of March 2021; at the start of the research GOTO consisted of only 4 telescopes on a single mount and a corresponding field of view of $\sim$ 20 deg^2. It has the capacity to use either R, G, B, or wide-band L filters (Steeghs, 2017; Dyer et al., 2020), though the data we have collected uses only the wide-band L filter. Using this broader filter, deeper magnitudes can be reached: up to mag 20.5 with a 180 s exposure. While the main aim is the identification of optical counterparts to events like neutron star/black hole mergers, it will also act as an all-sky survey. Chapter 4 describes in detail our work using this catalogue to search for and classify variable sources with light curve characteristics that match those of spiders, and 2.4.2 introduces simulations of GOTO photometry.

2.2.2 Data reduction

The ULTRACAM pipeline

The ULTRACAM reduction pipeline uses comparative aperture photometry, where the flux from the target source in each exposure is compared to the flux of some set of comparison stars, assumed to be of constant brightness. Through this, any variation not intrinsic to the target source, such as atmospheric variation, is removed. Note that each step in the reduction occurs simultaneously for each of the three CCDs. This process is covered in detail for the sources PSR J1227–4853 and PSR J1023+0038 in chapter 3, but here we will describe the general calibration and

reduction process. An example ULTRACAM frame from 2019 observations of PSR J1227–4853 is shown in figure 2.2.

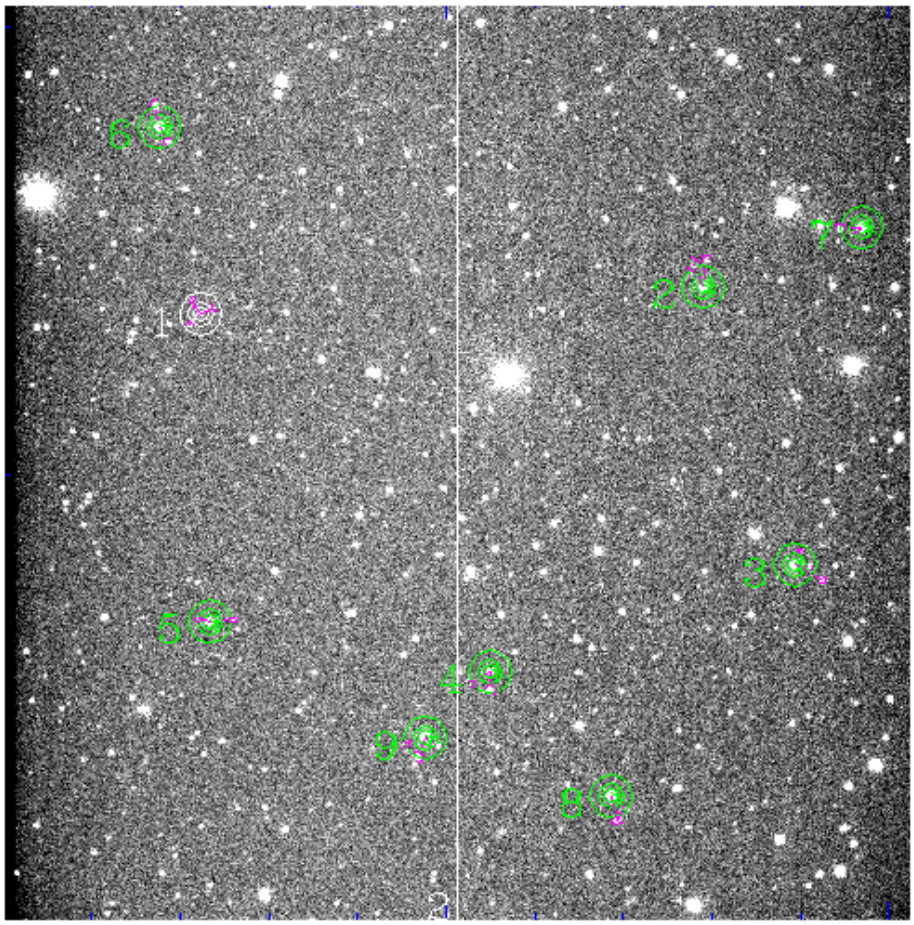

FIGURE 2.2: ULTRACAM field of PSR J1227-4853. The target aperture is shown in white, while the comparison stars are in green. Note the central vertical line; this is the division between the two readout channels.

Before reduction, the images are calibrated through the use of bias, flat, and dark frames (as opposed to the so-called science frames, which contain the source images). Dark frames are the least important type of calibration frame for ULTRACAM, as the extremely low dark current is negligible even for u'-band exposures, as mentioned in section 2.2.1. However, the use of flat field images (which correct for optical non-uniformities such as vignetting and non-uniform pixel responses on the CCD) and bias frames (which remove the offset voltage applied to pixels to ensure electron counts are not negative) is essential for achieving properly calibrated images. Dark frames are captured with the telescope dome closed and an opaque slide covering the instrument aperture. Ideally, these frames should have the same exposure time as the science images. Bias images are also captured in the same way, but have

the shortest possible exposure time (ideally zero seconds). Flat field images are captured at twilight, with the dome open and instrument exposed as usual for science frames. However, the sky at twilight is approximately uniform, such that any non-uniformities in the telescope or instrument optics are revealed. This usually includes vignetting from the telescope or dome and any imperfections in the mirror. Flat field and dark frames all must be calibrated to produce *master flats* and *master darks*. Master dark frames are a bias-subtracted mean of a large number of dark frames, and master flats are a bias- and dark-subtracted mean of a large number of normalised flat frames. Note that the dark frames must be scaled by the ratio of exposure times of science to dark frames. From these, science frames are calibrated via the process

$$\text{calibrated science} = \frac{\text{raw science} - \text{master bias} - \text{master dark}}{\text{master flat}}, \tag{2.1}$$

on a per-pixel basis. However, if the dark frames have the same exposure times as the science frames, then subtraction of a bias frame is not necessary as the master dark will also contain the bias.

With the calibration complete, the reduction apertures can be set: one for the source and at least one for the comparison stars used to calibrate the flux. It is desirable to use the maximum number of comparison apertures possible. These apertures are composed of three concentric circles: one to contain the source, and the other two corresponding to outer and inner sky apertures to measure the sky flux. Sections of these aperture regions can be masked if there is an unwanted source close to a target or comparison source. Figure 2.3 illustrates these apertures and masks. Due to the way apertures are tracked across frames, it is common for it to jump to an unmasked, unwanted source if the target object is at optical minimum. As such, apertures can also be linked together such that the pixel vector between them is kept constant. Finally, during reduction, the number of pixels and the total number of counts in those pixels are recorded for each aperture for each exposure, of the calibrated science images. An example graphical output from the ULTRACAM pipeline is shown in figure 2.4. This output displays the light curve of each reduced band, as well as several other quantities. The x and y pixel positions of the target aperture centroid are displayed as the difference from those of the first observation: keeping track of this is important during observation to avoid tracking errors. The transmission percentage and FWHM (which represents the seeing in each band) are also displayed, useful for identifying unfavourable conditions. As well as this, a `.log` file containing the reduced data is created. Each line of this file represents an

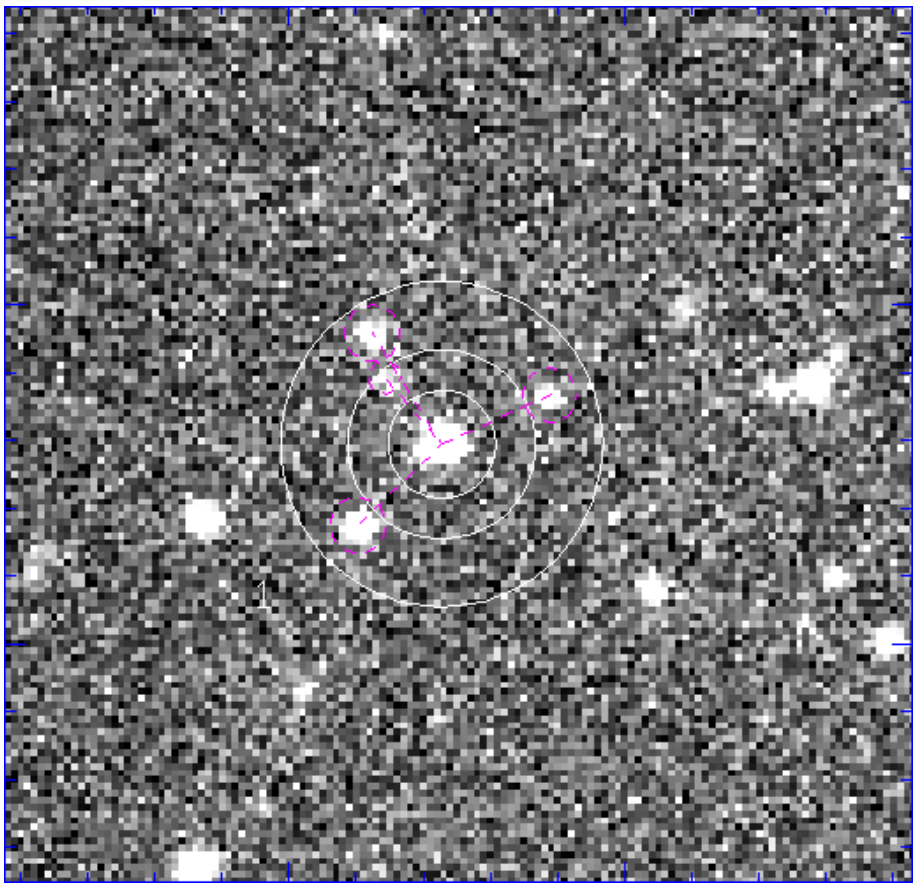

FIGURE 2.3: Section of figure 2.2, cropped around the target source. The concentric white circles represent the inner target and outer sky apertures, while the smaller pink circles mark pixels which have been masked. This is done to ensure the background count isn't contaminated by other sources.

exposure from one of the CCDs, such that each line contains the MJD (modified Julian date) of the exposure, the exposure time, in seconds, the number of object counts in the target aperture, the number of background counts in the sky aperture, 1-sigma uncertainty on the number of object counts, and the number of sky pixels available.

2.2.3 Calibration of photometric light curves

PyCam

A photometry reduction code, `PyCam` was written to obtain flux-calibrated light curves from the reduced ULTRACAM `.log` files. However, its use was superseded by a standard code used by the research group. Here we detail the procedure used in `PyCam`, though the principles used by both codes are the same.

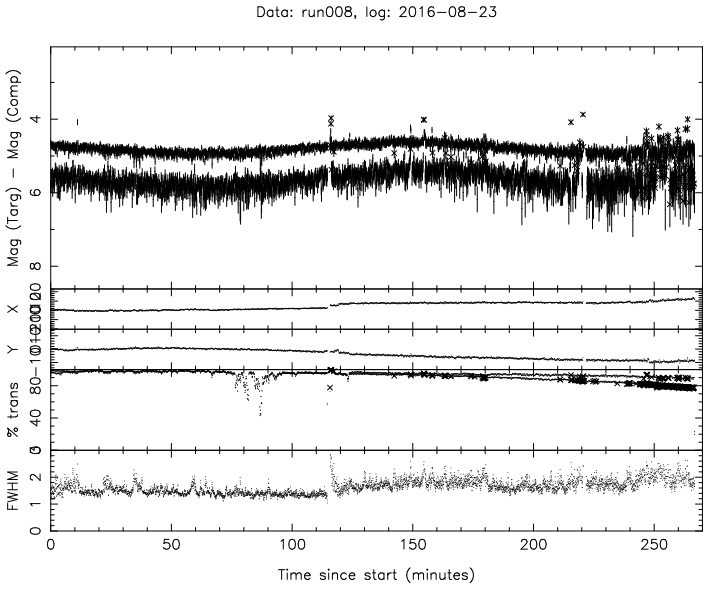

FIGURE 2.4: Graphical output of ULTRACAM reduction pipeline. From top to bottom: light curve, x-position, y-position, transmission %, FWHM (i.e. seeing). This is for source PSR J1628-3205, faint in the u'-band, such that reduction was only conducted for the i'- and g'-bands.. Note that the seeing becomes significantly variable from around 120 minutes onwards: the drastic impact this has on the data is clearly visible

Reduction, timing, and flux calibration The target and comparison star magnitudes and photon fluxes are related by the equation

$$m = m_{comp} - 2.5 \log_{10} \left(\frac{F}{F_{comp}} \right), \tag{2.2}$$

where m and m_{comp} are the magnitudes of the target and comparison stars respectively, and f and f_{comp} are their observed fluxes. Note that the particular units of these fluxes are not important as the ratio is calculated, assuming both were obtained under the same conditions: this is indeed the case for our observations. The magnitudes of the companion stars, m_{comp}, are obtained from all-sky surveys. Commonly used are the Sloan Digital Sky Survey (SDSS) or the AAVSO (American Association of Variable Star Astronomers) Photometric All-Sky Survey. This is performed separately for each band, for whichever filter was in the red, blue, and green arms of ULTRACAM (generally u', g', and i'/r'). These calculated magnitudes are stored along with the corresponding modified Julian date (MJD) of observation. However, it is generally more useful to consider time in terms of orbital phase. Before calculating the phases, it is necessary to correct to the SSB (solar system barycentric) time from the local time. This is done using functions in the `astropy`

package which implement the calculations detailed in Lorimer & Kramer (2004). Here, it is done by calculating the light travel time to the SSB and simply adding that to the raw MJD times, as well a correcting to the Barycentric Dynamic Time (TBD). Then, these SSB-corrected times are converted to orbital phases, ϕ. This is done using the known orbital period, P_{orb}, of the source and the known ascending node, $T_{\rm asc}$, via the calculation

$$\phi = (t_{\rm MJD} - T_{\rm asc}) \mod P_{\rm orb}, \tag{2.3}$$

where mod is the modulo operator. It is also possible to incorporate the orbital period derivative in this calculation, though this is rarely necessary given the observation timescales. The orbital period and $T_{\rm asc}$ are obtained from radio timing, though can also be informed by γ-ray timing. We employ the convention here that phase 0 occurs at the ascending node of the pulsar, such that inferior conjunction occurs at phase 0.25.

In cases where the comparison magnitudes cannot be obtained - typically this occurs for u' as there are fewer surveys which use this band - a zero point calibration must be used. This method uses the known zero point magnitude of the telescope and instrument in combination with the photon flux measurements and atmospheric extinction coefficient for each band to estimate the magnitude of the source. There are several shortcomings of this method, however. First, the extinction coefficient can vary significantly from night to night depending on conditions. Secondly, the presence of dust and dirt on the telescope mirror and optics can change the zero point by up to 0.5 mag in the case of ULTRACAM. As such this method is less accurate in general, but provides an acceptable magnitude calibration where the standard method is not possible.

Uncertainties The mean magnitude and uncertainty are estimated by calculating the signal-to-noise ratio of the source. First, the signal-to-noise ratio (SNR) is calculated using the equation

$$\mathrm{SNR} = \frac{N_{\rm star}}{\sqrt{N_{\rm star} + n\,(N_{\rm sky} + N_{\rm dark} + N_{\rm RON}^2)}}, \tag{2.4}$$

where $N_{\rm star}$, $N_{\rm sky}$, $N_{\rm dark}$, and $N_{\rm RON}$ are the photon counts per CCD pixel from the star, the sky, the dark current, and the read-out noise respectively, and n is the number of sky pixels. However, it is necessary to account for the conversion between

photon and electron counts, so we include a factor for the gain, G,

$$\mathrm{SNR} = \frac{N_{\mathrm{star}}G}{\sqrt{N_{\mathrm{star}}G + n\left(N_{\mathrm{sky}}G + N_{\mathrm{dark}} + N_{\mathrm{RON}}^2\right)}}, \tag{2.5}$$

where our N are now electron counts. For ULTRACAM, the gain is 1.2 e^-/ADU for both *slow* and *fast* modes, and the read-out noise for these modes is ~ 4 e^- and ~ 9 e^- respectively. For ULTRACAM, the dark current is 0.1 $\mathrm{e}^-\mathrm{pixel}^{-1}\mathrm{s}^{-1}$, which is negligible compared to the photon rate from the sky even in poorest conditions, so we can ignore it without penalty. It is also worth mentioning that the read-out noise is small compared to the sky noise such that ULTRACAM is sky-dominated in all cases in this work. So, the SNR equation simplifies to the form

$$\mathrm{SNR} = \frac{\sqrt{G}N_{\mathrm{star}}}{\sqrt{N_{\mathrm{star}} + n\left(N_{\mathrm{sky}}\right)}}. \tag{2.6}$$

Conveniently, all of these counts are in the ULTRACAM `.log` file. As such, the SNR for the target and comparison stars for each exposure can be calculated. The next step is translating this SNR to an error on the calculated magnitudes. From the magnitude equation (equation 2.2), we write the uncertainty as

$$\sigma_m = \frac{\mathrm{d}m}{\mathrm{d}F}\sigma_F. \tag{2.7}$$

We differentiate the $\log_{10}$ using the following identity,

$$\log_{10}(F) = \frac{\ln(F)}{\ln 10} \quad \longrightarrow \quad \frac{\mathrm{d}m}{\mathrm{d}F} = \frac{-2.5}{\ln(10)F} \approx \frac{1.086}{F}, \tag{2.8}$$

such that the error is

$$\sigma_m \approx 1.086\frac{\sigma_F}{F} \approx \frac{\sigma_F}{F} \approx \frac{1}{SNR}. \tag{2.9}$$

That is, the error on the magnitude is roughly the reciprocal of the SNR. Thus, the weighted mean is taken of the magnitudes calculated using each comparison star, and the uncertainty on this mean magnitude is taken as the standard error on the weighted mean. At this point it is important to note that we require the uncertainty on magnitudes to be better than approximately ± 0.3, as this corresponds to a SNR of 3, the minimum permissible significance of a detection. Below this we cannot separate our detection from the background, and as such these observations are usually discarded from the final light curves used for modelling. Note that fluxes are typically used in `Icarus`.

Colour terms

The ULTRACAM filters and the SDSS (and SSDSS) photometry system are not identical, so the passbands for a given filter from each will be slightly different. For example, the passbands may be slightly different widths, or have tails of different shapes. The result of this is that light from stars will pass through differently. A wider passband will allow photons through than a narrower one, which is not an issue if all stars are affected equally as the relative magnitude is calculated. However, since different stars have different spectra (and indeed some variables stars change colour during their cycles), this will cause a systematic error in the magnitudes of some stars. For example, if one ULTRACAM filter passband extends slightly farther towards shorter wavelengths, blue stars seen through this filter will appear brighter than red stars compared to the base SDSS system.

Corrections to account for this, known as colour terms, are applied during the magnitude calibration which transform the magnitudes to the SDSS system. Colour terms are determined by observing fields containing several standard stars with known magnitudes and colours. The measured and standard magnitudes are compared as a function of colour, with first order colour terms appearing as linear trends in the residuals. The colour terms between the ULTRACAM filters and the SDSS system can be up to 0.3 mag, so it is essential that these transformations are performed.

2.3 Observational work

2.3.1 Scheduler

An essential part of optical astronomy is scheduling observations: given a list of potential targets it is necessary to schedule where the telescope should be pointing and at what time. This must take into account both the priority of a target and the feasibility of observing that target on a given night, including considerations of the moon phase and the expected airmass at the source position. A brief note about the moon phase; we define three terms which describe the impact of light from the moon on observations. These are 'bright', where the light from the moon is a significant contribution to the sky brightness at 100% illumination, such as when full or close to the target source, 'dark', where the moon is either below the horizon or new at 0% illumination, or 'grey' where it is around 50%. Ideally, all observations would be performed in 'dark' time, however the observing time is naturally more competitive than bright time.

As part of the preparation for observing, a scheduler was written in Python. Given a source position, observation location, and observation date, the visibility of that source is calculated. Visibility here refers to the airmass along the source pointing, the path length that light must travel through the atmosphere, as this is a major factor in the quality of the observation. Sources observed with a high airmass, near the horizon, will be much more strongly attenuated than those observed at the zenith (the point directly overhead) and so appear significantly fainter. It is common for an airmass of 2 to be taken as the maximum permissible airmass. For comparison, the airmass at the zenith is 1 (such that the airmass at other inclinations is measured relative to this), and the airmass at the horizon is typically in the range 35 to 40. There are many formulae for calculating the airmass at a given zenith angle, z, with some empirical formulae yielding accurate results even very close to $z = 90°$. However, for the purposes of this scheduler, where sources with airmass above 2 are excluded, a simpler formula can be used. This corresponds to a maximum zenith angle of roughly $z = 60°$, via the approximation

$$\text{airmass} \approx \sec z, \tag{2.10}$$

which is accurate to around $z \sim 75°$. This approximation assumes a plane atmosphere of constant density on a flat Earth, however the approximation is of course valid for a round Earth.

Briefly, the software works as follows. The telescope to be used, observation dates, and a list of sources are read from either hard-coded lists, user input, or online catalogues. The latitude, longitude, and height of the telescope are calculated. The list of sources includes the name, *ra-dec* position, and observation priority. RA (right ascension) and Dec (declination) are the coordinates of the equatorial coordinate system. From the observation dates, the sunset and sunrise times are calculated for each night at the observing location. For example, the sunset time is found by dividing each day into hour divisions, calculating the position of the Sun in the sky at each time, then iterating over these until the Sun's position drops below the horizon. Then, this is repeated for each 5 minute interval in the hour before the first time at which the Sun is below the horizon. The sunrise time is calculated in a similar way, but instead for when the sun first rises above the horizon. By this method the sunset and sunrise times are calculated to 5 minute-accuracy. This level of accuracy is roughly what is required, since re-pointing the telescope and beginning observations takes on average 5 minutes.

Following this, each night (between consecutive sunset and sunrise times) is divided into 5-minute intervals. For each source, the zenith angle is calculated at

these times and the times at which it passes above and below 60° from zenith, as well as the minimum zenith angle, are recorded. It is necessary to run the calculations in parallel, since for a large number of sources this process can take some time. The angle between the source and the moon is calculated and recorded, as well as well as the sunset, sunrise, and twilight times. Twilight is defined as the time at which the sun is at a zenith angle of 108°. Finally, this information is plotted, and the result can be seen in figure 2.5.

For example, we see that PSR J1023+0038 is an ideal candidate, as the entire 4.75 hour period (Archibald et al., 2009) can be covered in one observation and the source is positioned sufficiently far from the moon. Conversely, sources like PSR J1628–3205 are only visible in acceptable conditions for a few hours.

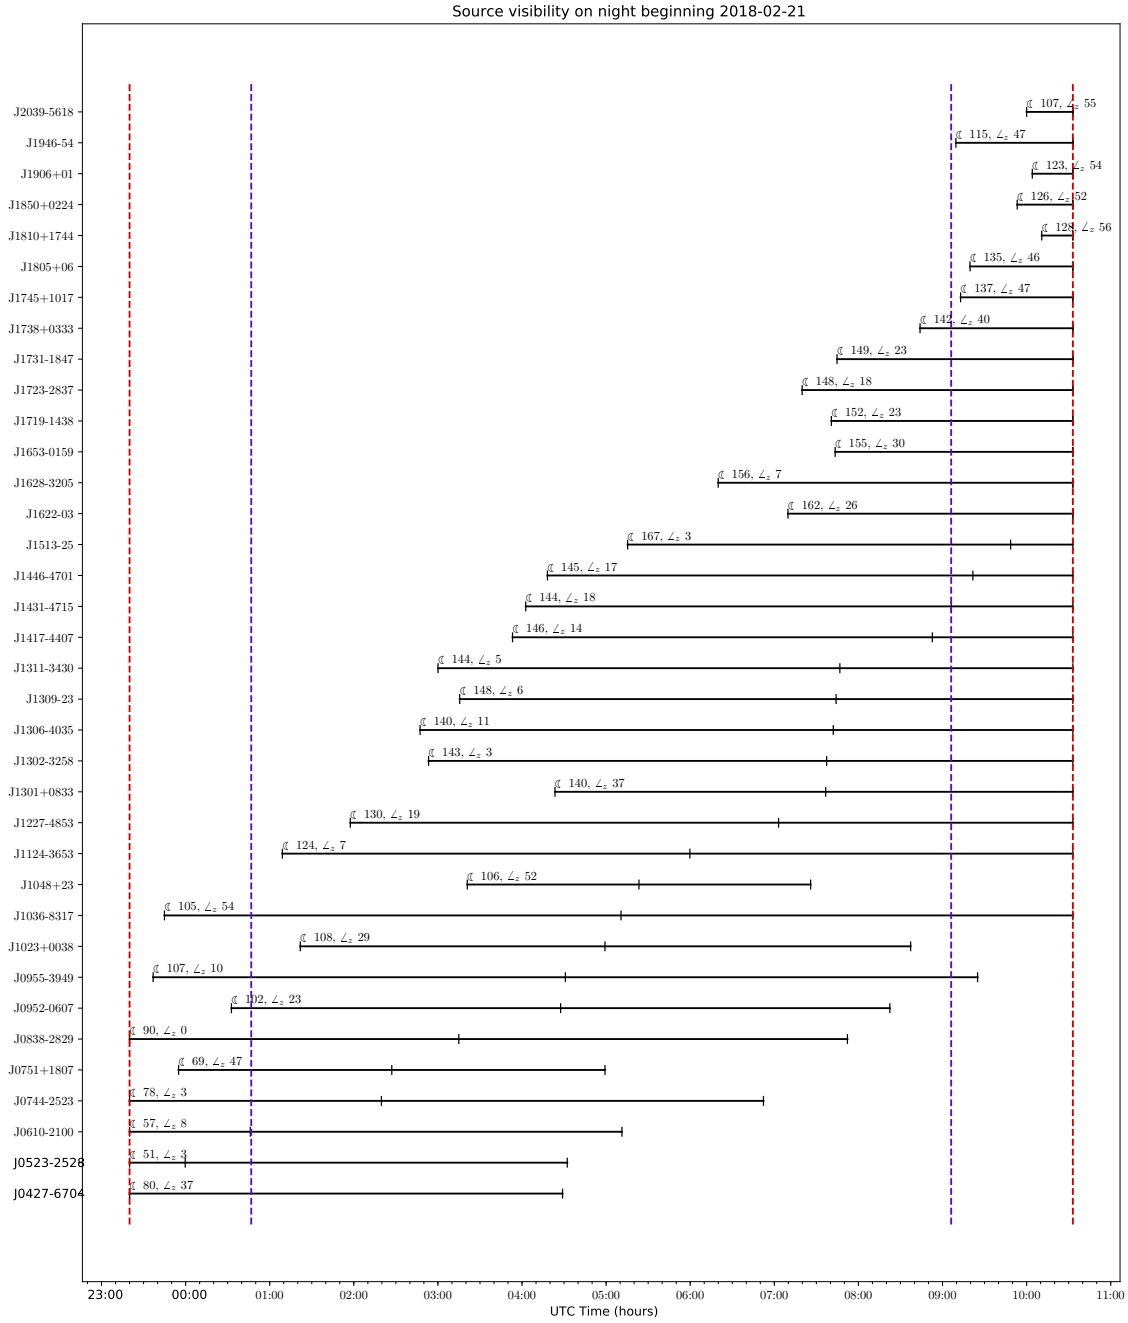

FIGURE 2.5: An example schedule created with the Scheduler, for my 23rd birthday, 2018-02-21. The red lines are sunset and sunrise and the blue lines are twilight. The angle to the moon and zenith are labelled. The three points on each line indicate the time the sources 'rises' below 2.0 airmass, its minimum zenith angle, and the time it 'sets' below 2.0 airmass.

2.3.2 ULTRACAM on the NTT: October 2017

Observing with ULTRACAM

In October 2017, we were awarded a 5 night-long block of observing time during the ESO (European Southern Observatory) P101 observing period. The observations took place between the 13th of October and the 17th of October, inclusive, at the La Silla Observatory, Chile. These observations used the ULTRACAM instrument, introduced in section 2.2.1, with the NTT. The primary aim of this observing mission was to search for and obtain photometry of the optical companions in spider systems. The first of these aims was accomplished by 'fishing'; simply obtaining exposures of the target location in order to first identify if a companion is present, and if so to measure the colour and magnitude information. Repeated, short observations over the run can be used to build up a light curve. The second of these aims is accomplished by capturing long blocks of exposures to build up a light curve over the orbit. However, two and a half of the nights we were unable to observe due to strong winds forcing closure of the telescope dome on the first night and heavy cloud cover on the final two. As such the observing program was adapted from a total of 10 targets to only five; 2 were fishing targets, and 3 were high priority sources.

P101 targets

PSR J2241–5236 This is a MSP in a black widow system, detected as a Fermi γ-ray source and follow-up with the Parkes radio telescope (Keith et al., 2011). It is in a 3.5 hour orbit, and there have been no previous optical detections. As such, this was a 'fishing' target. We detected an extremely faint optical counterpart at the radio timing position, by stacking 100×11 second exposures.

3FGL J2017–16 This is a candidate black widow pulsar, first detected as a Fermi γ-ray source (Lyapin et al., 2017). It has an orbital period of 2.3 hours, and there is an optical counterpart with i-band magnitude 20 (Sanpa-arsa, 2016). However, the pulsar position is not known publicly. As such, this was a fishing target, with the aim that a detection would allow for future observations. We detected a source at the Fermi position, but with only a single observation this source cannot be tested for variability.

3FGL J0427.9–6704 This is an eclipsing binary pulsar in an accreting state with an orbital period of 8.88 hours, an I-band magnitude of approximately 16.8 out of eclipse and approximately 18 in eclipse (Strader et al., 2016a). We observed this source in the (u', g', i') bands with a cadence of approximately 10 seconds over

2 nights. Instead of the characteristic two-component light curve, we instead saw constant flaring and flickering in all three bands, indicative of an accretion disc.

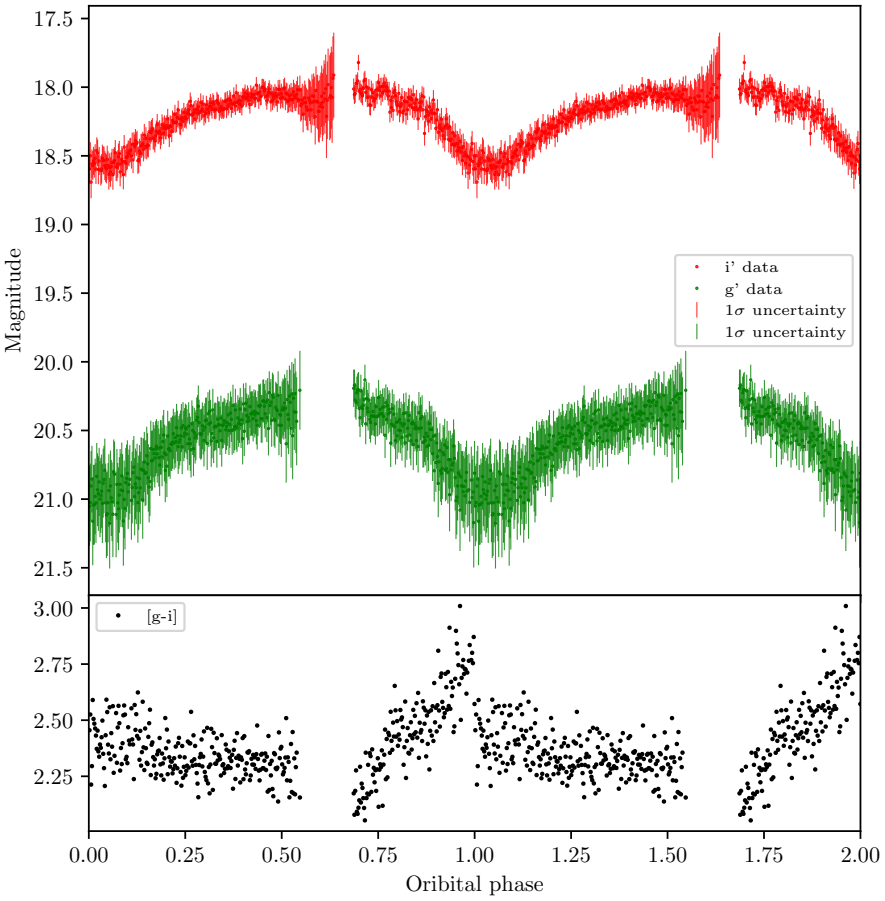

FIGURE 2.6: *Top:* Phased i'- and g'-band light curves of 3FGL J0744-2523. These data were reduced using the ULTRACAM pipeline and PyCam (introduced in section 2.2.3). The incomplete coverage around phase 0.6 is a result of PyCam eliminating data points with a signal-to-noise ratio of less than 3. The data here have significantly larger errors than those at other points in the orbit. This is due to high winds at the time of observation. *Bottom:* g'-i' colour index, in phase with the above light curve. Note the sharp discontinuity at phase 0/1. This discontinuity appears to be physical as it occurred during a single observation run, and 5 comparison stars were used in reduction. The uncertainty on the colour index is ±0.2.

3FGL J0744.1–2523 This was discovered as a possible eclipsing binary MSP from an unassociated Fermi source by Salvetti et al. (2017), with an orbital period of 2.76 hours. The nature of this source is not certain, as the lack of published spectroscopy means that we cannot confirm that the primary is not optically bright. If this were to be the case, it instead suggests that this source is a W Ursae Majoris-type (W UMa) binary, with twice the current orbital period. W UMa-type stars are contact binaries

with a common envelope that exhibit double-peaked light curves (Lucy, 1968). We obtained slightly over one orbit of photometry data in SDSS u', g', and i' bands with a cadence of approximately 20 seconds. However, a significant part of the coverage is poor due to strong winds on one of the nights. A flux-calibrated light curve can be seen for this source in figure 2.6, showing the patch of poor data. This light curve shows some clear deviations from the typical quasi-sinusoidal shape, notably around phase 0.4 and phase 0.8. It is unclear if these features are physical, however they display similarities to the asymmetries displayed in the tMSPs presented in chapter 3. Calculating the $g' - i'$ colour index reveals an interesting feature around phase 0: there appears to be a discontinuity in colour index coincident with the optical minima. If this discontinuity is a constant, and indeed real, orbital feature, it could be used as a marker to determine the true orbital period of this optical source with additional photometry. Recently, the Fermi gamma ray source was associated with an isolated neutron star, suggesting that this optical source is not a binary MSP and is likely a W UMa binary instead (C. Clark, private communication).

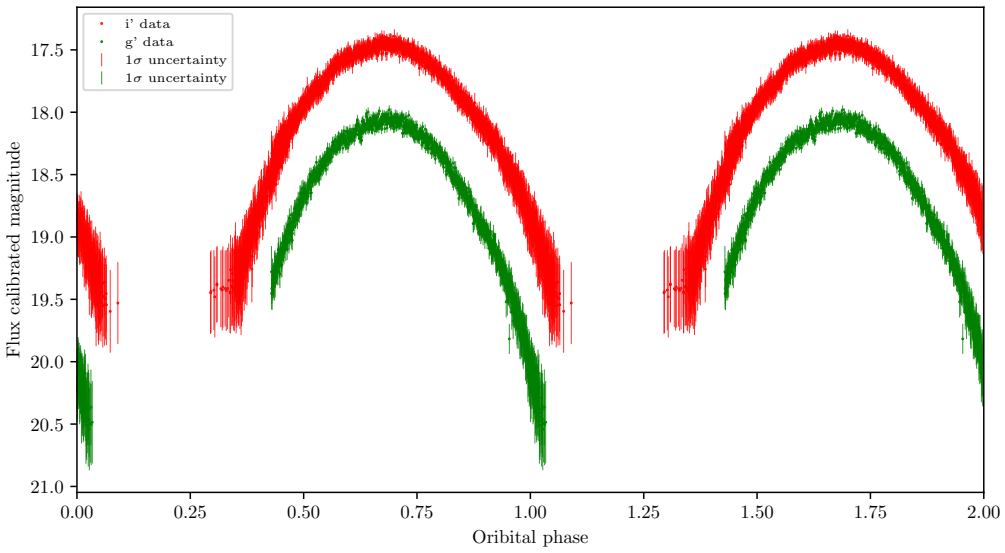

FIGURE 2.7: Phased i'- and g'-band light curves of PSR J2339-0533. These data were reduced using the ULTRACAM pipeline and PyCam. The incomplete coverage around phase 0.25 is a result of PyCam eliminating data points with a signal-to-noise ratio of less than 3.

PSR J2339–0533 This is a millisecond pulsar in a redback system, discovered from an unassociated Fermi source by Ray et al. (2014) and further analysed by Pletsch & Clark (2015). The source has an orbital period of 4.6 hours. We obtained over one orbit of data in SDSS u', g', and i' bands with a cadence of approximately 8 seconds. The sources showed strong variability in all three bands, however only the brighter i' and g' bands were reduced. Figure 2.7 shows the reduced light curve of

this source from our observations. This source was too faint to detect at its minima in u'-band, and the presence of a very bright star close to the ULTRACAM field of view created a brighter than usual background. This background star is the likely cause of the 'flattening' of the i'-band data around phase 0.3.

2.3.3 Other Sources

This section covers J0427 in more detail, and introduces data other than that gathered in the P101 October run. Three sources have been the main focus; PSR J1628-3205 (hereafter J1628), the previously-mentioned 3FGL J0427.9-6704 (hereafter J0427), and the redback PSR J1306–40 (hereafter J1306).

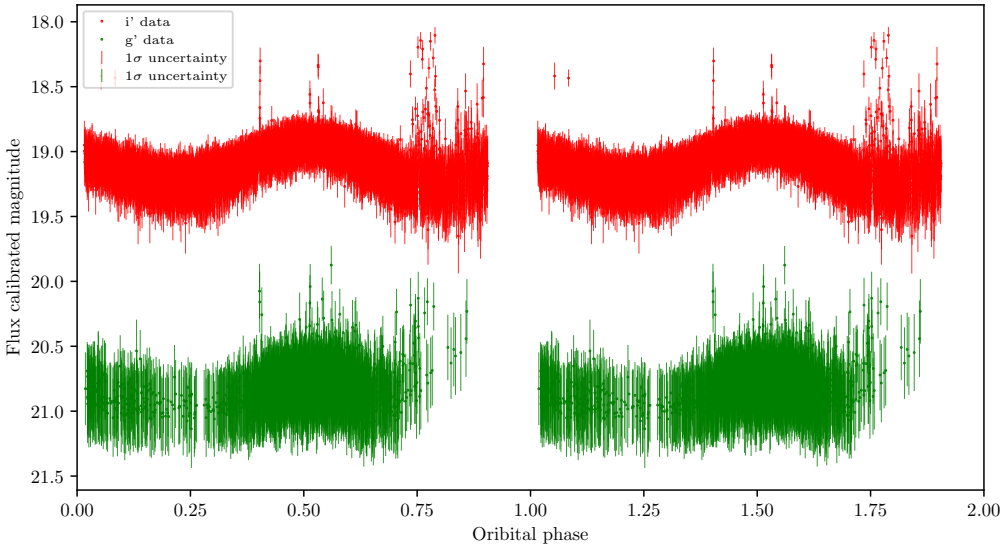

FIGURE 2.8: Flux and phase calibrated light curve of PSR J1628-3205, using data from the nights of 2016-07-06, 2016-08-23, and 2017-06-17. The poor signal-to-noise and 'flaring' towards the end of the night are clear. Note that this 'flaring' is not physical, but is instead a result of highly variable seeing.

PSR J1628–3205

Identified as a radio pulsar from an unassociated Fermi source by Ray et al. (2012), this source is a redback with an orbital period of 5.0 hours. Optical observations were first reported by Li et al. (2014a), where a strongly ellipsoidal modulation was found. We have obtained 90% of an orbit of ULTRACAM/NTT g'- and i'-band photometry spread over 3 observations, on the nights of 2016-07-06, 2016-08-23, and 2017-06-17. The phase and flux calibrated light curve is shown in figure 2.8. The longest observation, 2016-08-23, covers the whole 90% phase range, with the observations from the other two nights being much shorter. Despite careful phase

and flux calibration, the small, long-term modulation reported by Li et al. (2014a) cannot be seen. However, the observing conditions during these observation were poor and this effect may be hidden by large uncertainties. The signal-to-noise of the g'-band is very poor, to the extent that the variability is only barely distinguishable from the noise. As well as this, a band of thin cloud appeared towards the end of the longest observation which caused large variability in the seeing. This manifests as a 'flaring' in the light curve, rendering significant portions of the data unusable, though it is possible that further reduction with variable aperture sizes may address this. Note that this is not flaring in the astrophysical sense, such as the bright flares associated with accretion discs, merely a label due to the similar appearance. The usefulness of this data can be improved by first binning the data in time (by a factor of 10) to increase signal-to-noise, and then performing the trimming process outlined in section 2.2.3 to produce the light curve shown in figure 2.9. The variability in both the i' and g' bands is now significantly clearer, but the small, long-term variability is still not visible. The variability of the g' band is significantly smaller than that of the i'-band, and the maximum in g' at phase 0.4 appears to lead the maximum in i' at phase 0.5. This asymmetric temperature distribution agrees with the observations of Li et al. (2014a), who suggest the cause is off-centre heating by the pulsar wind. Properly quantifying the effect will require more g'-band observations under better conditions.

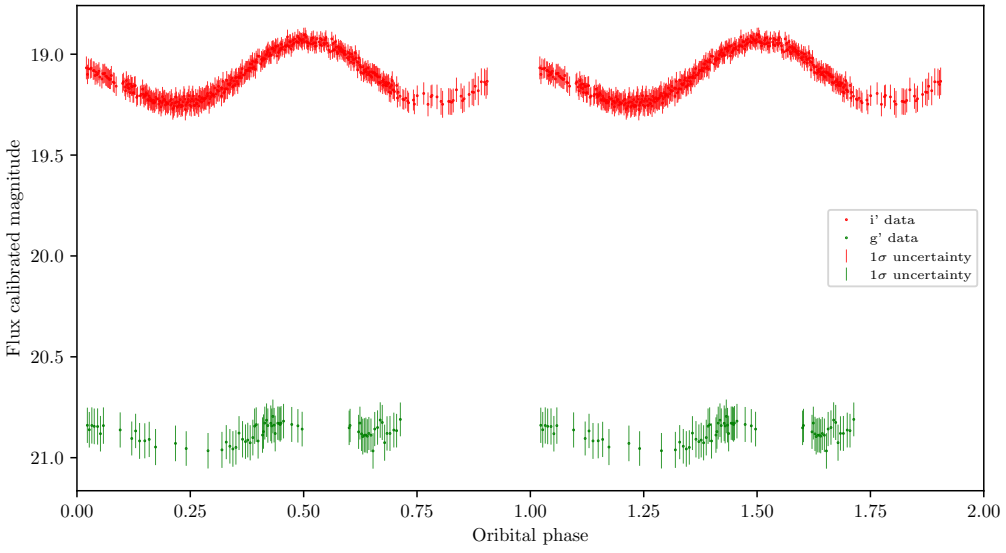

FIGURE 2.9: Flux and phase calibrated light curve of PSR J1628-3205, with 10-factor binning and trimming. The variability is now very clear in the i'-band, and visible in the g'-band. However, considering the few g' data points remaining, it is clear that better data is needed.

3FGL J0427.9–6704

J0427, introduced in section 2.3.2, is the subject of a paper recently released (Kennedy et al., 2020). Identified as an unassociated Fermi γ-ray source, Strader et al. (2016a) present NuSTAR X-ray light curves showing periodic eclipses, suggesting the presence of an accretion disk around a compact primary, with a low-mass near main sequence companion. A γ-ray light curve is also shown, with signs of an eclipse at the same phase as the X-ray eclipse, though this is not statistically significant. Optical photometry using 6 years of OGLE I-band data show a deep primary eclipse, and an uncharacteristically large scatter out of eclipse. Optical spectra show a dramatic switching of emission and absorption lines on approximately 20-minute timescales. The cause of this switching is unclear. Strader et al. (2016a) conclude that the source is a LMXB with a neutron star primary, however the lack of radio pulsations and no γ-ray eclipse throw some doubt on this. Following the investigation presented in Kennedy et al. (2020), the nature of the source remains unclear. It displays properties consistent with tMSPs, indicating this is an eclipsing binary with a neutron star primary, but since radio pulsations have not been detected a white dwarf primary cannot be ruled out.

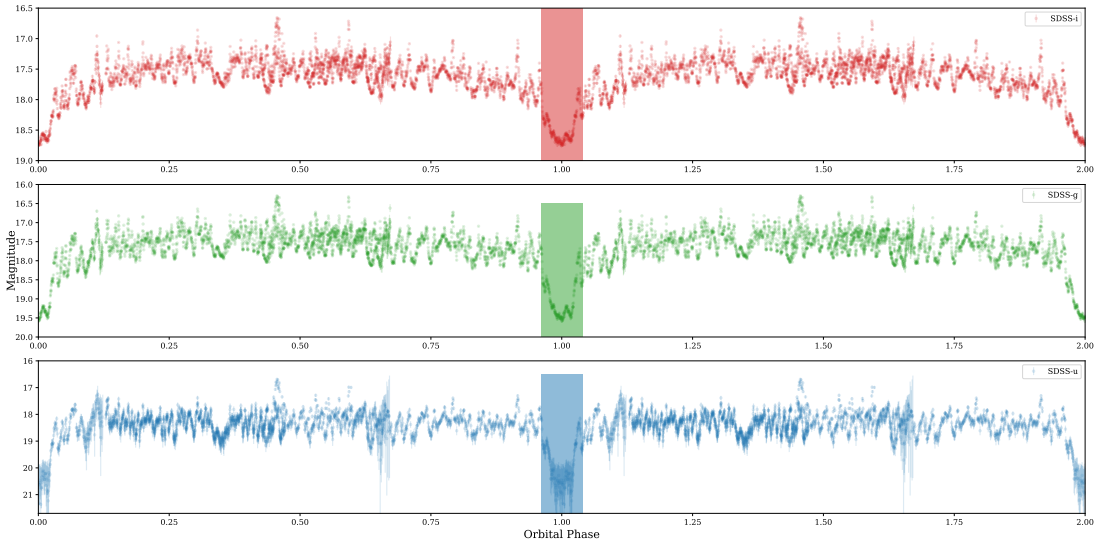

FIGURE 2.10: Phased and flux-calibrated SDSS u'-, g'-, and i'-band light curves of J0427. At phase 0.0/1.0 is the primary eclipse, highlighted in this figure. The small dip around phase 0.3 is not a secondary eclipse. The constant flickering and flaring is clear, as well as the underlying shape of the i'- and g'-band light curves in comparison to the relatively flat u'-band light curve. Image credit: Kennedy et al. (2020)

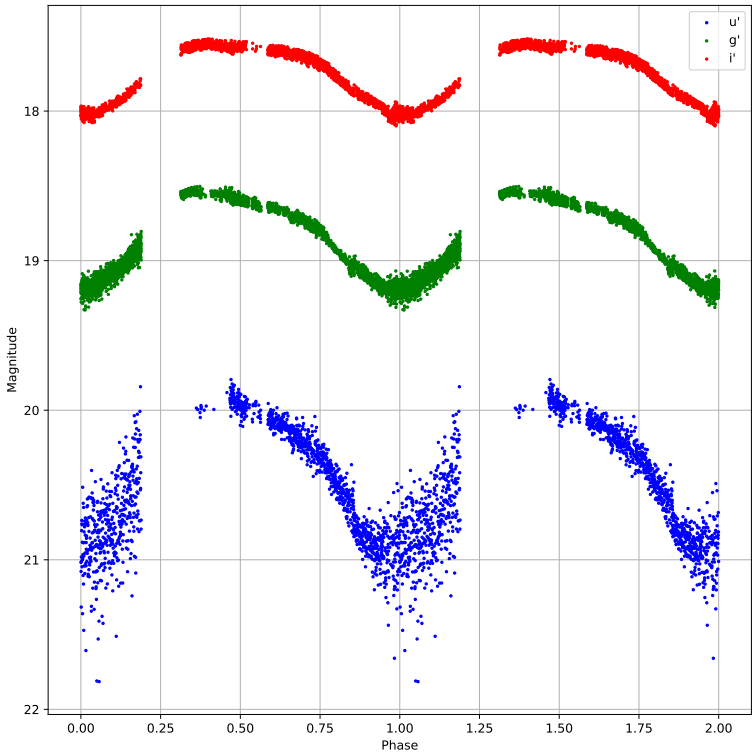

FIGURE 2.11: Phased light curve of J1306, repeated for clarity. The ~ 0.1 gap around phase 0.2 was not covered by observations. The observations of phases $0.9 - 1.0$ and $0.0 - 0.2$ show a comparatively lower signal-to-noise ratio, most noticeably in the u' band, due to poorer conditions.

PSR J1306–40

PSR J1306–40 (hereafter J1306), is a MSP in a redback system with an orbital period of $P_{orb} \approx 26.3$ hr, a spin period of 2.2 ms and a DM of 35 cm^{-3}pc, discovered in the radio timing survey, SUPERB (Keane et al., 2018). Keane et al. (2018) suggest this is a redback system with a significant radio eclipse, as no orbital solution was found. Subsequently, optical and X-ray followup presented in Linares (2018) corroborated the redback nature of the system. Linares (2018) used archival V-band photometry from the Catalina Sky Survey (CSS) data to determine that the orbital period of J1306 is 1.09716(6) days, and XMM-Newton X-ray data to confirm this period.

We acquired new, multi-band optical photometry of J1306, obtained in June 2018 with the ULTRACAM photometer on the NTT, using Super-SDSS u', g', and i' filters. The observations were performed over 8 consecutive nights from 01/06/2018 to 09/06/2018, totalling approximately 90% phase coverage. However, the weather conditions were often poor, with bands of thin clouds and poor seeing being the most common problems. As such, observations at some phases have significantly larger uncertainties and poorer signal-to-noise ratios than those observed with good conditions. We used exposures of approximately 10 s, chosen as J1306 is a relatively bright source with a mean r-band magnitude of 17.8.

The ULTRACAM images were reduced using the ULTRACAM pipeline, performing ensemble comparative photometry with 5 comparison stars. These comparison stars have known i' and g' magnitudes from the The AAVSO Photometric All-Sky Survey (APASS). However, no other sources in the field had known u' magnitudes. Instead, a zero-point photometric calibration was performed instead, as the u' zero-point and extinction are known for ULTRACAM on the NTT. This method is less accurate, as the extinction can vary by more than 0.1 mag per unit airmass from night to night, and the zero point can change by up to 0.5 mag depending on the condition of the telescope mirror. After initial review of the light curves, shown in figure 2.11, the effect of the poor weather was clear. Large drops in transmission percentage caused by passing clouds caused drastic artefacts in the reduced light curves. These phases were carefully masked and omitted from the data to be modelled. This masking was performed only on the u' channel, as in the i' and g' channels the calibrated magnitudes do not seem to significantly depart from the trend as the use of comparison stars with known magnitudes mitigates these fluctuations. Instead we only observe a decrease in signal-to-noise, compared to the erratic behaviour seen at the same phases in the u' channel.

We phased the data using the ephemeris from Linares (2018), folding at the orbital period. The resulting light curve shows a single-peaked structure, indicating that the heating of the day side of the companion is dominant, compared to a

relatively small ellipsoidal variation. The light curve is considerably asymmetric, with the optical maxima occurring at approximately phase $\phi = 0.3$, as opposed to the $\phi = 0.5$ as expected from the pulsar superior conjunction. This is reminiscent of the light curve of PSR J1023–0038 (hereafter J1023), the canonical transitional millisecond pulsar (tMSP) in its non-accreting radio powered state, as well as PSR J1227–4853.

2.4 Simulations

2.4.1 Signal-to-noise ratio

When simulating optical light curves, it is essential to first be able to estimate the signal-to-noise ratio (SNR) of your observations. There are a number of external factors that influence this, such as the sky brightness, seeing, and airmass, but also several variables intrinsic to the instrument in question, including the specifics of the CCD and telescope dimensions. As such, calculations are usually performed on a per-telescope and per-instrument basis, since it is not possible to generalise. The SNR calculators currently used in this project focus on the ULTRACAM and HiPERCAM instruments, detailed in section 2.2.1, on the WHT (William Herschel Telescope), GTC (Gran Telescope Canarias), and NTT (New Technology Telescope), each detailed in section 2.2.1. In addition to this a similar SNR calculator has been developed for GOTO, though the accuracy is not optimal as some key characteristics of the telescope have not yet been determined from commissioning. This is because the work in this thesis was carried out during the prototype phase of the telescope.

In all cases, the SNR calculations produce the number of object counts, number of sky (background) counts, and the resulting SNR for a source of a given magnitude, for given observing conditions, with a certain telescope and instrument pairing. The observing conditions include the sky brightness (the magnitude of the sky background in a given band, for a given moon phase), the filter, the exposure time, the airmass, and the seeing. The telescope and instrument used determine the overall sensitivity, defined by the so-called zero point (the source magnitude at which the CCD records one count per second per pixel) and the pixel scale (the resolution of the CCD, in arcsec pixel^{-1}).

2.4.2 Light curve simulations

The simulation

Here we introduce the simulation of optical light curves using the example of simulating an orbit of ULTRACAM/NTT data of PSR J1227–4853 (hereafter J1227), a tMSP now in a rotation-powered state having transitioned from an accretion powered state when observed in 2013 (de Martino et al., 2014a; Roy et al., 2015a). The aim of this simulation was to estimate the uncertainties on derived parameters, such as the inclination, i, and the mass of the neutron star primary, M_{NS}, given an optical light curve and some prior parameters. In practice, a full treatment would be performed using `Icarus`, though this method is suitable

Estimates of the mean and minimal magnitude of the source in i' band were made from the results in de Martino et al. (2014a), as well as the amplitude of the fundamental, irradiation component, A_{irr}. Values of the filling factor, f, the mass ratio, q, and the inclination, i, were also noted in order to calculate the magnitude of the ellipsoidal variation, A_{ell}, via the relationship

$$A_{\rm ell} \sim q f^3 \sin^2 i, \tag{2.11}$$

where A_{ell} is the fractional variation of flux with respect to the mean flux. These parameters are as described in table 1.1 and section 1.6.2. Note that this relationship requires a constant of proportionality, which is of order unity (Breton et al., 2012). It must be noted that these parameters were not obtained to high precision in de Martino et al. (2014a), and as such the magnitude of the ellipsoidal component is very much an estimate. Since A_{ell} is in terms of flux it was necessary to calculate the expected flux, using the SNR calculations introduced in section 2.4.1. We used an exposure time of 10 seconds (as J1227 is not a particularly faint source), typical values of airmass and seeing, 1.5 and 1.2 arcsec respectively, and assumed a dark moon for optimal sky brightness. By assuming the observations as sky-dominated, as they are for ULTRACAM on the NTT, it is possible to scale the signal-to-noise, S, at some flux, F, from the minimum SNR, S_{min} at minimum flux, F_{min}, via the equation

$$S(F) = \frac{F}{F_{\rm min}} S_{\rm min}. \tag{2.12}$$

As such, the SNR can be calculated for the minimum brightness of the source and simply scaled in order to estimate the uncertainty for all other simulated data points.

To simulate the light curve, we assumed a two term cosine series (to incorporate the irradiation and ellipsoidal variation), such that the flux is described by the

function

$$f(\phi, A_{\rm irr}, A_{\rm ell}) = -A_{\rm irr}\cos(2\pi\phi) + A_{\rm ell}\cos(4\pi\phi) + C, \tag{2.13}$$

where ϕ is the phase of the orbit and C is some offset to scale the function to the mean magnitude of the source. This function assumes the orbital modulation components are completely in phase. In practice, there is likely to be some offset between them, however. Then, for each exposure covering the orbital phase, a point was drawn from a Gaussian distribution with mean $\mu = f(\phi, A_{irr}, A_{ell})$ and standard deviation $\sigma = \mu/S(\mu)$. Then, the fluxes were converted to magnitudes to plot, and errors scaled using the same procedure outlined in section 2.2.3. The simulated light curve is is shown in figure 2.12

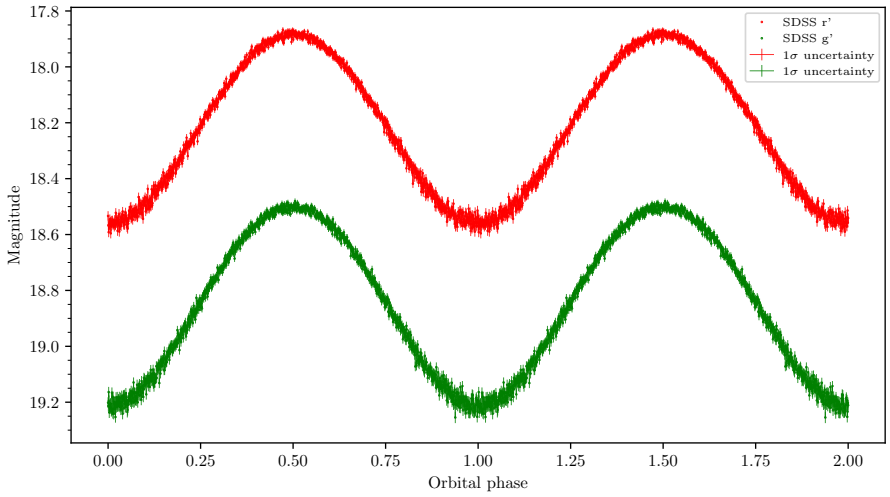

FIGURE 2.12: Two channels (Super SDDS r' and g') of simulated data of J1227–4853. At the magnitudes, ULTRACAM on the NTT allows for excellent sensitivity, even at a high cadence of 10 seconds.

Following this, Monte Carlo methods were used to estimate a value and error of the inclination, using a rearranged equation 2.11,

$$i = \sqrt{\arcsin\left(\frac{qf^3}{A_{\rm ell}}\right)}, \tag{2.14}$$

by using a large number of normally distributed values of f, q, and A_{ell} drawn using prior knowledge of their mean values and uncertainties. Following this, the estimated value of the inclination was used to determine the pulsar mass, $M_{\rm NS}$ by the same method, using the mass function,

$$\frac{M_{\rm NS}\sin^3 i}{(1+q)^2} = \frac{P_{\rm orb}K_2^3}{2\pi G}, \tag{2.15}$$

where P_{orb} is the orbital period, K_2 is the radial velocity of the companion star, and G is the gravitational constant. This equation can be rearranged in terms of M_{NS},

$$M_{\rm NS} = \frac{P_{\rm orb} K_2^3 (1+q)^2}{\sin^3(i) \cdot 2\pi G}. \tag{2.16}$$

The orbital period, P_{orb} is known to very high precision from pulsar timing (Roy et al., 2015a), and the companion radial velocity, K_2 is known from spectroscopy (de Martino et al., 2014a). The $V \sin i$ and K_2 estimates obtained from modelled spectra and the recovered amplitude of the ellipsoidal term were used to obtain estimates of the errors of $\sin i$ and $M_{\rm NS}$ using standard error propagation and Monte Carlo sampling. Here, $V \sin i$ is the projected companion spin velocity Given ULTRACAM data with the same quality as the simulation, we would obtain fractional uncertainties of 3.5% on the inclination, and 11% on the pulsar mass, compared to 18% and 36% (that is, $M_{\rm NS} = 1.4 - 3.0$ $\rm M_\odot$) respectively from de Martino et al. (2014b). As can be seen in chapter 3, this estimate does not take into account that highly asymmetric nature of the true optical light curves and the difficulties in modelling that this results in.

A similar simulation was performed for PSR J0952–0607 (hereafter, J0952) for a similar proposal. This source is far fainter, with minima reaching 25th magnitude. However, the larger diameter of the GTC compared to the NTT, and the increased sensitivity of HiPERCAM over ULTRACAM help to counteract this. As such, we may use the same exposure time to maximise our time resolution. The simulated light curve shown in figure 2.13 was produced.

Recovery of parameters

Until now, we have discussed simulations of light curves observed by highly sensitive telescopes and instruments. We also investigated the performance of smaller telescopes, such as the Thai Robotic Telescope (TRT) or GOTO. As a precursor to the work covered in chapter 4, we investigated how effectively the period could be recovered from light curves obtained using GOTO, and what factors affected this recovery ability. The signal-to-noise calculations outlined in section 2.4.1 were adapted to be used to estimate the SNR for GOTO observations. The instrumental magnitude (zero point), the sky brightness (at different lunar phases), and the extinction were estimated for the R, G, B, and extrapolated to the L filters. We note that the zero points were estimated from the limiting magnitude of the telescope, as opposed to using values from the instrument specification; the GOTO instrument is at this point still in the commissioning stage so these values are not exactly known.

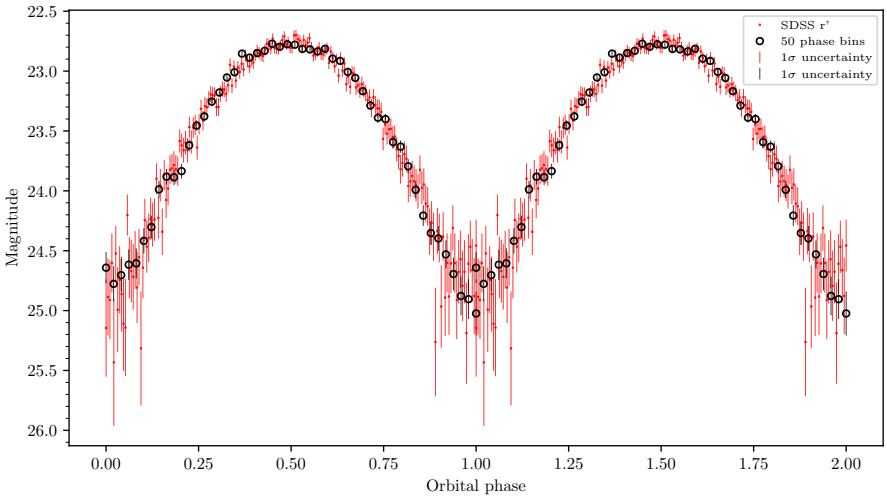

FIGURE 2.13: Simulated light curve of Super SDSS r' data of PSR J0952–0607, with HiPERCAM on the GTC. The noisy data around 25th magnitude is where the signal-to-noise ratio approaches 1. The black points are 50 phase bins.

Similarly the extinction and sky brightness were also estimated; data from the same observatory (ING - Isaac Newton Group of telescopes - at La Palma) were used. As opposed to the regularly spaced, high cadence observations capable with ULTRACAM, the nature of the GOTO observing schedule is such that we only obtain up to a few data points during a given night. For this preliminary investigation we assumed that a light curve was composed of N observations spread randomly across a year. We used an exposure time of 120 seconds with an R-band filter, and 'ideal' observing conditions of dark sky (no moon), an airmass of 1.5, and seeing of 1.0. In practice, some or all of the observations in a light curve will be in non-ideal conditions, such as with poor seeing or a full moon, so the results of this section should be taken as the best case outcome.

We aimed to determine to what extent certain factors influenced how well the period of a light curve could be recovered using a standard Lomb-Scargle periodogram. We considered the effect of the source brightness, period, the number of samples N, and the amplitude of modulation. We found that there was little to no variation in recovery ability across the range of orbital periods found in spider pulsars, around $1 - 30$ hours, as this is shorter than the total timespan of the observations.

We approached this problem by calculating the fraction of periods correctly recovered out of k light curves at each of a set of 5 SNRs over a range of modulation amplitudes. In this case, we vary the irradiation amplitude (parametrised as fractions of the mean flux) while keeping the ellipsoidal modulation fixed at 0.0 such that the test light curves are sinusoidal. We consider two cases of $N = 10$ and $N = 40$ separately. Each one of the k light curves was created with a period drawn

from a uniform distribution over the range $[1.5, 27]$ hours, with a phase offset drawn from a uniform distribution over the range $[0.0, 1.0]$. An observing schedule was created for each of $N = 10$ and $N = 40$, with N samples drawn randomly over two stages; N nights were chosen, and to each a random scatter was applied from a Gaussian distribution with mean and standard deviation $(0, 2)$ hours. This was done to minimise the effect of aliasing due to periodic observations. Indeed, aliasing is a significant problem with light curves of real GOTO observations; this is discussed in depth in chapter 4. We consider the following SNRs: $[3, 5, 10, 20, 80]$, which encompass worst-case through to best-case observations. An observation with an SNR of 3 is only just a significant detection. Since the telescope configuration is fixed, these correspond to observations of a source with mean R-band magnitudes R= $[20.9, 20.3, 19.5, 18.7, 16.8]$. The (fractional) modulation amplitudes considered are 50 logarithmically spaced fractions between 0.01 and 1.0. We used $k = 100$ light curves per set of parameters.

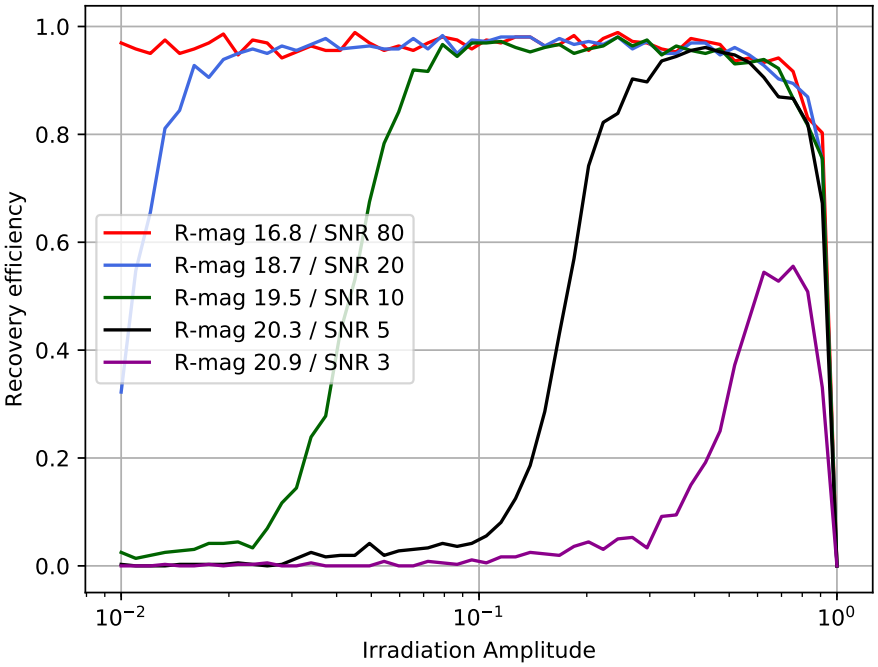

FIGURE 2.14: Effect of light curve modulation amplitude on period recovery efficiency for light curves with 10 samples. Each colour represents a different mean magnitude, showing how brighter sources provide significantly better results than those near the detection limit of the instrument. The irradiation amplitude is defined as a fraction of the mean flux.

The dependence of the fraction of periods recovered on each of these parameters can be seen in figure 2.14 for $N = 10$ samples and in figure 2.15 for $N = 40$ samples. In both cases, the period recovery efficiency (i.e. the fraction of periods correctly

recovered) is positively correlated with the SNR and the amplitude of modulation. Rather than a gradual change, the recovery efficiency drops off sharply at a threshold where the modulation amplitude becomes comparable to the random scatter due to noise. The ability to recover periods is significantly better with $N = 40$ samples compared to $N = 10$; this is to be expected as with Lomb-Scargle periodograms, better sampling generally results in more accurate period estimation and better robustness to noise. Note that the thresholds at which the recovery efficiency sharply decreases is not consistent between each figure; for example for an SNR of 5 this occurs at an irradiation amplitude of approximately 0.2 for $N = 10$ and 0.05 for $N = 40$. This suggests that sources with relatively small modulation amplitudes, such as redbacks, will require a large number of samples to accurately recover their orbital periods.

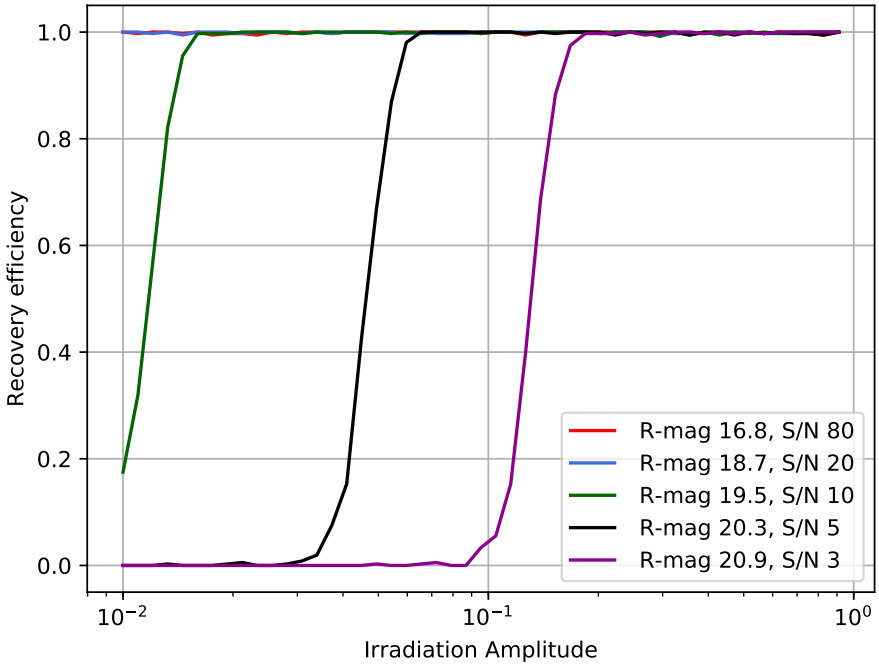

FIGURE 2.15: Effect of light curve modulation amplitude on period recovery efficiency for light curves with 40 samples. We see a clear improvement with an increased number of samples, including substantially improved performance as the modulation amplitude approaches unity. Note that for SNR 20 and SNR 80, the period recovery efficiency remains near 100% at all times, suggesting that light curves of even smaller modulation amplitude may be recovered.

For $N = 10$, it can be seen that the recovery efficiency quickly drops to zero as the irradiation amplitude approaches unity. This represents a light curve where the optical minima is below the detection limit of the telescope, such as with many black widows, despite a large amplitude of modulation. In this case a substantial fraction of the observations are of SNR ~ 0, that is to say they do not contribute

to the periodogram. For $N = 10$ this can leave too few observations to determine a period, while for $N = 40$, figure 2.15 illustrates that the period can still be accurately recovered.

2.5 Conclusion

While there are few key results to discuss, the work in this chapter lays important groundwork for the research in chapters 3 and 4. However, there are several notable outcomes in addition to this. The optical observations using ULTRACAM have formed the basis for published papers in the case of J0427 (Kennedy et al., 2020), as well as on-going work such as with J1306 and J1628. As well as this, the data reduction pipelines introduced have been used to produce the light curves used in chapter 3. The scheduler has been used frequently for planning telescope time proposals and observations throughout the course of the research, and has also formed the basis for an updated version of the code which continues to be used for the same purposes. Likewise, the ULTRACAM and HiPERCAM simulations have been used extensively to provide technical justification for telescope time proposals, while the GOTO simulations have provided the basis for the synthetic light curve code used in chapter 4.

Chapter 3

Optical photometry of two millisecond pulsars

3.1 Introduction

Transitional millisecond pulsars (tMSPs) are a class of neutron star binary containing a recycled millisecond pulsar (MSP), spun up by accretion from a low-mass, semi-degenerate companion to spin periods of the order of milliseconds (Alpar et al., 1982). tMSPs are unique in that they are observed to transition between an accretion-powered (AP) Low-Mass X-ray Binary (LMXB) state and a rotation-powered (RP) radio pulsar state, the latter so far associated with the 'redback' class of pulsars (Archibald et al., 2009). Redbacks are a sub-class of the eclipsing 'spider' binaries, in which a low-mass ($0.2M_{\odot} \lesssim M_c \lesssim 0.4M_{\odot}$) quasi-main sequence companion star in a tight (∼few hour), tidally locked orbit is irradiated by the wind of a MSP. This results in the ablation of the companion's surface into a tail of ionised matter, causing long eclipses at radio frequencies, and distinctive quasi-sinusoidal optical modulation caused by heating of the inner face of the companion, e.g. Breton et al. (2013); Roberts (2011). Spider binaries host some of the most massive and fastest spinning neutron stars (Linares, 2020).

As summarised in Britt et al. (2017), observations of the three confirmed tMSP systems have revealed several shared characteristics, though it is important to note that due to the small sample size these could be coincidental. In the AP state, pulsed X-ray, UV, and optical emission has been detected (Jaodand et al., 2016; Papitto et al., 2019; Jaodand et al., 2021), as well as bimodal flickering and the flaring in X-ray emission typical of an accretion disc (Patruno et al., 2014; Linares, 2014). The optical emission in the RP state is indistinguishable from that of non-tMSP redback systems as described in the previous paragraph, while the AP state emission also exhibits bimodal flickering and flaring (Kennedy et al., 2018; Shahbaz et al., 2018).

Strong emission lines are seen in the AP state (Archibald et al., 2009; Bassa et al., 2014b) which fully disappear in the RP state, suggesting they originate from the accretion disc. In the AP state, tMSPs exhibit a flat radio spectrum suggesting self-absorbed synchrotron emission, while in the RP state the radio emission is pulsed with a spectrum characteristic of synchrotron emission with a steep power law (Archibald et al., 2009; Patruno et al., 2014), typical of rotation-powered MSPs.

tMSPs present a unique opportunity to not only study the accretion mechanism of LMXBs, but also gain insight into the evolution of pulsar binary systems. It is generally agreed that LMXBs are the predecessor to spiders and several other types of MSP binary, but the mechanism by which the accretion is 'switched off' is not known (Chen et al., 2013), nor is the mechanism by which the MSP magnetic field decays as it gets recycled (Konar & Bhattacharya, 1997; Cumming et al., 2001). As such, the study of tMSPs is important in uncovering the evolutionary history of spiders and LMXBs: they may be a missing link between these two populations (Archibald et al., 2009). However, it is possible that they are themselves a distinct population; in this case they remain important astrophysical laboratories to study the accretion process. Since the timescale of their transitions is on the order of weeks or months, with transitions occurring every few years, we can study the entire accretion process on human timescales.

We present new optical light curves of two tMSPs, both in the radio pulsar states; PSRs J1023+0038 and J1227–4853. These are two of the three confirmed tMSPs; the third is PSR J1824 − 2452I, although its location in a globular cluster prevents a detailed study in optical wavelengths (De Falco et al., 2017; Coti Zelati et al., 2019). We note that there are a few 'candidate' tMSPs, such as 3FGL J0427.9 − 6704 (Strader et al., 2016b; Kennedy et al., 2020), which show similar AP state properties to confirmed tMSPs but lack a radio MSP association and have not yet been seen to transition.

Often referred to as the canonical tMSP, PSR J1023+0038 (herafter J1023) was initially classified in 2001 as a cataclysmic variable system with a binary period of 0.198 days (4.75 hours) (Bond et al., 2002). The double-peaked emission lines and blue optical spectrum indicated an accreting binary with a white dwarf primary, with optical photometry showing the flickering and flaring typical of an accretion disc. Woudt et al. (2004) and Thorstensen & Armstrong (2005a) presented the first evidence for a state change, respectively showing optical photometry and spectroscopy which lacked the usual signatures of an accretion disc. The strong emission lines in the optical spectra were replaced by absorption features, while the flickering and flaring in the light curve were no longer present. The state change was confirmed in

2007 with the detection of a radio pulsar with a spin period of 1.69 ms (Archibald et al., 2009; Wang et al., 2009). In June 2013, the radio pulsations from the MSP could no longer be observed (Stappers et al., 2014b), and were replaced by rapidly varying X-ray flux indicative of an accretion disc (Patruno et al., 2014). Kepler-K2 optical observations in 2017 also show clear evidence of an accretion disc (Kennedy et al., 2018; Papitto et al., 2018).

Identified as a variable X-ray source with XMM-Newton (Bonnet-Bidaud et al., 2012; de Martino et al., 2013), XSS J12270 − 4859 (now PSR J1227–4853, hereafter J1227) was initially classified as an LMXB due to the presence of flares and 'dips' in the X-ray light curve. Between 2012 and 2013, the X-ray and optical fluxes of J1227 were observed to decrease to new minima (Bogdanov et al., 2014; Bassa et al., 2014b), and the spectral emission features of an accretion disk disappeared. Radio observations revealed a MSP with a period of 1.69 ms at the source coordinates (Roy et al., 2015b), showing that J1227 had transitioned from an LMXB state to radio pulsar state displaying a redback-like optical modulation with an orbital period of 0.288 days (6.91 hours). Gamma ray pulsations at the radio MSP period were discovered using data from the Large Area Telescope (LAT) on the *Fermi Gamma-ray Space Telescope* (Fermi-LAT), which indicated an LMXB to a tMSP transition epoch of 2012-11-03 (Johnson et al., 2015).

In this study we first present our new photometry, outlining the reduction and calibration procedure, in section 3.2. In section 3.3 we discuss the nature of and the potential mechanisms behind the asymmetry of the light curves. We discuss our modelling of these light curves using the Random Forest binary light curve synthesis code in section 3.4, in particular our constraints on the orbital parameters of the systems, and implement two extensions to the Random Forest model. The first extension accounts for an additional hot spot on the surface of the companion, and the second is a new description of the temperature distribution of the companion which takes into account diffusion and convection in the outer shell. Then, we outline the results of the modelling and discuss their validity and the implications on the tMSP transition mechanism in sections 3.5 and 3.6.

3.2 Optical observations

3.2.1 ULTRACAM on the NTT

Our observations were performed using the ULTRACAM instrument mounted on the 3.5 m New Technology Telescope (NTT) at the La Silla observatory, Chile.

ULTRACAM (Dhillon et al., 2007) is an optical imaging photometer capable of simultaneous 3-band observation. We used filters from the ULTRACAM Super Sloan Digital Sky Survey (Super-SDSS) $u_s g_s r_s i_s z_s$ photometric system (Dhillon et al., 2018), with u_s and g_s filters on the first two CCDs, and either of i_s or z_s for the third. Our typical integration time was 10 seconds with 25 ms dead time between each frame. Our observations are summarised in table 3.1.

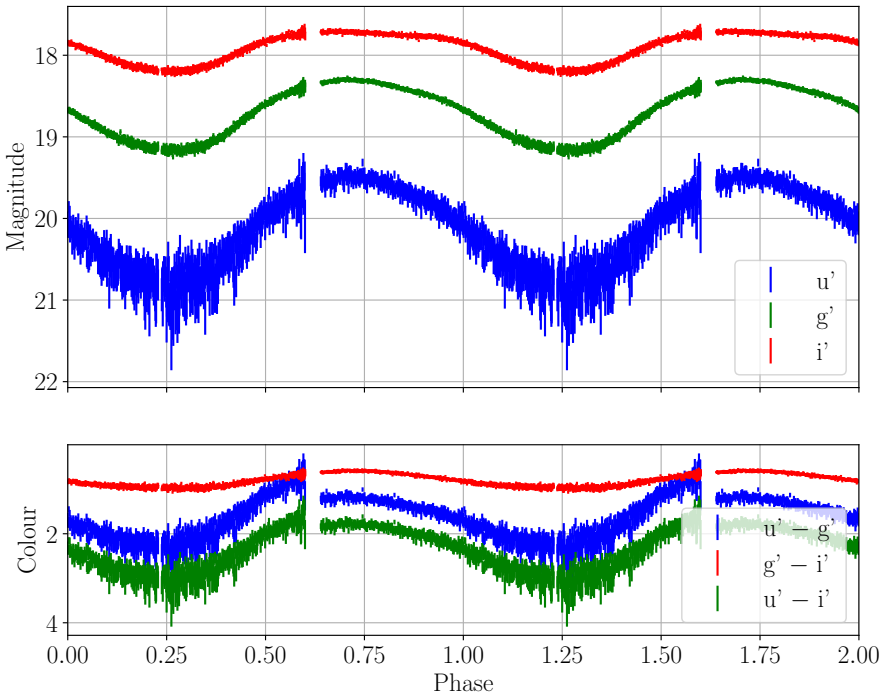

FIGURE 3.1: Phased light curve of J1227, repeated over two cycles for clarity, with each colour corresponding to a different filter as in the legend. The asymmetry can clearly be seen in the i_s and g_s bands around phase 0.75. At phase 0.6, the Sun rising is responsible for the large scatter, particularly so in the u_s band. The bottom panel shows the colour information. We adopt the phase convention where the pulsar is at superior conjunction at phase 0.25.

3.2.2 Observations of J1227

J1227 was observed on 2019-02-27 beginning at 03:09:53 UTC, during its radio pulsar state. The observations were completed in one night, providing more than 90% orbital phase coverage in mostly photometric conditions, although some clouds were present near the end of the observation, decreasing the SNR of these images. We

TABLE 3.1: Table of observations of J1227 and J1023, with information gathered from the ULTRACAM online logs and reduction files. The seeing, displayed as the median, was calculated from the FWHM of the observations. From the observer comments, we note cloudy conditions on 2010-05-05, 2020-01-27, and 2020-01-28. Further to the co-adds shown here, no further binning in time was performed during reduction.

Date	Start time (UTC)	Source	Filters	Phase coverage	Seeing (")	Airmass	u_s co-adds	Exposure time (s)
2010-05-04	23:03:48	J1023	u_s / g_s / i_s	$\sim 105\%$	1.23	1.15-2.08	18	9.56
2010-05-05	22:58:12	J1023	u_s / g_s / i_s	$\sim 90\%$	1.14	1.15-2.16	6	9.56
2010-05-06	23:14:26	J1023	u_s / g_s / z_s	$\sim 105\%$	1.17	1.15-2.26	6	5.76
2019-02-27	03:04:41	J1227	u_s / g_s / i_s	$\sim 90\%$	0.90	1.06-1.42	2	10.02
2020-01-27	04:59:54	J1227	u_s / g_s / i_s	$\sim 60\%$	1.14	1.06-1.45	3	10.01
2020-01-28	05:10:59	J1227	u_s / g_s / i_s	$\sim 55\%$	1.04	1.06-1.39	3	10.01

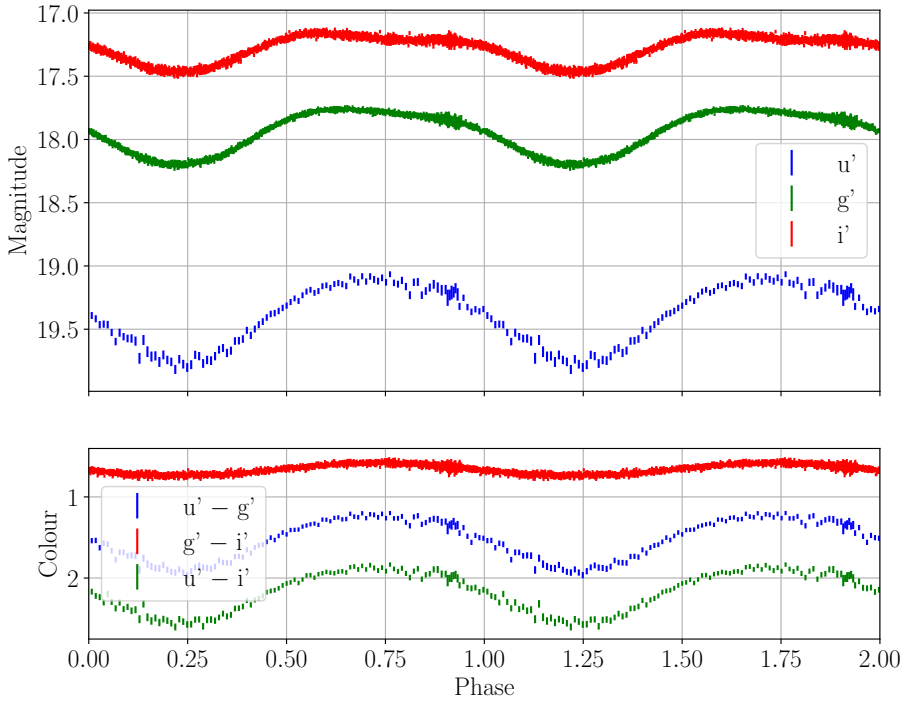

FIGURE 3.2: Phased light curve of J1023, clearly showing the asymmetry around phase 0.75. Similarly to J1227, this asymmetry is more pronounced in the i_s and g_s bands. Note that the artefacts around phase 0.8-0.9 are due to poor seeing conditions.

reduced the data with the ULTRACAM pipeline using an ensemble aperture photometry method (Honeycutt, 1992). We used 8 comparison stars common to u_s, g_s, and i_s to correct for atmospheric transmission variations. We employed the same 8 calibration stars of known i' and g' magnitudes from the AAVSO Photometric All-Sky Survey (APASS) to calibrate the i_s and g_s magnitudes to the absolute photometry system. A colour term is used during reduction to transform between the two systems. The same comparative photometry was also performed for the u_s band, but as there were no objects with known u_s magnitude in the field we used the zero point of this band, calculated from separate observations of SDSS standard stars, to calibrate the magnitudes instead. While the zero point and typical extinction coefficients are known for ULTRACAM in this configuration, this is a less accurate calibration method than comparative photometry, so we included a larger band calibration offset for the u' band in the modelling.

During observations a temporal co-addition factor of 2, where the CCD is read out every other exposure, was used for the u_s band and the resulting SNR of the data was sufficient that we did not need to perform any further temporal averaging of any

of the bands. The final step of our reduction was to discard any observations with error flags from the pipeline or SNR below a threshold of 3. This second condition was used as several observations near the optical minimum were impacted by cloud cover. The resulting dataset contains a total of 5899 good data points: 2360 in i_s, 2359 in g_s, and 1180 in u_s. These data were folded at the orbital period using the ephemerides from radio timing (Roy et al., 2015b). We will apply the following convention thorough the paper to define the orbital phase: zero phase corresponds to the pulsar ascending node, and therefore, the companion inferior conjunction (optical minimum) occurs at phase 0.25.

The phased light curve of J1227 (see fig. 3.1) displays single-peaked sinusoidal modulation, due to the irradiation of the companion. The peak-to-peak amplitude of modulation is approximately 0.6 mag in i_s, 0.8 mag in g_s, and 1.4 mag in u_s, with mean magnitudes of 18.0 mag, 18.7 mag, and 20.4 mag respectively. Considering the colour information, the companion star becomes redder during the pulsar superior conjunction (optical minimum), in line with the expectation that the night side of the star is cooler than the day side. The light curve shows the asymmetric nature of the modulation, and a 'flattening' of the optical maxima most noticeable in the i' band due to a significant ellipsoidal modulation contribution.

3.2.3 Observations of J1023

J1023 was observed over 3 consecutive nights starting on 2010-05-04, during the object's radio pulsar state. While the observations provide nearly 100% phase coverage in u_s, g_s, i_s, and z_s, the latter two nights suffered from cloudy skies and so the quality of the first night of observations far exceeds that of the second and third, both in terms of usable phase coverage and SNR. Additionally, due to the cloud cover, the magnitude calibration is not completely consistent between nights. As such, we performed the modelling using just the data from 2010-05-04 to ensure that these potential calibration issues did not affect the results.

The observations were reduced in the same way as with J1227, using ensemble aperture photometry with the ULTRACAM pipeline. However, as the position of J1023 has been covered by the Sloan Digital Sky Survey (SDSS), calibration stars were available for all four bands, including u'. In total, 12 comparison stars were used for i_s and z_s, 11 for g_s, and 6 for u_s, with the same number being used to calibrate the magnitudes. We obtained a total of 3819 good data points; 1859 in i_s, 1858 in g_s, and 102 in u_s, which were folded on the orbital period. The phased light curve, shown in figure 3.2 shows asymmetrical modulation in all 3 bands, with a single irradiation peak in u_s. The relative contribution of the ellipsoidal variation to the light curve shape, which produces the double-peaked modulation per orbit

(see e.g. Li et al. (2014b) for a clear example of this), is higher for redder bands (see e.g. i_s band compared to u_s band). We measure mean magnitudes of 17.3 in i_s, 17.9 in g_s, and 19.4 in u_s, with modulation amplitudes of 0.3 mag, 0.4 mag, and 0.7 mag respectively.

3.2.4 Radial Velocities

To help constrain the projected companion radial velocity, K_2, of J1023 we combined our photometric data with spectroscopic radial velocity measurements from Shahbaz et al. (2019) and McConnell et al. (2015). We used radial velocity curves obtained from metallic line spectra captured with the ISIS instrument on the 4.2 m William Herschel Telescope (WHT) in 2016 for Shahbaz et al. (2019) and 2009 for McConnell et al. (2015). Both radial velocity curves had been produced using broadly the same set of metallic lines, over the same range of wavelengths.

3.3 Asymmetries

While observed in both these tMSP systems, asymmetric light curves appear to be a feature of redback systems in general (e.g. PSR J2215+5135, Schroeder & Halpern (2014)) and are not specific to tMSPs. As such it is unlikely that they arise from, for example, reprocessing or obscuring of the pulsar wind by some disc remnant. Indeed there is no mechanism driving the asymmetry that is widely accepted and evidenced, though there are a number of possible theories.

Considering the work of Romani & Sanchez (2016), a swept-back intra-binary shock (IBS) between the pulsar and companion winds could be responsible for the asymmetry via heating of the companion by non-thermal X-ray emission produced in the wind shock. In that work the modelling includes the effect of the IBS heating on the companion and finds good agreement with data. More recently, Kandel et al. (2020) performed modelling of the asymmetric redback PSR J2339-0533 using a hot spot model which aims to describe the ducting of high-energy particles (such as those shed from the IBS) onto magnetic caps on the companion star's surface. However, the results of Zilles et al. (2020), which estimate the penetration depths of high energy photons in the companion photosphere, suggest that the X-rays reprocessed by the shock could not sufficiently heat the companion to the observed asymmetry temperatures.

Dynamics on the companion surface may instead produce an asymmetric temperature profile. As the day side of the companion is strongly heated, we may expect strong circulatory winds and thermal structures similar to those observed on

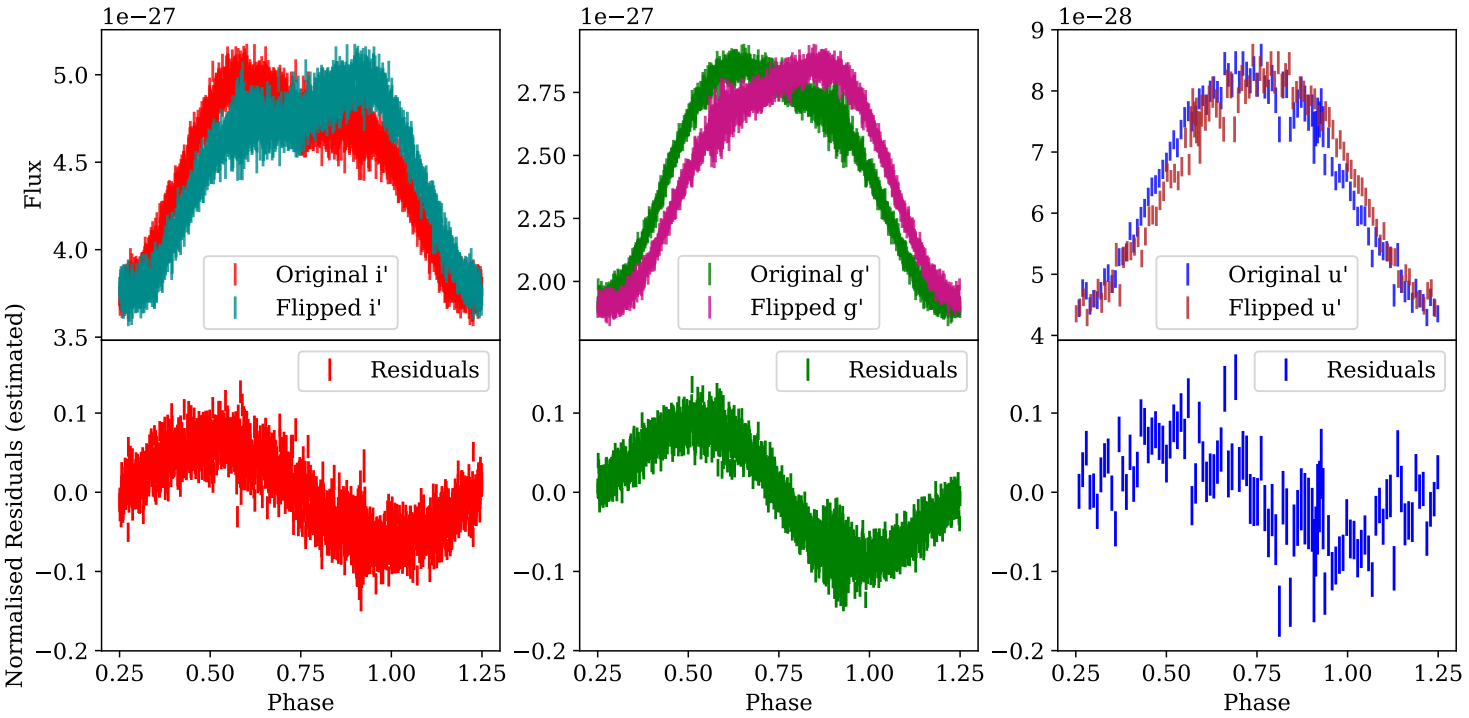

FIGURE 3.3: *Top:* i_s, g_s, and u_s light curves of J1023 overlaid with the same light curves mirrored about phase 0.5, illustrating the asymmetries. *Bottom:* Estimated residuals between original and mirrored light curves, normalised to the mean band flux. The mirrored light curve was interpolated over the original phases in order to calculate these residuals.

hot Jupiters (see, e.g., Lothringer & Casewell (2020); Komacek & Showman (2020); Jackson et al. (2019)) or cataclysmic variables (Martin & Davey, 1995). The lack of fusion in hot Jupiters allows these global winds to form complex meteorologies, which are unlikely to form on redback companions as the radial convection in the envelope would disrupt it. The large temperature gradient between the day and night side would be sufficient to fuel the winds and so allow heat to flow through the atmosphere, with circulation driven by the spin of the companion (Tan & Showman, 2020). We implement this as presented in Voisin et al. (2020), however we note there are possible alternatives. In de Wit et al. (2012), an alternative hot spot model is presented, with the temperature distribution motivated by these thermal flows. Or, in Demory et al. (2013), a longitudinal temperature map of the surface of the companion is used, using a number of fixed bands.

A recent, novel approach in Romani et al. (2021) models the asymmetry of the optical light curve of the black widow pulsar PSR J1810+1744 by acting on the gravity darkening parameter. However, this method was used in the case of a very highly irradiated companion, while these tMSPs display only modest irradiation compared to their internal energy output.

We note the similarity to the HW Vir class of compact binaries, consisting of a hot sub-dwarf primary and cool, close companion (typically a white dwarf) (Schaffenroth et al., 2019), which do not display asymmetric light curves despite the otherwise similar orbital parameters. In these systems the sub-stellar point on the companion is heated to temperatures of over $\sim 10^4$ K, significantly more than the tMSPs in this work. We suggest that the depth to which the irradiation in HW Vir binaries is significantly shallower than with tMSPs, due to the different source of heating, so the heat redistribution layer may not be sufficiently deep to produce asymmetries.

To illustrate the asymmetry in J1023 we overlay the light curves of each band, in fluxes, over the same light curve mirrored about phase 0.5. We are then able to analyse the asymmetry by calculating the residuals. We normalise these residuals to the mean flux of each band then interpolate the flipped light curve at the phases of the original in order to calculate residuals between each curve. Note that this means that these are not residuals in the traditional sense due the interpolation, however they clearly demonstrate the difference between the original and mirrored light curves. Seen in figure 3.3, the residuals follow a nearly sinusoidal shape across all three bands. The amplitude of these sinusoids is also comparable across the bands, largest in the g_s band, followed closely by i_s and then u_s at around 60% the amplitude of i_s. Note that there is significantly increased scatter in u_s compared to g_s and i_s, which causes the amplitude of the sinusoid to appear larger than it is.

Performing the same analysis on the light curves of J1227, we find that the shape

of the residuals is again consistent between each band. However, they more closely follow the shape of a sinusoid at the second harmonic; this is shown in figure 3.4 for the g_s band. Additionally, the amplitude of this modulation is much more varied between bands; strongest in u_s, decreasing to 60% in g_s and finally to roughly 25% in i_s.

3.4 Numerical modelling

3.4.1 Icarus

We modelled our optical light curves using the `Icarus` code (see Breton et al. (2012) for a thorough introduction) in order to constrain the orbital parameters of the system and the temperature profile of the irradiated companion. Briefly, the companion star surface is deconstructed into a tessellation of surface elements. These surface elements define a temperature distribution of the companion surface informed by a model atmosphere for each observational band. Heating from irradiation, plus any hot spots, is then applied, then the optical flux is calculated at each orbital phase. At each orbital phase, the model is evaluated using a given set of input parameters and the χ^2 residual between the model and the observation is calculated. We use atmosphere grids created using the ATLAS9 synthesis code (Castelli & Kurucz, 2003). In the standard heating model, the input parameters the model uses are as follows: the orbital inclination angle, $\cos(i)$, the Roche lobe (RL) filling factor, f, the base (night side) temperature of the companion, T_0, the irradiation temperature of the companion, $T_{\rm irr}$, the distance modulus, DM, the companion's projected radial velocity amplitude, K_2, the mass ratio (defined as the ratio of the pulsar mass to the companion mass), q, the co-rotation coefficient, Ω, the gravity darkening coefficient, $\beta_{\rm g}$, and the V-band extinction coefficient, A_v. Note that while the `Icarus` code uses the DM as a model parameter, we often discuss the distance derived from this; the distinction between the two will clearly be made.

The filling factor is defined as the ratio of the stellar surface radius in the direction of the pulsar to the distance to the L1 point. We also derive a volume-average filling factor which is a representation of the volume of the star to the volume of the Roche lobe, $f_{\rm VA} = \langle R \rangle / \langle R_{\rm RL} \rangle$. We fix the co-rotation coefficient to $\Omega = 1$, as we assume both systems are tidally locked, and for both sources we fix $\beta_{\rm g} = 0.08$ as we assume the late-type companion stars have large convective envelopes (Lucy, 1967). As well as the volume-averaged filling factor and distance, we also derive the pulsar mass, $M_{\rm psr}$ and the blackbody-equivalent temperatures of the day and night side of the companion.

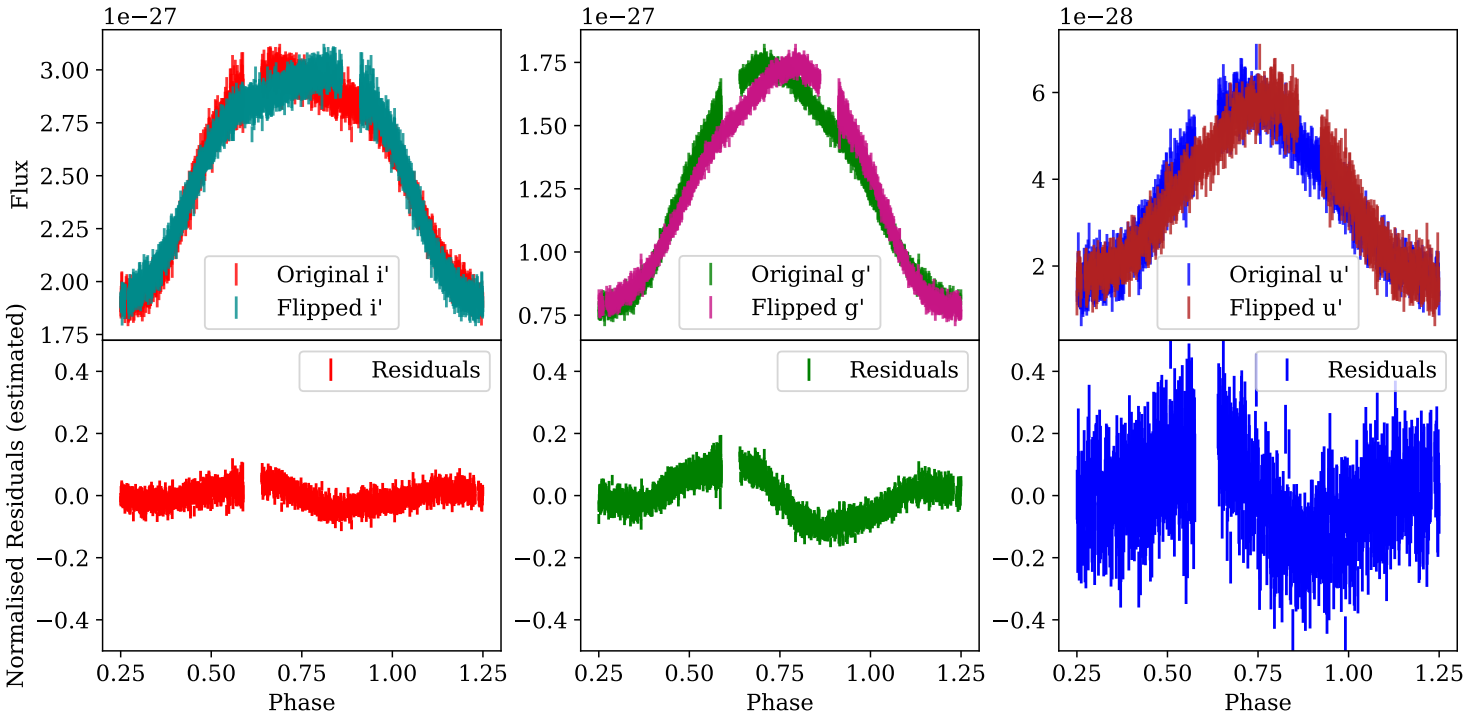

FIGURE 3.4: As figure 3.3, for J1227. The light curves were trimmed to remove the region of large scatter around phase 0.6 to 0.65. Note the shape now reflects that of a second-harmonic sinusoid, compared to the fundamental sinusoid seen in J1023.

Relations connecting some of the parameters previously introduced and the orbital ephemerides are also considered in our modelling. In particular, the mass ratio follows the relation

$$q = \frac{K_2 P_B}{2\pi a_1}, \tag{3.1}$$

where P_B is the orbital period and a_1 is the projected semi-major axis of the pulsar, derived from radio timing (Roy et al. (2015b) and Archibald et al. (2009) for J1227 and J1023 respectively). While we expect tMSPs to potentially host massive neutron stars (recycled MSPs) we do not put any tight constraints on the pulsar mass, as we use only the mass ratio of the system as a parameter. Finally, we allow the systemic velocity, Γ_v to be a free parameter as our radial velocity data are not mean-subtracted. At each step in the MCMC chain, we use the `Icarus` code to determine the effective centre-of-light radial velocity, $v_{\rm eff}$, of the companion, evaluated using a model atmosphere corresponding to the wavelength range of the spectroscopic radial velocity curve. In this case, this corresponds to the SDSS r' band. The amplitude of this effective radial velocity curve is defined as $K_{\rm eff}$. We then fit for and subtract any linear offset between the modelled and measured radial velocity curves, and calculate the residuals. These residuals are combined with the residuals from the modelled light curve to calculate the posterior probability at each step in the chain.

To determine the model parameters we used the multi-nested sampler, `MultiNest` (Feroz et al., 2009), implemented in Python as `pymultinest` (Buchner, J. et al., 2014). This method also allows us to directly compare the Bayesian evidence of each choice of model. In this paper this is quoted as the natural logarithm of the model evidence, $\log Z$, and the reduced χ^2 is determined from the best posterior solution (as opposed to the best likelihood solution). The selection of priors is very important as the model is extremely degenerate; where possible, we use priors on parameters informed by published measurements. For these we use Gaussian priors, centred on the literature value with standard deviation equal to the given uncertainties (i.e. the 68% significance). For parameters with no known constraints, we either use top-hat priors over a range of physically sensible values (such as constraining the temperature to the range of the atmosphere grids) or leave the parameter unconstrained.

3.4.2 Standard symmetrical direct heating model

Initially, we used a symmetrical direct heating model to act as a benchmark. This model assumes a constant base temperature, T_0, across the companion, then models the effect of heating by assuming that the additional flux is thermalised and locally

re-emitted such that we can express the day side temperature, $T_{\rm day}$ as

$$T_{\rm day}^4 = T_0^4 + T_{\rm irr}^4, \tag{3.2}$$

where $T_{\rm irr}$ is the so-called irradiation temperature. The effect of gravity darkening is applied prior to the irradiation to give the temperature distribution of the companion. Irradiation effects take into account the distance between the companion and the pulsar and the incidence angle of the irradiation.

For J1227 the free parameters in this initial model were the orbital inclination, $\cos i$, the base temperature, T_0, and irradiation temperature, $T_{\rm irr}$, of the companion, the distance modulus, DM, the companion velocity K_2, and the filling factor, f. While de Martino et al. (2014b) suggest an inclination of between $43^\circ < i < 73^\circ$ and the modelling of de Martino et al. (2015) constrains $46^\circ < i < 65^\circ$, we opted to leave the inclination unconstrained to perform independent modelling (i.e. not biased by previous studies). We used top hat priors on the temperatures, setting the limits in accordance with the range of temperatures covered by our model atmosphere grids; 1300 K to 10000 K. We left the filling factor mostly unconstrained, with limits $0.0 < f < 1.0$, as we expect the Roche Lobe to be mostly full, but have no physically imposed minimum filling factor. The priors on the distance modulus were calculated following the method described in Luri et al. (2018), using the joint probability distribution of distances derived from the GAIA parallax measurement combined with the model of galactic MSP densities and velocities from Levin et al. (2013). The GAIA parallax was significant, at 0.623 ± 0.168 mas, and so this dominated the distribution. The resulting prior distribution for J1227, showing the contribution of each component, is shown in figure 3.5.

We constrain the companion velocity, K_2, based on the radial velocity amplitude inferred from spectroscopy by de Martino et al. (2014b); $K_2 = 261 \pm 5$ km s^{-1}. Note that the radial velocity data are not publicly available at the time of writing. Rather than using a simple Gaussian prior on this value, we instead use the radial velocity curve method described in the previous section. Our radial velocity 'curve' consists of this single value at phase $\phi = 0.25$, the pulsar inferior conjunction, where the maximum projected velocity occurs. This method ensures the radial velocity amplitude from spectroscopy is corrected to the centre of mass of the companion. We also use Gaussian priors on the derived parameter $V \sin(i) = 86 \pm 20$ km s^{-1}, the companion rotational velocity, obtained from the same spectroscopy, using the relation derived in Wade & Horne (1988),

$$V \sin(i) = (K_1 + K_2) R_2(f), \tag{3.3}$$

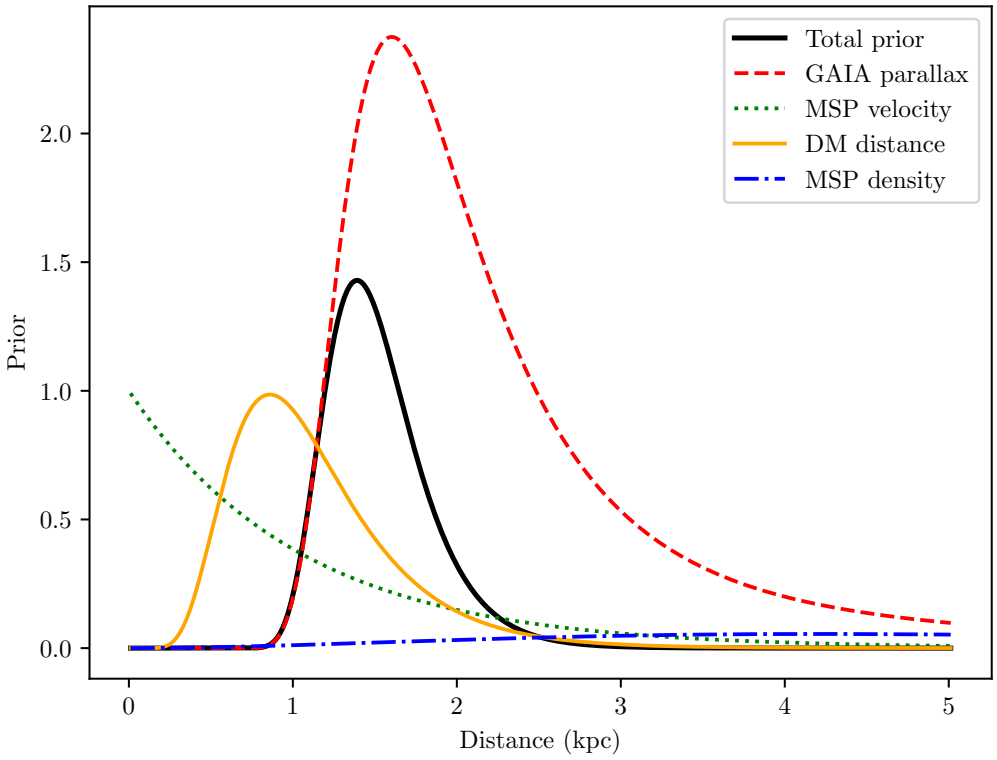

FIGURE 3.5: Distance prior distribution for J1227, in black, showing component distributions. The total prior distribution is dominated by the GAIA parallax term.

where K_1 is the radial velocity of the pulsar, and $R_2(f)$ is the volume-averaged radius of the companion star (in units of a) evaluated at a given RL filling factor, f. We therefore use this to constrain both the filling factor and companion system velocity. We obtain K_1 from radio timing (Roy et al., 2015b), while K_2 and $R_2(f)$ will be calculated by the `Icarus` model.

For J1023, we used the same top hat priors on T_0, $T_{\rm irr}$, and f as with J1227. While a well-constrained inclination can be derived from the results of Deller et al. (2012), $42 \pm 2°$, these calculations assume that the companion is Roche-lobe filling (Thorstensen & Armstrong, 2005a). As such we do not use any priors on the inclination. Deller et al. (2012) accurately determined the distance to the system from parallax measurements using long baseline radio interferometry to be $d = 1.368^{+0.042}_{-0.039}$ kpc. This is much more precise, with a lower uncertainty, than estimating the distance using the GAIA parallax method and so this was used to inform our distance modulus priors. For K_2, there are radial velocity measurements available from Shahbaz et al. (2019), however the heating of the companion distorts the radial velocity curve. This results in variable measurements of the radial velocity semi-amplitude. We take the measurements obtained from 2009 ISIS observations in the pulsar state of metallic absorption lines which correspond to a radial velocity

semi-amplitude of $K_2 = 276.3 \pm 5.6$ km s^{-1}. Instead of using these to inform a Gaussian prior on K_2, we use the data presented in their work to fit a radial velocity curve using the method described in section 3.4.1. From the same work we used the $V \sin(i)$ measurement of $V \sin(i) = 77.7 \pm 2.7$ km s^{-1} to constrain f and K_2, also from the 2009 ISIS observations of metallic absorption lines.

We also fit the V-band interstellar extinction separately for each source. For J1227 we use the prior value $A_V = 0.341$, calculated using the relationship $A_V = 3.1E(B-V)$ (Cardelli et al., 1989) from the colour excess presented in (de Martino et al., 2014b), in turn calculated from the N_H column density presented in de Martino et al. (2010b). We allow for a 20% uncertainty on this value. For J1023 we test two methods. First, we used the colour excess determined in Shahbaz et al. (2015) (again using the N_H column density) to calculate the extinction and the same relationship, obtaining $A_V = 0.2263$ and allowing for the same 20% uncertainty. Second, we use the Pan-STARRS dust maps of Green et al. (2018) to obtain $A_V = 0.109$ As these methods do not produce consistent extinction coefficients, we compared the evidence of models with each A_V prior and otherwise identical parameters. For both sources we use the reddening coefficients of Schlafly & Finkbeiner (2011) to calculate the appropriate extinction for each band.

3.4.3 Single-spot heating model

We extended the symmetrical model to include a single hot spot on the companion's surface. This was motivated by the successful modelling of similar asymmetries in other redback sources using a single-spot model (e.g., Shahbaz et al. (2017); Romani & Sanchez (2016); Nieder et al. (2020), with Clark et al. (2021) describing the implementation in `Icarus`). Our motivation for the inclusion of a hot spot is largely empirical, however we discuss two possible physical origins of the asymmetric heating in section 3.3; heating from X-rays reprocessed by a swept-back shock, and thermal winds on the companion surface.

Each spot introduces four free parameters: the spot temperature, $T_{\rm spot}$, the spot radius, $R_{\rm spot}$, and the spot position angles, θ_s, ϕ_s. θ and ϕ are the polar angle and azimuth angle, such that $\theta_s = 0^\circ$ is the North pole of the companion and $\phi_s = 90^\circ$ is the direction towards the L1 point. The spot temperature is added to the base temperature of the companion after the effects of gravity darkening, but before the irradiation such that $T_{\rm star}(\theta_s, \phi_s)^4 = (T_0 + T_{\rm spot}(\theta_s, \phi_s))^4 + T_{\rm irr}(\theta_s, \phi_s)^4$. In practice, there is little difference between this configuration and applying the spot after the irradiation; only the width of the spot would change. The spot geometry is defined by a 2D, axially symmetric Gaussian profile with a central maximum temperature of $T_{\rm spot}$ and width of $R_{\rm spot}$, with $R_{\rm spot} < 90^\circ$. We use a single spot as using more than

one is likely to over-fit the data, while also being more computationally expensive and significantly increasing the degeneracy of the parameter space. The priors on the spot temperature and radius are uniform over the ranges 0 K $< T_{\rm spot} <$ 10000 K and $0° < R_{\rm spot} < 90°$ respectively.

We note that the asymmetric heating caused by a hot spot on the leading edge of the companion (the side of the companion moving 'forwards' through the orbit) can also be described by a 'cold' spot on the trailing edge (the side moving 'backwards'). As such, we also performed fitting with a cold spot with a negative spot temperature, $T_{\rm spot} < 0$. We use uniform priors on $\cos\theta_s$ and uniform priors on the ϕ_s angle. To constrain the spot temperature and radius, we use Gaussian priors on the intensity of the flux from the spot, $I \propto T^4 R^2$, with a mean of $I = 0$ and a width of $\sigma_I = 10^{12}$. This was chosen in order to avoid a very small ($R_{\rm spot} \leq 5$ deg), very hot spot, since this prior favours cooler, larger spots.

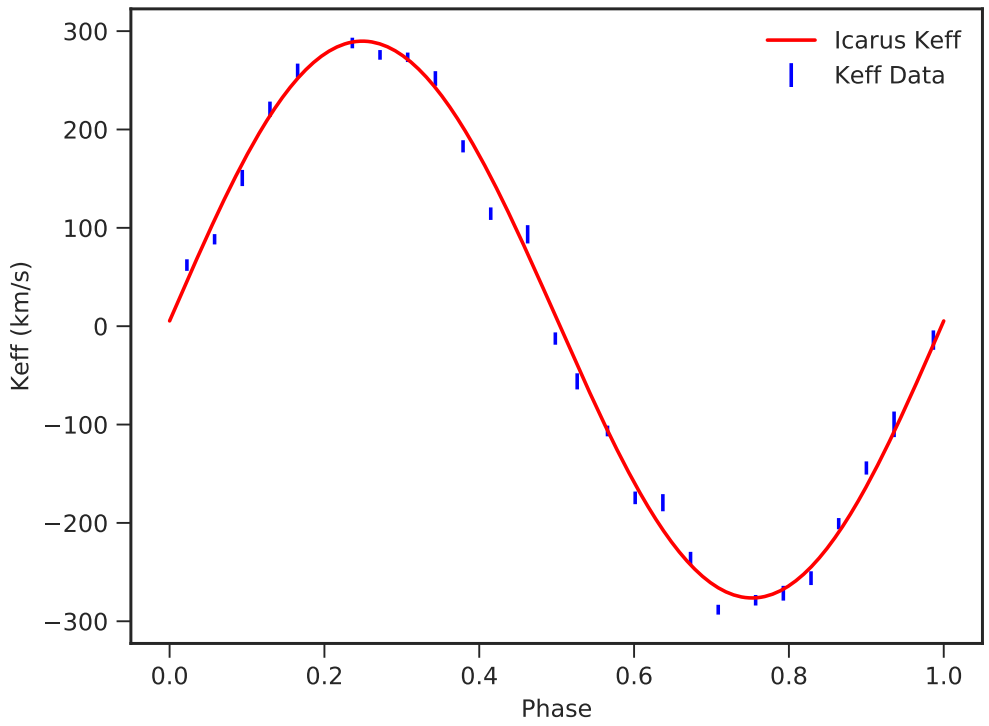

FIGURE 3.6: Radial velocity fitting curve for J1023. The points in blue are the radial velocity measurements from spectroscopy (Shahbaz et al., 2019), and in red is the best-fitting radial velocity curve from `Icarus`. The red RV curve is calculated at each step in the MultiNest sampler and the χ^2 value from the fit to the data is added to the posterior distribution. $K_{\rm eff}$ is the effective centre-of-light radial velocity of the companion.

3.4.4 Heat redistribution

As detailed in section 3.5, our modelling with both the symmetric and hot spot models produced results that were not reliable, and indeed indicated that neither model sufficiently describes the highly asymmetric light curves of these systems. As a result, we additionally used a further extension of the Random Forest code which directly models heat redistribution via diffusion and convection within the outer envelope of the companion. We implement the treatment in Voisin et al. (2020), though we acknowledge also the treatment of wind circulation in Kandel & Romani (2020) as a specific case of the latter. We continue our use of MultiNest evidence sampling to constrain the model parameters and heat redistribution laws. While the direct heating model assumed a constant companion base temperature (save for the effects of gravity darkening) and the hot spot extension assumes an additional, fixed, temperature source, the heat redistribution model allows for parallel (that is, with no radial component) energy transport within the outer shell. Energy transport follows the model (Voisin et al., 2020)

$$\vec{\nabla}_{\parallel} \cdot \vec{J}_{\parallel} = -\left(\sigma_{\rm sb}\left(T_*^4 - T_0^4\right) - L_{\rm w}\right), \tag{3.4}$$

which reduces to direct heating when the right-hand side is zero, and

$$\mathbf{J}_{\parallel} = -\kappa\left(\frac{T_*}{T_{\rm max}}\right)^{\Gamma}\nabla_{\parallel}T_* - T_* f(\theta)\sin(\theta)\mathbf{u}_{\phi}, \tag{3.5}$$

which is a generalisation of the parallel energy transport law derived in Voisin et al. (2020). $\mathbf{J}_{\parallel}$ is the surface energy flux, $\nabla_{\parallel}$ is the 'surface gradient', κ is the diffusion coefficient, T_* is the surface temperature of the companion, Γ is the diffusion index, $f(\theta)$ is the polar convection profile, and $\mathbf{u}_{\phi}$ is the unit vector of the longitude. $T_{\rm max}$ is an arbitrary constant which we define as $T_{\rm max}^4 = T_{\rm day}^4 = T_0^4 + T_{\rm irr}^4$. For this extension to the model, note that the spherical coordinates are defined differently to those in the hot spot model. θ is the colatitude, with the spherical coordinates defined such that the polar axis is the spin axis of the star, with $\phi = 0$ intersecting the binary axis on the night side of the star. Out of several forms of polar convection profile, $f(\theta)$, we initially chose $f(\theta) = v$, where v is the strength of the convection current in energy flux per unit temperature and is the first additional model parameter. This profile describes constant longitudinal advection. We also chose a convection profile of the form $f(\theta) = v\exp\left(-\frac{\theta^2}{2w^2}\right)$ which localises this flow to a Gaussian region around the equator with angular width w. We allow for the diffusion coefficient $\kappa = \kappa_{\rm max}\left(\frac{T_*}{T_{\rm max}}\right)^{\Gamma}$ to depend on the local temperature following a power law of index Γ. Here, $\kappa_{\rm max}$ is the value of κ when $T = T_{\rm max}$. We performed fits with both

a $\Gamma = 0$ for linear diffusion, and with the diffusion index as a free parameter. The diffusion coefficient, κ, is the last additional model parameter. When $\kappa = 0$, the diffusion is switched off and the model becomes convection only. In this case, this model is equivalent to that in Kandel & Romani (2020) when used with a 'bizone' convection profile. However, we do not use that profile in this work. We performed fits with both $\kappa = 0$ and κ as a free parameter.

3.5 Results

3.5.1 J1227

Standard model

As expected, this model was unable to account for the asymmetries in the light curves, reflected by a reduced chi-squared value of $\chi^2_\nu = 2.37$ for our best-fitting model and evidence of -5177.6. We note that the evidence provides little information on its own, but is included as a means to compare each model. As a result, the model is a poor fit and the best-fit parameters are likely erroneous; for example the sampler favours a nearly edge-on inclination of $i = 88.3^{+0.7}_{-1.0}{}^\circ$ converging at the upper limit of the prior. An inclination this close to 90° is highly unlikely as X-ray eclipses have not been observed (de Martino et al., 2014b). In light of this, we include an additional constraint in subsequent fits; the pulsar (and hence the inner region of the accretion disc) must not be eclipsed by the companion star at any orbital phase. This results in an upper limit on the inclination, around $i_{\max} \sim 77^\circ$, though the exact value depends on the filling factor and mass ratio parameters.

The companion velocity is derived from the inclination and K_2 (and therefore $K_{\rm eff}$) in our model, such that the fit with velocity $K_{\rm eff} = 269^{+5}_{-5}$ km s^{-1}, consistent with the de Martino et al. (2015) velocity, corresponds to an unrealistically low pulsar mass of $M_{\rm psr} \sim 0.9 M_\odot$. Loosening our prior on K_2 results in a more acceptable pulsar mass of $M_{\rm psr} \sim 1.2 M_\odot$ but a companion radial velocity of $K_{\rm eff} = 308^{+12}_{-20}$ km s^{-1}, which is clearly not consistent with the literature. The companion temperature ($T = 5452^{+19}_{-20}$ K) and irradiation temperature ($T_{\rm irr} = 5230^{+26}_{-28}$ K) do agree with those determined in (de Martino et al., 2015), as these parameters are primarily influenced by the colour information rather than the shape of the light curves. However we do not reproduce their filled RL ($f = 1.0$); instead we determine $f = 0.825^{+0.002}_{-0.002}$, corresponding to a volume-average filling factor of $f_{\rm VA} = 0.95$. Note also that their modelling uses a symmetric model. Considering

TABLE 3.2: Numerical results for the modelling of J1227; including from top to bottom the model parameters, selected derived parameters, and model statistics. From left to right; standard symmetric model, single hot spot model, heat redistribution with linear diffusion and constant advection profile (HR1), and heat redistribution with convection only (HR2). $\log Z$ is the natural logarithm of the model evidence and χ^2_ν is the reduced chi squared value with $\nu = 5908$ datapoints. Note that the inclination of the symmetrical model is higher than the others as the constraint described in section 3.5.1 is not applied.

Parameters	Symmetrical	Single spot	HR 1	HR2
Inclination, i (°)	$88.3^{+1.3}_{-2.0}$	$76.6^{+0.6}_{-1.2}$	$77.1^{+0.1}_{-0.15}$	$77.0^{+0.2}_{-0.3}$
Mass ratio, q	$5.6^{+0.1}_{-0.1}$	$5.8^{+0.1}_{-0.1}$	$5.48^{+0.07}_{-0.07}$	$5.5^{+0.1}_{-0.1}$
Radial velocity, K_2 (km s^{-1})	282^{+5}_{-5}	294^{+4}_{-5}	277^{+3}_{-4}	277^{+5}_{-5}
Filling factor, f	$0.825^{+0.002}_{-0.002}$	$0.838^{+0.003}_{-0.003}$	$0.850^{+0.002}_{-0.002}$	$0.852^{+0.002}_{-0.002}$
Base temp., T_0 (K)	5452^{+19}_{-20}	5556^{+14}_{-15}	5584^{+11}_{-11}	5585^{+15}_{-15}
Irradiation temp., $T_{\rm irr}$ (K)	5230^{+26}_{-28}	5312^{+21}_{-21}	5479^{+16}_{-15}	5489^{+22}_{-21}
Spot temp., τ (K)	-	2100^{+200}_{-200}	-	-
Spot radius, ρ (°)	-	$7.8^{+0.6}_{-0.5}$	-	-
Spot polar angle, θ (°)	-	95^{+2}_{-2}	-	-
Spot azimuth angle, ϕ (°)	-	-27^{+1}_{-1}	-	-
Diffusion coeff., κ (W K^{-1} m^{-2})	-	-	-95^{+7}_{-3}	0
Diffusion index, Γ	-	-	0	-
Convection amp., v (J m^{-2} K^{-1})	-	-	3230^{+43}_{-45}	3260^{+22}_{-21}
Volume-averaged f, $f_{\rm VA}$	$0.952^{+0.001}_{-0.001}$	$0.959^{+0.001}_{-0.001}$	$0.965^{+0.002}_{-0.002}$	$0.966^{+0.001}_{-0.001}$
Effective radial velocity, $K_{\rm eff}$ (km s^{-1})	269^{+5}_{-5}	277^{+4}_{-5}	263^{+3}_{-3}	263^{+5}_{-5}
Distance, d (kpc)	$1.64^{+0.09}_{-0.09}$	$1.82^{+0.09}_{-0.09}$	$1.84^{+0.06}_{-0.06}$	$1.79^{+0.1}_{-0.1}$
Pulsar mass, $M_{\rm psr}$ ($M_\odot$)	$0.93^{+0.04}_{-0.04}$	$1.13^{+0.05}_{-0.04}$	$0.96^{+0.03}_{-0.03}$	$0.96^{+0.05}_{-0.04}$
Blackbody day temp. (K)	6040	6170	6200	6210
Blackbody night temp. (K)	5310	5410	5430	5430
Reduced chi-squared, χ^2_ν	2.37	1.11	1.38	1.38
Model evidence, $\log Z$	-5177.6	-3345.1	-3906.4	-3906.0

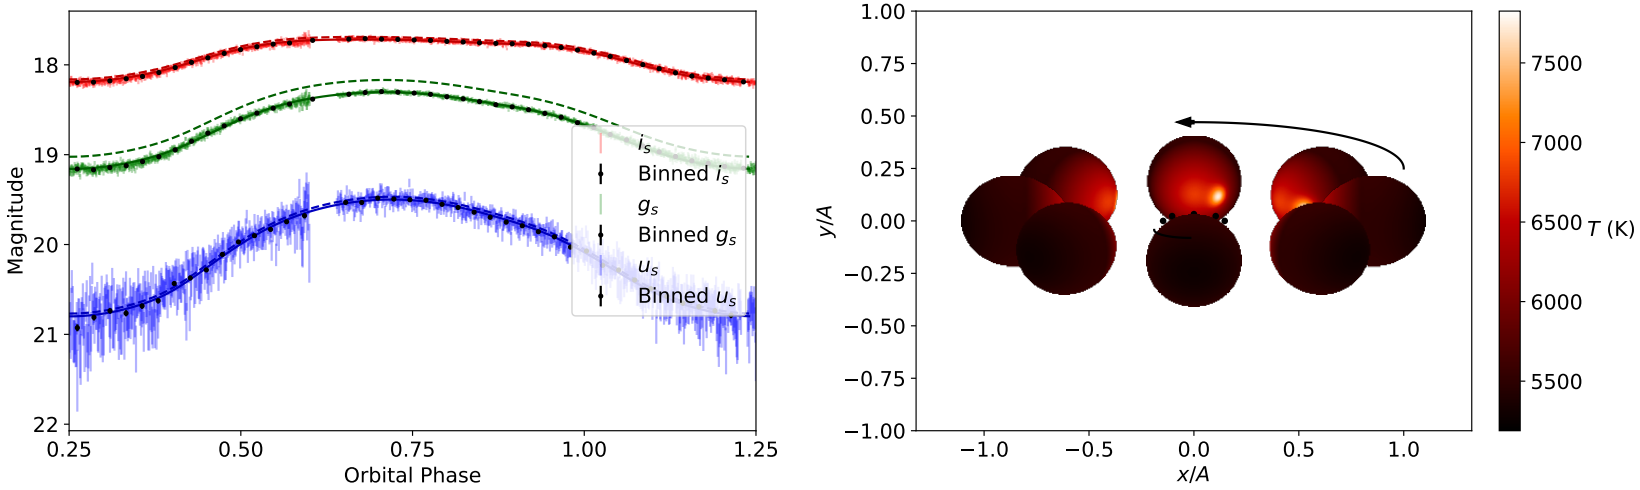

FIGURE 3.7: *Left:* i_s, g_s, and u_s light curves of J1227 overlaid with the best fit hot spot model. The dashed lines show the model light curve without band calibration corrections. *Right:* Temperature distribution of companion - note how the inclination is such that the companion only just avoids eclipsing the pulsar at inferior conjunction.

these discrepancies in addition to a poor overall fit, we conclude that the standard `Icarus` heating model is not appropriate for this source.

Hot spot model

The hot spot heating model was able to account for the majority of the asymmetry with $\chi^2_\nu = 1.11$ and evidence of -3345.1, indicating a significantly improved fit over the symmetric model. We find that the filling factor determined by our best-fit model, $f = 0.838 \pm 0.003$ ($f_{\rm VA} = 0.958 \pm 0.003$), also indicates that the companion is not Roche-lobe filling and is consistent with the symmetric model result. The system distance, $1.82^{+0.09}_{-0.09}$ kpc and best-fit temperatures, $T_0 = 5556^{+14}_{-15}$ K and $T_{\rm irr} = 5312^{+21}_{-21}$ K, are also broadly similar to the symmetric model.

As with the direct heating model, the model prefers a $K_{\rm eff}$ comparable to that obtained from the spectroscopy, while the pulsar mass is unreasonably low, at 0.93 ± 0.04 $M_\odot$. We therefore constrained the pulsar mass to the range $1.0{\rm M}_\odot < M_{\rm psr} < 3.0{\rm M}_\odot$ for a repeat of this fit. However, the companion effective radial velocity with this constraint, 307^{+10}_{-10} km s^{-1}, is unacceptably large compared to the spectroscopic measurement of de Martino et al. (2014b). This occurs as K_2 increases to compensate for the high inclination, increasing q as well.

The inclination again indicates a nearly edge-on system with a best-fit value of $76.6^{+0.6\circ}_{-1.2}$. The shape of the posterior distribution is skewed, showing that the model has converged with an inclination very close to the limit imposed by the eclipse limit. This is illustrated in the corner plot in figure 3.8.

While the hot spot model provides an improved fit to the data, the best-fit values of the inclination, pulsar mass, and companion radial velocity are not suitable.

Cold Spots

The asymmetry of the light curve of J1227 is stronger in the cooler i_s and g_s bands than the u_s band, suggesting that a cold spot may be better suited to model the asymmetry. We repeated the analysis in presented section 3.5.1 using a cold spot model and found that the same trends were present. However, the distances and masses are larger than with the hot spot and for all inclinations the fitting is poorer. Notably, the effective $K_{\rm eff}$ velocities are consistently more than 1-sigma larger than the spectroscopic K_2, whereas these velocities were consistent when using the hot spot. Comparing the evidence for the $i = 60^\circ$ run as an example, this was -3933.7 for the cold spot compared to -3600.8 for the hot spot, indicating a less favourable model. These factors indicate that the hot spot model is preferred over the cold spot for J1227.

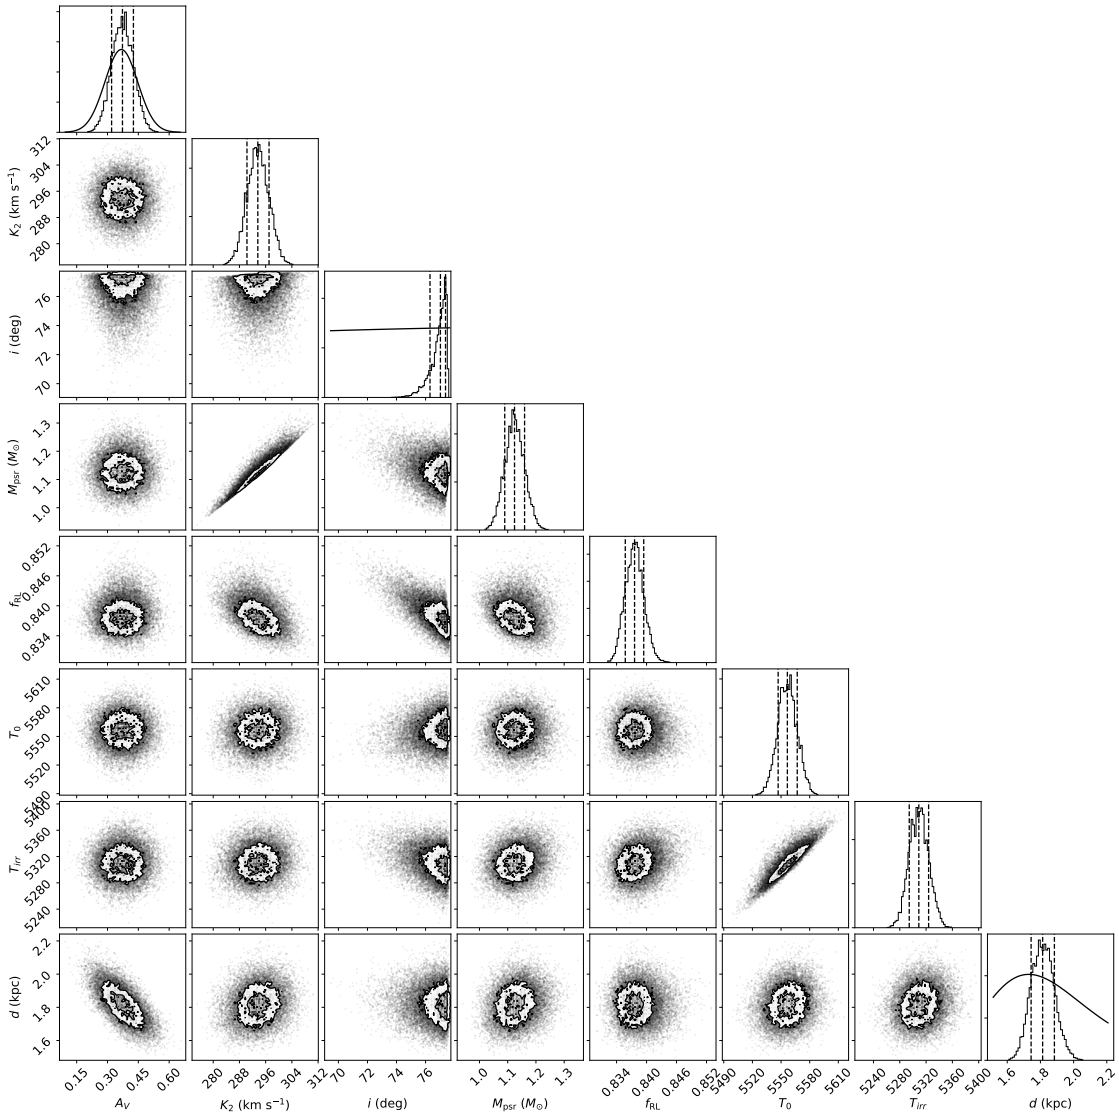

FIGURE 3.8: Corner plot of selected parameters of the J1227 hot spot model. Not shown are the hot spot parameters and the mass ratio. The plots along the diagonal are the posterior distributions of the parameters; the solid black lines on these are the prior distributions. The remaining plots show the position of walkers, illustrating covariance between parameters. Note the posterior distribution of the inclination, with the walkers converged against the upper limit set by the zero eclipse width.

Heat redistribution

We performed several fits using the heat redistribution model introduced in section 3.4.4, though none improved on the model parameters or evidence of the best-fitting hot spot model; that is, the edge-on inclination and small pulsar mass were still favoured. These results are presented in table 3.2 alongside the results from the symmetric and single spot models. Note the similarity between the parameters of the two models, which results in indistinguishable model light curves. While the diffusion coefficient (κ) must be positive, we used a lower bound of -100 to avoid boundary effects around 0. As such, the negative value of diffusion coefficient in model HR1 (linear diffusion and convection) is unphysical and suggests a true value of 0. Additionally the uncertainty of this value is likely underestimated as the posterior distribution converged on the boundary. Indeed, the model parameters of HR1 and HR2 are otherwise consistent within uncertainties, suggesting that the model with convection only is a better description of the system. However, the model evidence and χ^2 in both cases favours the hot spot model.

Modelling assuming a filled Roche lobe

We also performed fits where the Roche Lobe is filled, with $f = 1.0$ as in de Martino et al. (2014b). Under this assumption, the symmetric model provided an unsatisfactory fit similar to the f-free symmetric model and so was discarded. Modelling with a hot spot returned an acceptable fit with a evidence of -3871.8 and a reduced χ^2 of 1.32; comparable but not better than the f-free case. However, the best-fit value of K_2 corresponds to a velocity of $K_{\rm eff} = 217 \pm 5$ km s^{-1}, which is more than 8 standard deviations from the spectroscopic measurement.

Modelling with fixed inclination

With the inclination otherwise unconstrained, both the standard and hot spot models strongly favour an edge-on system with an inclination close to 90°. A system this edge-on is ruled out by the non-detection of X-ray eclipses, suggesting an inclination of $i \lesssim 73°$ (de Martino et al., 2014b), and the lack of eclipses seen in spectra. Furthermore, when the radial velocity constraint from de Martino et al. (2014b) is enforced the model returns a pulsar mass in the range $M_{\rm psr} \sim 0.8 - 1.0 M_\odot$ which is clearly inappropriate. Conversely, relaxing this constraint we obtain a reasonable pulsar mass of $M_{\rm psr} \sim 1.4 M_\odot$ but a value for K_2 which is too large, i.e. ~ 330 km s^{-1}. To attempt to overcome these discrepancies, we performed modelling of the system with the inclination fixed at each of $i = 40°$, 50°, 60°, 70°, using the hot spot model.

Broadly, we observe that the pulsar mass, T_0, and filling factor are negatively correlated with inclination, while the system distance and irradiation temperature are positively correlated. The other parameters are not affected within uncertainties. We summarise these results in table 3.3, though we do not cover the full results of the fit with $i = 40^\circ$ as the pulsar mass of $M_{\rm psr} \geq 3.0 M_\odot$ and the distance of $d = 2.41 \pm 0.12$ kpc are both unreasonably large, suggesting that inclinations this low can be safely discarded. It is worth noting that while the fit value of K_2 does not seem to be correlated with the inclination, the effective value appears to have a slight positive correlation with the inclination. However, the $K_{\rm eff}$ is consistent with the de Martino et al. (2014b) radial velocity for inclinations 50°, 60°, and 70°, suggesting that the prior on the radial velocity is still tightly constraining. Likewise, the RL is consistently under-filled at all inclinations. We note that the evidence of each fit is also correlated with the inclination in a direction that suggests the model favours a more edge-on system, echoing what we observed when modelling with the inclination unconstrained: an edge-on inclination is favoured despite strong penalties from priors. We attempted modelling with the inclination tightly constrained rather than fixed, at $60 \pm 1^\circ$, however the inclination did not converge to a solution after $\sim 5\times$ the usual computing time, suggesting this configuration is not appropriate. While this investigation shows the behaviour of the model at more face-on inclinations, the reason that the unconstrained inclination consistently converges to $i \sim 90^\circ$ is still unclear.

We constructed a mass-mass plot using the results in table 3.3, shown in figure 3.9. This plot reveals constraints that we can apply to the masses of the companion and pulsar. We calculate a lower limit to the mass ratio, $q_{\rm min}$, from the centre-of-light $K_{\rm eff}$ radial velocity. Since in `Icarus`, the mass ratio is calculated from the larger, centre-of-mass K_2, this $q_{\rm min}$ acts as a lower bound to the pulsar mass at each inclination. These K_2 velocities at each inclination therefore correspond to the best fit mass ratio, and are consistent across the whole range to within the uncertainties. To further constrain the object masses, we interpolated the model distance estimates at each inclination in the allowed region of the plot. Then, by using the same prior distribution used during the fitting (a combination of GAIA parallax, galactic MSP density, and velocity distributions and the DM distance), we calculated the confidence interval over the mass-mass plane. This shaded region indicates that the more edge-on inclinations produce more favourable distances. Indeed, these also correspond to better χ^2_ν values and evidence. Using the evidence alone, the most favoured inclination is $i = 70^\circ$, however this is a comparatively worse evidence than the best fitting free inclination model. However, we can conclude that

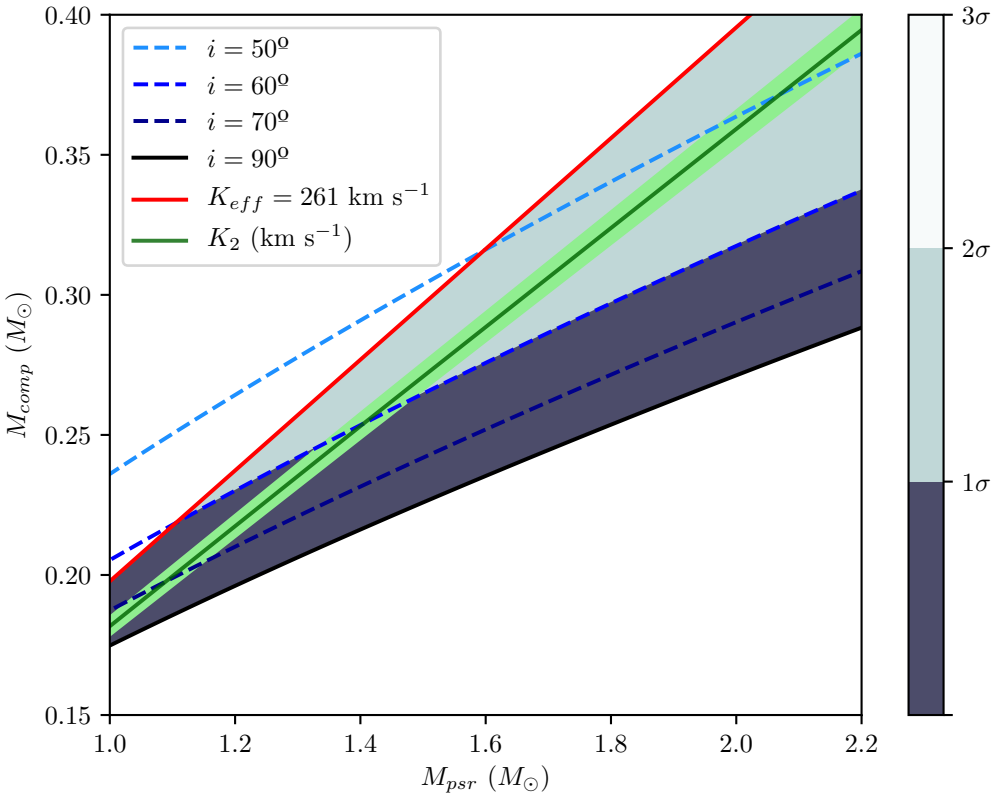

FIGURE 3.9: The mass ratio calculated from the centre-of-light $K_{\rm eff}$ value, in red, provides us with a lower limit of the pulsar mass for each inclination, while the best-fit mass ratio from each fit's K_2 is in green. The mass function at each inclination is shown by the series of blue dashed lines, with the absolute limit at $i = 90°$ shown by a solid black line. As such, the excluded region of the mass-mass plot is in white. The best-fit model distance at each inclination is interpolated over the mass-mass plane. The shaded area then represents the corresponding confidence interval of the distances with respect to the distance prior.

TABLE 3.3: Model parameters and MultiNest evidence for fixed-inclination modelling of J1227 at $i = 50°, 60°, 70°$. This model has $\nu = 5908$ datapoints.

Model Parameters	$i = 50°$	$i = 60°$	$i = 70°$
f	$0.901^{+0.02}_{-0.02}$	$0.866^{+0.002}_{-0.002}$	$0.841^{+0.002}_{-0.002}$
$M_{\rm psr}$ ($M_\odot$)	$2.1^{+0.1}_{-0.1}$	$1.41^{+0.06}_{-0.06}$	$1.09^{+0.05}_{-0.05}$
d (kpc)	$2.1^{+0.1}_{-0.1}$	$1.8^{+0.1}_{-0.1}$	$1.67^{+0.09}_{-0.09}$
T_0 (K)	5233^{+16}_{-16}	5288^{+15}_{-17}	5329^{+17}_{-19}
$T_{\rm irr}$ (K)	5356^{+26}_{-26}	5184^{+23}_{-26}	5080^{+26}_{-28}
K_2 (km s^{-1})	282^{+5}_{-5}	280^{+5}_{-5}	280^{+5}_{-5}
$K_{\rm eff}$ (km s^{-1})	258^{+5}_{-5}	262^{+5}_{-5}	264^{+5}_{-5}
χ^2_ν	1.43	1.30	1.24
$\log Z$	-4099.9	-3837.3	-3690.3

for the range of inclinations $i \sim 50° - 70°$, we obtain a pulsar mass in the range $M_p \sim 1.09M_\odot - 2.0M_\odot$ and a companion mass in the range $M_c \sim 0.2M_\odot - 0.37M_\odot$. The lower-bound pulsar mass of $M_{\rm psr} \approx 1.09M_\odot$ at $i = 70°$ would make J1227 one of the least massive known MSPs (see, e.g., Romani et al., 2016; Clark et al., 2021 and indeed very close to the lowest possible pulsar mass under current formation mechanisms.

Since the mass estimates at more face-on inclinations more closely resemble those found in spiders, this is at odds with the better evidence and distance estimate at 70°. Considering these discrepancies we may surmise that the hot spot model is not a complete description of the system.

3.5.2 J1023

Standard model

As with J1227, the symmetric heating model did not provide a good fit to the data, with a best-fit reduced χ^2 value of $\chi^2_\nu = 6.20$ ($\nu = 3821$) and evidence of -5295.0. This is a comparatively worse fit than the same modelling of J1227, in part because smaller error bars from a brighter source and better observing conditions make the asymmetry more significant compared to the noise, with the u_s-band fit especially poorly. The best-fit parameters are summarised in table 3.4. We determine an inclination angle of $i = 46.4^{+0.5°}_{-0.7}$, which is consistent with the range of possible inclinations of Archibald et al. (2009), $34° < i < 53°$, and consistent with the inference

TABLE 3.4: Numerical results for the modelling of J1023; including from top to bottom the model parameters, selected derived parameters, and model statistics. This model has $\nu = 3821$ datapoints. From left to right; standard symmetric model, single hot spot model, heat redistribution with linear diffusion and constant advection profile (HR1), and heat redistribution with convection only (HR2). Note the similarity between the blackbody temperatures of each model, despite the large range of T_0 and $T_{\rm irr}$ temperatures.

Parameters	Symmetrical	Single spot	HR 1	HR2
Inclination, i (°)	$46.4^{+0.5}_{-0.7}$	$45.1^{+0.8}_{-0.9}$	$48.7^{+0.20}_{-0.23}$	$45.3^{+0.6}_{-0.5}$
Mass ratio, q	$7.8^{+0.05}_{-0.05}$	$7.8^{+0.1}_{-0.1}$	$7.87^{+0.04}_{-0.04}$	$7.89^{+0.05}_{-0.05}$
Radial velocity, K_2 (km s^{-1})	295^{+3}_{-3}	295^{+3}_{-3}	$297.0^{+2.3}_{-2.2}$	$298.0^{+2.5}_{-2.5}$
Filling factor, f	$0.86^{+0.01}_{-0.01}$	$0.94^{+0.02}_{-0.02}$	$0.808^{+0.003}_{-0.003}$	$0.864^{+0.007}_{-0.007}$
Base temp., T_0 (K)	5580^{+14}_{-13}	4885^{+31}_{-30}	4477^{+107}_{-43}	5500^{+10}_{-10}
Irradiation temp., $T_{\rm irr}$ (K)	4903^{+26}_{-23}	4677^{+48}_{-45}	7351^{+60}_{-150}	4867^{+15}_{-15}
Spot temp., τ (K)	-	1134^{+32}_{-32}	-	-
Spot radius, ρ (°)	-	48^{+2}_{-2}	-	-
Spot polar angle, θ (°)	-	$8.2^{+1.0}_{-0.8}$	-	-
Spot azimuth angle, ϕ (°)	-	-57^{+5}_{-5}	-	-
Diffusion coefficient, κ (W K^{-1} m^{-2})	-	-	93500^{+4000}_{-8000}	0
Diffusion index, Γ	-	-	0	-
Convection amplitude, v (J m^{-2} K^{-1})	-	-	52000^{+2000}_{-4000}	8630^{+60}_{-60}
Volume-averaged f, $f_{\rm VA}$	$0.969^{+0.009}_{-0.008}$	$0.994^{+0.004}_{-0.003}$	$0.942^{+0.002}_{-0.002}$	$0.971^{+0.004}_{-0.004}$
Effective radial velocity, $K_{\rm eff}$ (km s^{-1})	287^{+3}_{-3}	283^{+3}_{-3}	$291.0^{+2.3}_{-2.2}$	$290.0^{+2.5}_{-2.5}$
Distance, d (kpc)	$1.26^{+0.02}_{-0.01}$	$1.28^{+0.02}_{-0.02}$	$1.19^{+0.01}_{-0.01}$	$1.30^{+0.02}_{-0.02}$
Pulsar mass, $M_{\rm psr}$ ($M_\odot$)	$1.76^{+0.06}_{-0.05}$	$1.89^{+0.10}_{-0.09}$	$1.69^{+0.03}_{-0.03}$	$1.62^{+0.06}_{-0.06}$
Blackbody day temp. (K)	5820	5700	5750	5750
Blackbody night temp. (K)	5470	5380	5420	5410
χ^2_ν	6.20	1.20	1.26	1.45
$\log Z$	-5295.0	-2293.9	-2462.6	-2633.5

in Thorstensen & Armstrong (2005a) that $i < 55°$. Note that these literature values were calculated under the assumption that the companion is Roche lobe-filling. We again obtain a RL filling factor of significantly less than 1.0 though there is a strong negative correlation with the inclination, such that at an inclination consistent with the Archibald et al. (2009) calculation, the filling factor approaches $f = 1.0$. The system temperatures indicate that the companion is not as strongly irradiated as in J1227. A best-fit distance of $d = 1.26^{+0.02}_{-0.01}$ kpc broadly agrees with the radio parallax measurement of $1.368^{+0.042}_{-0.039}$ kpc from Deller et al. (2012) Note that across all models the distance is consistently underestimated; this is discussed in section 3.6.4. This fit produces a radial velocity of $K_{\rm eff} = 287 \pm 3$ km s^{-1}, corresponding to a mass ratio of $q = 7.8 \pm 0.05$ which is comparable but not consistent with the result from radio timing in Archibald et al. (2009). This is as expected, as the radial velocity parameter appears to be consistent with the upper bound of measurements from the spectral lines in Shahbaz et al. (2019), and the mass ratio is calculated directly from this velocity. Unlike with our modelling of J1227 the posterior distributions of these model parameters, notably the mass, generally converged within the range of expected literature values. However, the large χ^2 value leads us to again conclude that the symmetric model is insufficient.

Hot spot model

With $\chi^2_\nu = 1.20$ and evidence of -2293.9, the hot spot model provides a much better fit to the data than the symmetric model, capturing the asymmetries to a good degree. The parameters of the best-fitting model are shown in table 3.4. The filling factor remains below 1.0 and the inclination is again consistent with that inferred from radio timing from Archibald et al. (2009). The distance is also broadly consistent with the interferometry distance in Deller et al. (2012); 1.28 ± 0.02 kpc compared to 1.368^{+42}_{-39} kpc. Considering the best-fit spot parameters, the spot radius of $48 \pm 2°$ and position near the companion pole is strikingly similar to the spot properties seen in Kandel et al. (2020). The temperature distribution of the companion is shown alongside the best-fitting model in figure 3.10. We expect tMSPs to host massive neutron stars, and the constrained mass of $M_{\rm psr} = 1.89^{+0.10}_{-0.09}$ is no exception. Note however that there is a moderate inconsistency between the determined pulsar masses.

Heat redistribution

Choosing a diffusion index of 0 (linear diffusion) and a constant advection polar convection profile, we ran an initial model of the light curve. We found that the

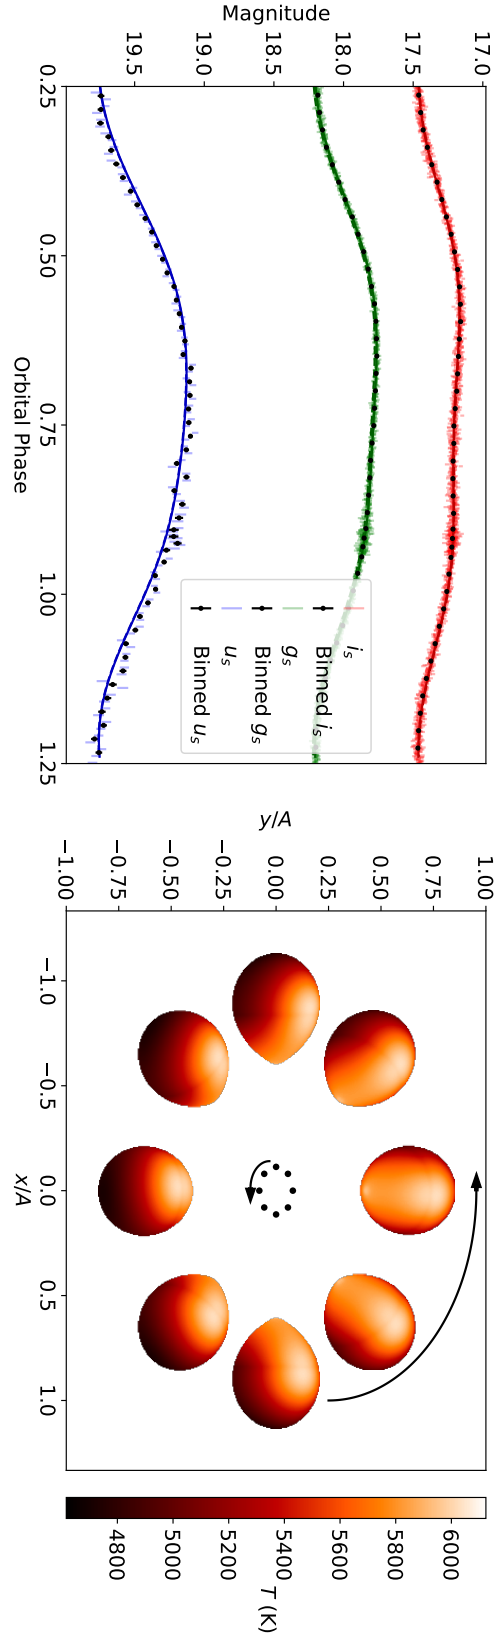

FIGURE 3.10: *Left:* i_s, g_s, and u_s light curves of J1023 overlaid with the best fit hot spot model, with 40 phase bins in black. *Right:* Temperature distribution of companion, showing the asymmetry which manifests as a large, polar spot. The dark band around the star is an artefact due to the plotting only.

asymmetry was well-fit, but with a marginally worse evidence and reduced χ^2 than the best fitting hot spot model.

Since the heat redistribution model also allows for a range of combinations of diffusion and convection, the results of *(1)* a model with convection only, and *(2)* convection with linear diffusion (i.e. with a zero diffusion index) are shown in table 3.4. However, model *(3)*, with convection and temperature-dependent diffusion converged to a solution with a very large diffusion index which caused significant aliasing in the temperature distribution of the companion; as such these results have not been included. Comparing the other two models, the evidence favours a model with convection and linear diffusion, despite the fact that the best-fit distance for this model is significantly below the literature value.

3.6 Discussion

3.6.1 Filling factor

One consistency across all models is an underfilled Roche lobe, with our modelling returning values in the range $f \sim 0.825 - 0.90$ for J1227 and $f \sim 0.81 - 0.94$ for J1023. This is at odds with our expectations for these sources; as they are both tMSPs we would expect the Roche lobes to be full or nearly filled in the rotation-powered (RP) states given they have transitioned from Roche lobe overflow in their accretion-powered (AP) states. However, when considering the volume-averaged fill factor, the values of $f_{\rm VA} = 0.958$ and 0.994 for the best-fitting models of J1227 and J1023 respectively tell a story more consistent with our initial expectations; that the Roche lobe is indeed mostly full. Figure 3.11 illustrates the relationship between these two parameterisations. When considering the proximity of the companion star to a state of RL overflow, f is a more useful description as it gives the distance of the star surface to the L1 point. However, the volume-averaged filling factor gives a clearer picture of the size of the star relative to its RL, in these cases illustrating that a substantial increase in the volume-averaged radius of the star is not necessarily required for RL overflow to begin. For J1023 in particular, the volume-averaged fill factor approaches unity. With a nearly-full Roche lobe in the RP state, we may consider mechanisms for transitions.

We may assume that if the mass transfer is conservative, then the orbital separation should increase during the AP state in order to conserve angular momentum. This would correspond with an expansion of the companion's Roche lobe. While orbital period variations in the RP state have been observed (Archibald et al., 2013),

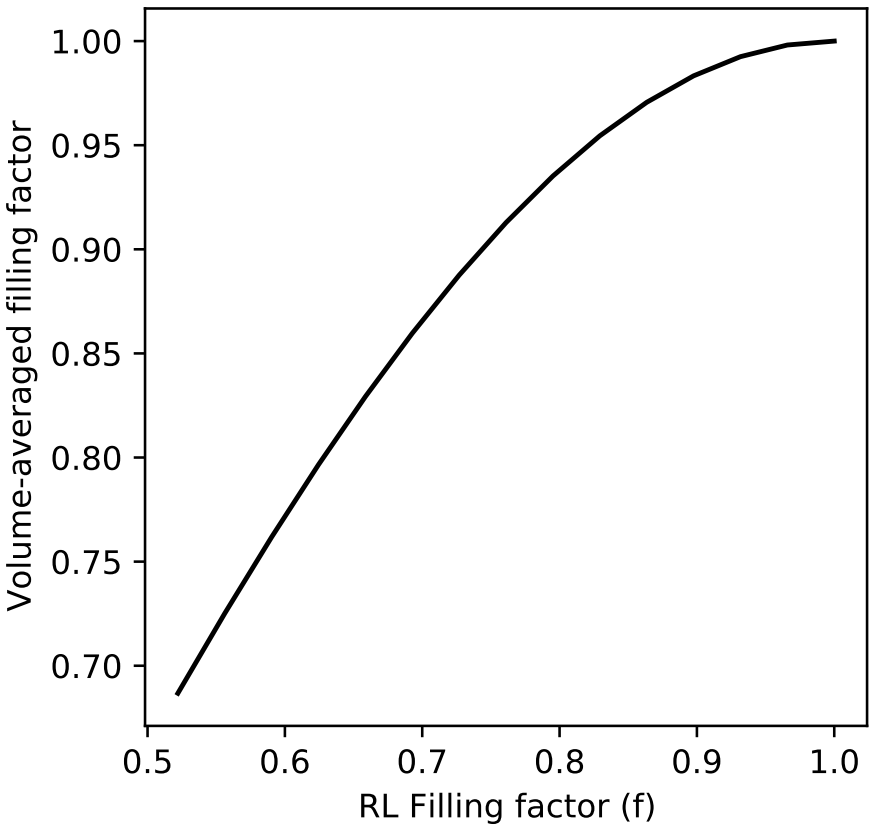

FIGURE 3.11: Relationship between filling factor, f, and the volume-averaged filling factor. At lower values of f ($f < 0.6$) the relationship is linear. Considering filling factors for $f > 0.8$, it can be seen how an apparently underfull Roche lobe can have a volume-averaged filling factor much closer to unity due to the tidally distorted shape of the companion.

these are over the timescale of ~ 100 days, shorter than the timescale between transitions of $\sim 1-10$ years. Furthermore, these variations are not of sufficient magnitude to cause the required changes in Roche lobe radius. The observed variation in $T_{\rm asc}$ are of order ~ 1 second. With the relation $\Delta T_{\rm asc} \sim \dot{P}_{\rm orb} t_{\rm obs}$, where $t_{\rm obs} \sim$ 100 d is the observation time, this corresponds to a change in a of $\Delta a/a \sim 10^{-5}$, a factor of $\sim 10^5$ too small. We assume that a fraction change in a is equivalent to the same fractional change in the Roche lobe radius.

Lastly, there appear to be no correlated changes in the orbital period over the several years of data, suggesting the changes are not gradual and continuous as would be the case with steady mass loss from ablation. Instead, a change in the structure of the star without significant mass loss could be explained by the size of the convective envelope decreasing or disappearing completely as the system transitions from the AP to RP state. We assume that the companion stars in redback systems have large convective envelopes, so it is possible that these are 'puffed up' while in the RP state. The irradiation of the companion by the pulsar wind is known to expand the companion photosphere, however it is uncertain if this could occur sufficiently within the transition timescale. Similarly, if an accretion disc shields the secondary from the pulsar irradiation, it could allow the companion photosphere to gradually contract during the AP state. Consider the Kelvin-Helmholtz mechanism,

$$\tau_{\rm KH} \sim \frac{GM^2}{2RL} \sim \frac{2G\rho M}{3\sigma_{\rm SB} T^4}, \tag{3.6}$$

where G is the gravitational constant, M is the mass of the star, R is its radius, L is its luminosity, ρ is the mean density of the star, T is its temperature, and $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant (Kippenhahn et al., 2012). Assuming a mean stellar density of 1 g cm^{-3}, a companion mass of 0.4 $M_\odot$, and a temperature of 5500 K, we obtain a timescale of $\tau_{\rm KH} \sim 2 \times 10^7$ yr. This is significantly longer than the $\sim$ yr timescales between tMSP transitions. If we consider the contraction of only the outer convective layer, some fraction k_R of the stellar radius, it is possible a thin layer of the companion could contract and expand within transition timescales. However, this is highly dependent on the depth of this necessarily thin layer. In light of this, we consider the discussion of the envelopes of asymptotic giant branch (AGB) stars in Soker (2015). Following equation 2 in Soker (2015), we separate the companion star into a core of mass $M_{\rm core} = f_c M_c$ and envelope $M_{\rm env} = (1 - f_c) M_c$, where f_c is a fraction between 0 and 1 Assuming the radius of the star is equal to the Roche lobe radius, calculated using the Eggleton approximation (Eggleton, 1983) with $q = 7.8$, we calculate the Kelvin-Helmholtz timescale for this envelope to be $\tau_{\rm KH} \sim f_c(1 - f_c) \times 100$ kyr. To obtain a timescale of order ~ 10 yr, as expected

from the tMSP transition timescale, we arrive at a mass fraction of $f_c \sim 10^{-4}$. This is a plausible result, suggesting that a fraction of the envelope can be expected to contract within tMSP transition timescales.

Considering the corner plot in figure 3.8, showing the J1227 hot spot model parameters, strong covariance between the filling factor and inclination can be seen. This covariance indicates a negative correlation, such that for a lower (more face-on) inclination, the filling factor would increase towards unity. Given that we suspect the inclination of J1227 to be over-estimated in our model, this may suggest that the star may be even closer to filling its Roche lobe than inferred. A similar relationship between the distance and filling factor is seen for J1023, where a distance closer to the interferometry measurement results in a filling factor closer to unity.

3.6.2 Asymmetries

For both sources the hot spot model clearly provides a better fit than the symmetric model — both in comparing the χ^2_ν values, the Bayesian evidence, and the model parameters — but it is clear that it is still not a complete description of systems with asymmetric light curves. This is evidenced by, for example, the model favouring an edge-on inclination for J1227 in all situations, despite strong penalties from priors. Further to this, for both sources the residuals in the u_s-band show systematic variations, implying that some aspects of the light curves are not captured by the model. The per-band residuals for the best-fitting hot spot model of J1023 are shown in figure 3.14. Examining these in more detail we find the model produces too much asymmetry in the u_s band compared to the data, while for the i_s and g_s bands the residuals show that the model is capturing the asymmetry well. Why this happens is not immediately clear While it may be interpreted as further evidence that the hot spot model is insufficient, similar systematics are also present in the best-fitting heat redistribution model.

As is described in section 3.4.4, our heat redistribution model incorporates the effects of diffusion and convection in the companion photosphere. Our results using this model are of comparable quality to the hot spot model, the Bayesian evidence indicating an improved fit over the direct heating model. However, the temperature distributions obtained differ significantly from those in the hot spot model, for example for J1023 for where the polar spot disappears in the heat redistribution model. The temperature distribution of the best-fit heat redistribution model is shown in figure 3.12. A possible model not tested in this work may be a combination of the two models; a hot spot with heat redistribution.

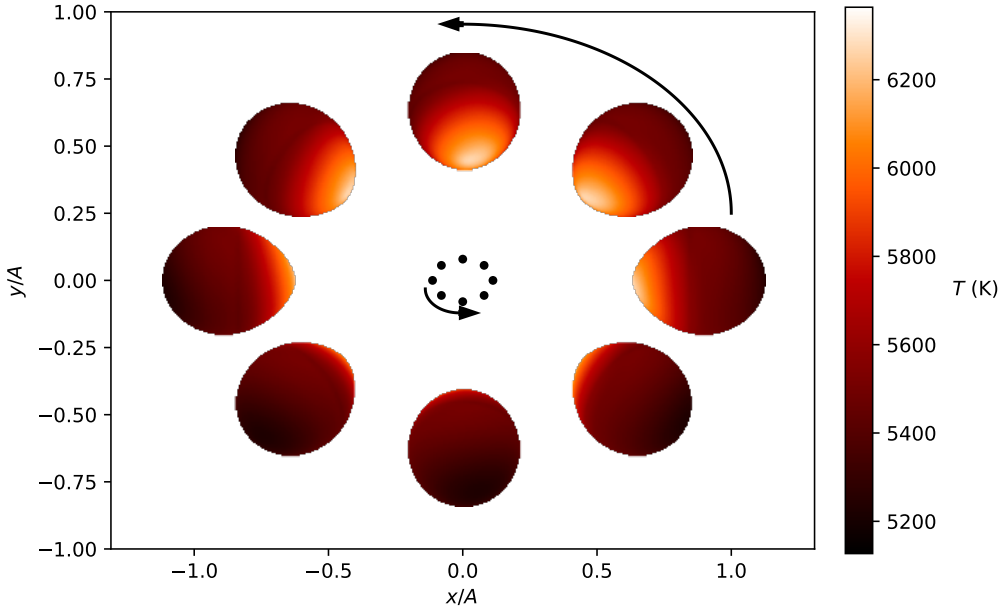

FIGURE 3.12: Companion star temperature distribution for the best-fit heat redistribution (HR1) model of J1023. The asymmetry of this distribution can be seen near the L1 point, and the lack of the polar spot seen in the hot spot model is no longer present. The model light curve and residuals are indistinguishable from the hot spot case.

3.6.3 J1227

Despite the modelling failing to provide a single best fitting solution, we can discuss several key findings. Our inclination range of $i \sim 50°-70°$ obtained by fitting at fixed inclinations agrees with the range proposed in de Martino et al. (2014b); de Martino et al. (2015), though these do not translate into a particularly strong mass constraint; a pulsar mass in the range $M_p \sim 1.09M_\odot - 2.0M_\odot$ and a companion mass in the range $M_c \sim 0.2M_\odot - 0.37M_\odot$. However, the tendency of the model towards edge-on inclinations is concerning. The lack of observed X-ray eclipses and the fact that no eclipse is observed in the spectra argue against this edge-on inclination. The preference for a high inclination in the model suggests that the model is attempting to increase the fraction of ellipsoidal modulation relative to the irradiation. The amplitude of the ellipsoidal modulation is also proportional to the filling factor cubed,

$$A_{\rm el} \sim f^3 q \sin^2 i \tag{3.7}$$

while the irradiation amplitude is proportional to the square of the filling factor (Breton et al., 2012). The high inclination may be compensating for a smaller filling factor, which may be caused by the model attempting to fit the asymmetries with a larger ellipsoidal term than is truly present. The source of asymmetry in the system is still uncertain; while the two proposed model extensions do offer

a significantly improved fit the discrepancies between the model parameters and observables, notably the inclination, suggest that there is still significant physics in the system that is not understood.

Before performing the modelling with fixed inclination described in section 3.5.1 we attempted to model the system with the pulsar mass fixed over a range of values, as with J1023. However, the K_2 velocities obtained with these models were unacceptably large; more than three standard deviations above the spectroscopic distances. While we might expect some systematic error in this spectroscopic value, it cannot be large enough to explain this discrepancy. Since in these fixed-mass models the mass ratio and inclination are derived from K_2, the high inclinations that the model consistently prefers necessitate high K_2 velocities. With the inclination fixed, q and $M_{\rm psr}$ are derived so this is no longer an issue and the model is well-behaved.

3.6.4 J1023

While several system parameters - namely the inclination, mass ratio, radial velocity, and filling factor - are broadly similar across the different models, some considerable differences remain. The temperature distributions and pulsar masses were starkly different in each case. Comparing the temperature distributions of the companion surfaces of each model we see that the large, polar spot seen on the hot spot model is not reproduced by the heat redistribution model, despite both models having very similar fit residuals. This could suggest that the large spot may be a sign of overfitting, which would indicate that the magnetic ducting theory proposed in Kandel et al. (2020) is not appropriate here. However, many cataclysmic variables (CVs) and rotating stars also display polar spots (Watson et al., 2007), which may suggest that the similar alignment of magnetic fields is not coincidental. It should be noted that the spots in Watson et al. (2007) are cold spots, as opposed to the hot spots which are favoured in this work. An additional explanation may be drawn from the fact that the viewing angle of this polar spot does not change much over the orbit, suggesting that this spot configuration is in fact fitting for an additional, constant source of flux.

As such, we introduce a third light in the system with an assumed power law index of unity derived empirically from broadband spectra (Hernandez Santisebastien, et al., private communication). This light is incorporated as a fraction, t, of the expected flux from the companion alone such that $t = F_{\rm 3rd\ light}/F_{\rm companion}$. For each band, these correspond to $t_i = 0.150$, $t_g = 0.196$, and $t_u = 0.625$, with the contribution from the third light most significant in the u_s-band. We then adjust the fluxes prior to fitting by multiplying each band by the ratio $1/(1+t)$, such that the

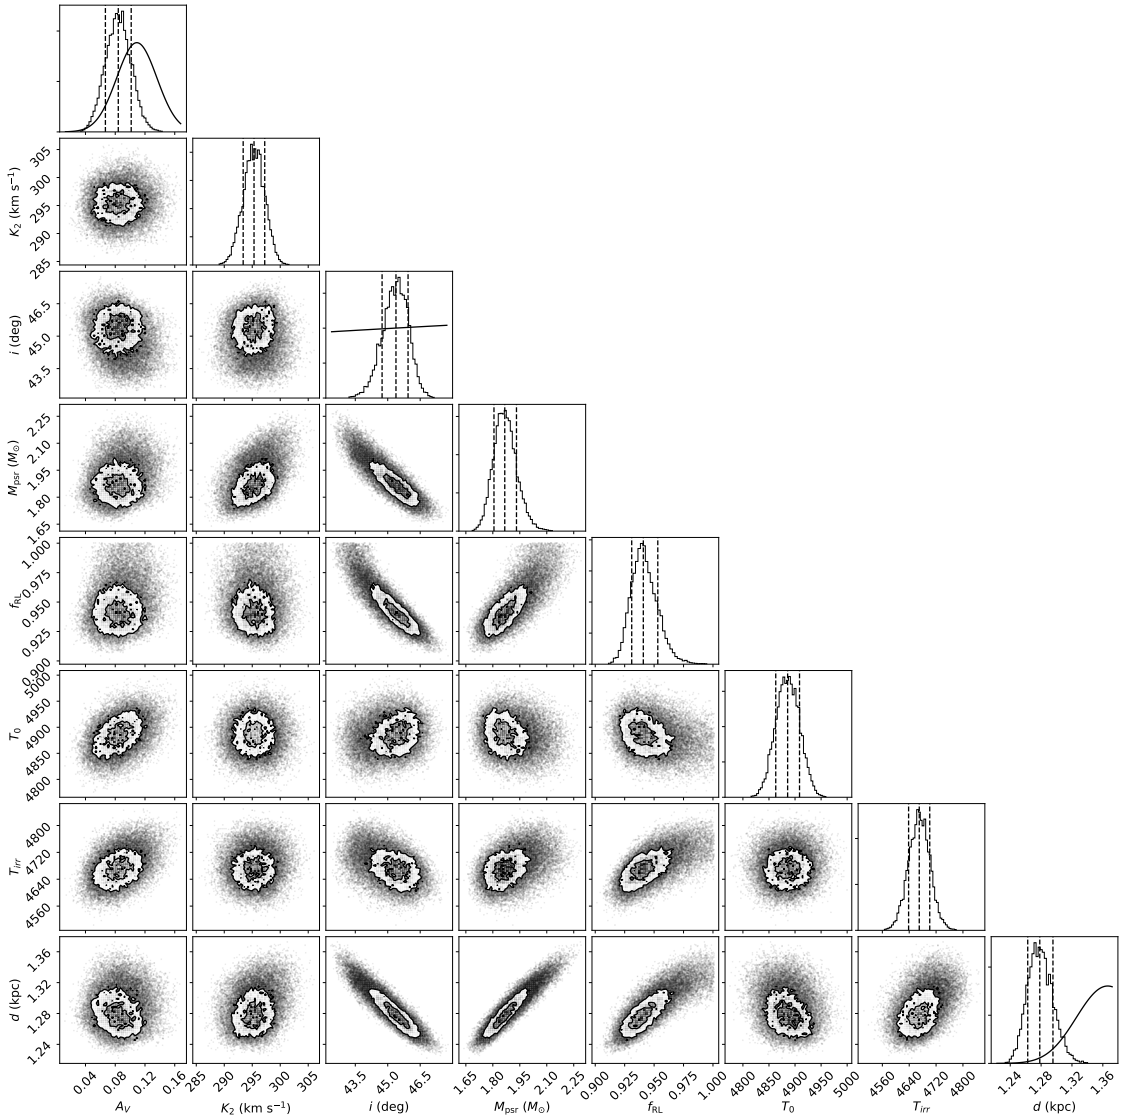

FIGURE 3.13: Corner plot of selected parameters of the J1023 hot spot model. Not shown are the hot spot parameters and the mass ratio. Note the discrepancy between the distance posterior and prior distribution, and the strong covariance of the distance with the inclination, pulsar mass, and filling factor.

modelled contribution of the third light is removed and the remaining flux represents only the light from the companion. We use a heat redistribution model with linear diffusion and constant advection profile. We chose this model over the better-fitting hot spot model as for that model the free temperature of the spot makes it difficult to clearly separate the flux contributions of each source. That is, the spot temperature can easily decrease to compensate for the reduced flux.

While the adjusted i_s and g_s bands were fit well by the model, the adjusted flux of the u_s band could not be matched. This suggests that either our power law model does not accurately describe the third light flux, or that a third light is not able to account for the distance discrepancy. Notably, the shape of the residuals in u' were unchanged compared to a fit using unaltered fluxes. We note that the spectra presented in Shahbaz et al. (2019) show no evidence for a stellar third light. A continuum emission source such as the synchrotron emission produced in the intrabinary shock (Romani & Sanchez, 2016) may be an alternative source of the flux, however this emission follows a negative power law. Shahbaz et al. (in prep. and private communication) further show no evidence for a third light, with the secondary star the sole source of flux from 6000 Å. However, a metal-rich secondary is observed, with an iron excess of Fe/H = 0.48 (Shahbaz et al. (2019) and private communication). This results in a u'-band flux of 82% the solar equivalent in our model, which may account for the closer distance that we obtain due to excess flux. A system distance of between 1.25 and 1.30 kpc, as we obtain, compared to the interferometry distance of 1.368 kpc corresponds with a decrease in magnitude of between 0.11 and 0.2 mag in the u_s band. This is equivalent to a decrease in flux of $83 - 90\%$, which is comparable to the expected decrease due to a higher metallicity. The atmosphere grids used in this modelling do not account for this high iron excess, and as such this avenue may help to explain the distance we obtain.

While the best fitting values are different for the heat redistribution and hot spot models, both show a strong covariance between the filling factor, f, and the distance. This can be seen in the corner plot shown in figure 3.13 This suggests that a better fit to the distance will bring the filling factor closer to unity, as we would initially expect.

Considering the competing prior values for the A_V extinction, our modelling shows that both values are almost equally favoured as no other parameters are affected by the change of prior within uncertainties. As such, we can surmise that the small difference in model evidence (-2297.1 for the value from N_H and -2293.9 for the A_V from dust maps) are due only to the differences in penalty from the A_V priors. The A_V value from the dust maps, which provides the marginally improved

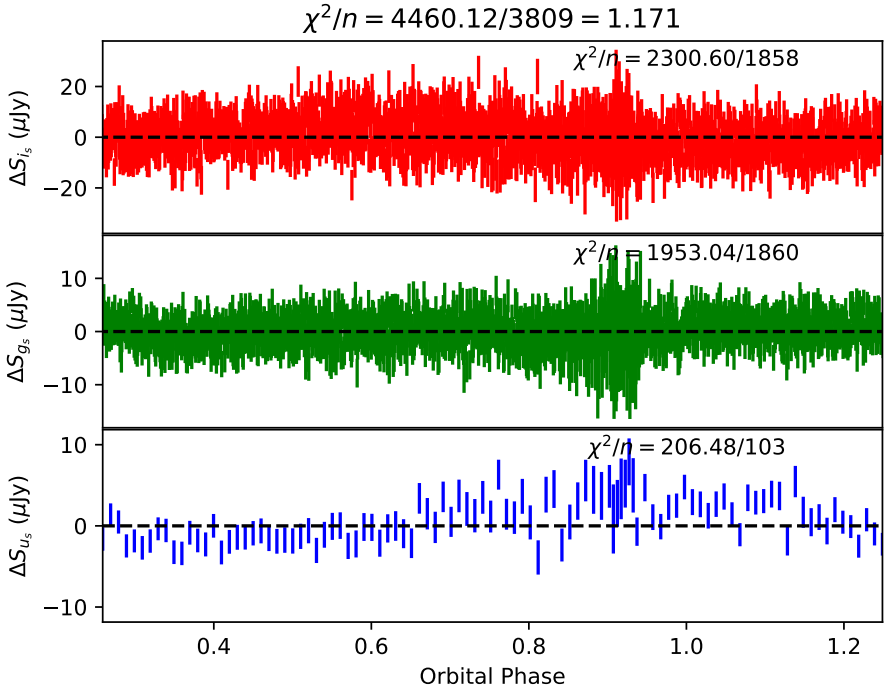

FIGURE 3.14: Residuals per band for the best-fitting hot spot model of J1023; note that the y-axis is in units of flux. The u' residuals show a clear systematic pattern that mirrors the sinusoidal shape of the asymmetry in figure 3.3. Note also that the reduced χ^2 value for the u' is significantly larger than that for the other two channels; 1.88 compared to 1.05 for g' and 1.31 for i'.

model evidence, was used for the modelling in this work, though in practice this investigation shows that the model parameter posteriors are nearly ideally distributed and any changes to the A_V have little effect on the results.

3.7 Conclusions

We present new, high time resolution optical photometry of the tMSPs PSR J1023+0038 and PSR J1227–4853, and discuss our numerical modelling of their light curves. Using a new extension to the `Icarus` code including the thermal contributions of a hot or cold spot on the companion surface, we modelled the asymmetric light curves of these sources and obtained significantly improved fits over the symmetric case. Using this model we constrained several key parameters of the systems, including the companion Roche lobe filling factor, temperature profile, and system distance. We also performed modelling using a further extension which considers the diffusion and convection on the companion star surface. This model also provided an improved fit compared to the symmetric model, though for both systems the evidence favours a model with a hot spot.

We found that the filling factor of both sources was less than 1.0, indicating that the Roche lobe is under-filled. This is at odds with other results showing a full Roche lobe in the system's RP state (de Martino et al., 2014b). We expect the companion in the AP state to fill its Roche lobe, so these results indicate that the filling factor plays an important role in the tMSP transition. However, when considering the volume-averaged filling factor, we find that the companion stars are only slightly underfilling their Roche lobes. Our results suggest that the companion stars may undergo an expansion and contraction between the AP and RP states of tMSP cycle, as there is no sign of a change in the size of the Roche lobe; changes in the orbital period are not large enough to account for the under full Roche Lobes. Taking our filling factors at face value, they indicate that between state transitions the stellar radius changes by an order of 5%, something which may be possible if the companion has a core-envelope structure with only a small fraction of its mass in the envelope.

Some of the limitations and errors in our results indicate that our model does not completely describe the asymmetry of the light curves, suggesting that improvements to the model are needed, or indeed a new, physically motivated approach. For example, significant systematics remain in the residuals of the u_s band of J1023. Using a further extension to the `Icarus` code which models diffusion and convection in the companion photosphere we performed a second round of modelling. This again provided an improved fit over the symmetric model, however the inconsistent results, such as the inclination of J1227, persisted, and indeed the fit overall was poorer than the hot spot case. However, modelling with some combination of the two extensions may be a good starting point or future investigations.

Even with these excellent data, these sources are not ideally modelled. Despite this, we believe the integrity of the key findings are intact and recommend additional study of these two systems. In particular for J1227, additional spectroscopy and radio interferometry would significantly improve the constraints on the distance and inclination.

Chapter 4

Periodicity search and classification pipeline for GOTO

4.1 Introduction

4.1.1 Background

The importance of wide-field surveys is clear in modern astronomy, from γ-ray source discovery using Fermi-LAT (Ray et al., 2012) to the measurement of parallaxes using GAIA (Gaia Collaboration et al., 2016). While the majority of these surveys use large, dedicated telescopes, such as the 2.5 m reflector at Apache Point Observatory used for the Sloan Digital Sky Survey (SDSS) (Gunn et al., 2006), a number of new optical surveys are emerging using 'off-the-shelf', consumer telescopes and instruments. These have the immediate advantage of being significantly cheaper than their specialist counterparts, and while they tend to be less sensitive this is often made up for by their scalability. This scalability is well demonstrated with the All Sky Automated Survey for Supernovae (ASAS-SN). Starting from a single 14 cm telescope, ASAS-SN now uses a worldwide network of 24 telescopes for constant all-sky coverage (Shappee et al., 2014). While this survey was conceived to search for new supernovae, a result of the near-daily observations is an archive of photometry of over 60,000 variable sources down to a V-band magnitude of $V \sim 17$ (Jayasinghe et al., 2018); this source list makes up the ASAS-SN catalogue of variable stars. Using periodicity search algorithms and machine learning, the survey has identified and classified a large catalogue of over 27,000 periodic variables. The Random Forest classifier used in this work achieved an overall accuracy of over 99% (Jayasinghe et al., 2019).

Archives like this are incredibly useful in the search and study of variable optical sources. We aim to achieve a similar goal with data from the Gravitational-wave

Optical Transient Observatory (GOTO), a fledgling wide-field survey designed to search for optical counterparts to gravitational wave detection events (Dyer et al., 2020). Notably, the larger telescopes used with GOTO compared to ASAS-SN result in a deeper survey due to the increased sensitivity.

Machine learning has recently been implemented for transient detection in the difference imaging produced by GOTO with over 95% accuracy (Mong et al., 2020; Killestein et al., 2021), however this is limited to real-bogus classification of images and does not extend to the classification of light curves. Burhanudin et al. (in prep.) introduce a classifier designed to distinguish variable stars from supernovae and active galactic nuclei, achieving an accuracy between 80% and 91%. We introduce a similar classifier designed to differentiate between the sub-classes of variable stars, and discuss its performance with the currently-available GOTO data.

Spider pulsars are notable for their distinctive optical and gamma-ray signatures, with the optical light curves exhibiting both ellipsoidal modulation and irradiation effects. The gamma ray signature has driven the bulk of discoveries of new spiders, with all but ~ 5 of the ~ 50 known spiders discovered in the last 10 years from FERMI gamma ray sources. However, there remain approximately 1300 FERMI sources with no optical association - the majority of these are AGNs, however many will be spider candidates. Of the 5064 fourth Fermi Large Area Telescope catalog (4FGL), 3130 are active galaxies while 239 are pulsars (Abdollahi et al., 2020). Therefore out of the ~ 1300 unassociated sources we may expect to find approximately 60 new pulsars. Surveys like GOTO are well suited to finding these optical counterparts. Given a Fermi error ellipse, sources observed by the telescope in this patch of sky can be searched for periodicity and classified. If any suitable candidates are identified, then deeper follow-up observations can be made to confirm the association. However, the optical magnitudes of many spiders, especially the fainter black widows, will be beyond the sensitivity of most smaller telescopes including GOTO. Despite this, we expect that many objects with magnitudes comparable with the brightest known spiders will be detectable, though how many of these remain to be discovered is unknown. In this section we discuss our implementation of a periodicity search and classification pipeline to synthetic and real photometry from the GOTO instrument, as well as results from applying these techniques to sources in the ASAS-SN variable catalogue.

4.1.2 The GOTO instrument

GOTO is a ground-based robotic telescope located at the Roque de los Muchachos Observatory observatory on La Palma (Dyer et al., 2020). Comprised of eight 0.4 m

telescopes on a single mount, it has a field of view of 40 square degrees. There are plans to extend the configuration with a further three telescope nodes; a second in La Palma and two at the Siding Spring Observatory in Australia. Each individual telescope, referred to as a Unit Telescope (UT), has its own filter wheel and CCD (charge coupled device) camera. There is capacity for R, G, and B Baader filters, though at the time of writing all observations have been performed using the wide-band (375 - 700 nm) L Baader filter (Gompertz et al., 2020). With this configuration, the telescope is sensitive down to a magnitude of 20.5 mag in 3 stacked 60 second exposures under ideal dark conditions. Each UT has a plate scale of 1.25"/pixel corresponding an individual field of view of $\sim$6.9 square degrees.

Images from GOTO are processed using the 'GOTOphoto' data reduction pipeline at a data centre at Warwick University, following the procedure outlined in Dyer et al. (2020). After initial calibration using master dark, flat, and bias frames, source detection with SExtractor (Bertin & Arnouts, 1996) and astrometry with astrometry.net (Lang et al., 2010) are performed on each image. Following this, detected sources are matched to the ATLAS-REFCAT2 catalogue (Tonry et al., 2018). The 3 stacked exposures are first processed individually, and then combined and reprocessed with the same procedure. This enables the detection of fainter sources, as well as the ability to track objects (such as asteroids) which move between the individual images. Finally, difference images are produced by subtracting the combined image from the corresponding master reference frame using HOTPANTS, enabling the detection of transient sources (Becker, 2015). In addition to this, the pipeline also performs so-called forced photometry, where non-detections of a variable sources are assigned a lower-bound magnitude. This method of photometry does not provide calibration as precise as with the reduction of images from dedicated photometers like ULTRACAM.

The primary goal of the GOTO project is to hunt from the optical counterparts to gravitational wave (GW) merger events, such as black hole or neutron star mergers. However, these events are observed relatively infrequently so the remainder of the observing time is used for sky surveys. In fact, nearly all of GOTO's on-sky time is spent in survey mode as of 2021. Survey telescopes like this are a key source of data for variable star astronomy, and not just as a means of discovering new sources. Long-term photometry of the whole sky can provide well-calibrated magnitudes for stars used in photometric reduction, like our use of the SDSS catalogue in calibrating ULTRACAM photometry, or for producing all-sky maps of interstellar reddening (Schlafly & Finkbeiner, 2011). Of course, these surveys also provide key photometry for the identification and classification of sources detected in other wavelengths; photometry from the Catalina Sky Survey (CSS) was used to identify and measure

the orbital period of the optical counterpart to the redback system PSR J1306–40 (Linares, 2018). It is this last case that best aligns with our aim to identify the optical counterparts to unassociated FERMI sources.

However, the volume of data produced by survey telescopes is immense; for GOTO each frame contains of order 10,000 individual sources over 40 square degrees. Given that GOTO surveys the entire visible Northern sky every $\sim$ 14 days, we can expect to observe at least $\sim$ 10 million individual sources (Gompertz et al., 2020). Each source will have a light curve of $\sim \ 10 - 100$ observations, though the number of samples will increase with time. The need for automated detection and classification pipelines is clear; the rate of arrival of data is simply too great to perform the required operations manually. As such, we turn to techniques from the machine learning ecosystem to accomplish the necessary classification, as well as incorporating existing technologies into the pipeline.

4.1.3 Summary of the project

We created and tested a periodicity search and light curve classification pipeline for photometry from the GOTO instrument. This pipeline was developed based on the work detailed in chapter 2, whereby we adapted the ULTRACAM simulation code to the specification of GOTO. Initially, we used synthetic light curves to quantify how well the periodicity and modulation amplitude of spider-like light curves, as sampled by GOTO, could be recovered. This process is detailed in section 2.4.2. Once this simulation code was established, the synthetic light curves were also used to test a periodicity search pipeline. This pipeline, detailed in section 4.3, is the main focus of this chapter. This section introduces the constituent parts of the pipeline, including data acquisition, reduction and cleaning, period search algorithm, and finally the ML classifier. Lastly, we discuss some of the key results obtained during the development of the pipeline, and their consequences for the future of this project.

4.1.4 Introduction to star types

Table 4.1 introduces some of the abbreviations used to describe the star classes used in the classifier and other parts of this project, while figure 4.1 shows a typical ASAS-SN light curve for each. Spiders typically fall into the EW (W Urase Majoris-type binary) class.

Semi-regular variable stars (SR) are late spectral type giant or supergiant stars which display moderate periodicity in the variability of their brightness, with periods longer than 30 days (Chinarova & Andronov, 2000). These stars often

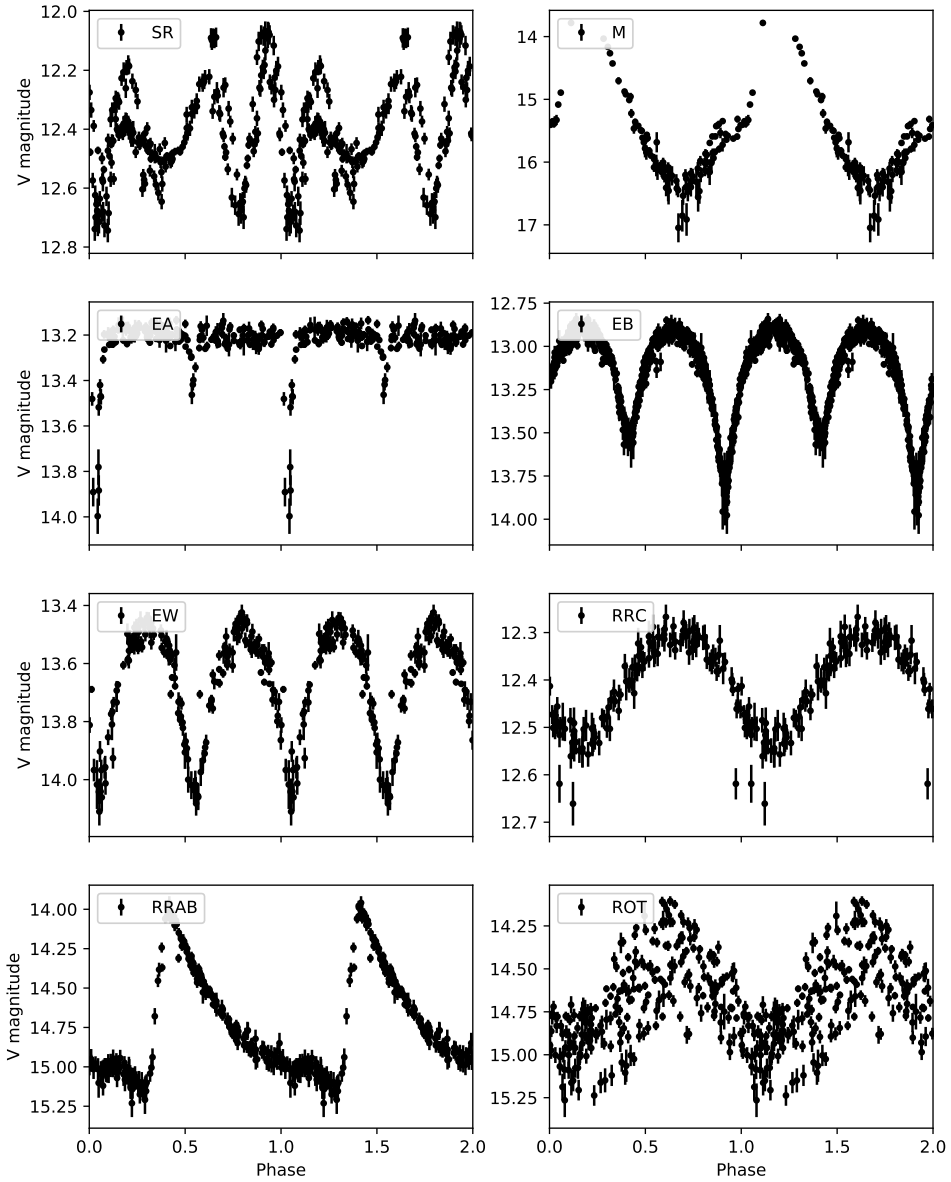

FIGURE 4.1: Example light curve for each of the 8 star types considered in this work. Each light curve shows the ASAS-SN V-band magnitude. Variation within each class can be substantial, however the light curves shown here were chosen as they display the characteristic features of each class.

TABLE 4.1: ASAS-SN classes and descriptions, as described in Jayasinghe et al. (2018).

Abbreviation	Star type
SR	Semi-regular variables with $P > 30$ d
M	Mira variables
EA	Detached Algol-type binaries
EB	β Lyrae-type binaries
EW	W Ursae Majoris-type binaries
RRC	RR Lyrae variables (first overtone)
RRAB	Asymmetric RR Lyrae variables
ROT	Spotted variables with rotation modulation

display multiple or variable periods, generally with an amplitude of modulation of less than 2.5 mag in V-band (Lebzelter et al., 1995). Where there is no regular period, the mean is often taken.

Mira Variables (M) are long ($P \gtrsim 100$ d) variable stars which experience thermal pulsations, named after the type star Mira (Bedding & Zijlstra, 1998). These stars are typically very red, owing to their location at the top of the asymptotic giant branch, and are generally between 1 $M_\odot$ and 2 $M_\odot$ (Ireland et al., 2004). Mira variables exhibit a colour-dependent period-luminosity relationship (Feast et al., 1989).

Detached Algol binaries (EA) are a subclass of the Algol-type binary, containing an early-type, main-sequence primary and a less massive, RL-filling secondary above the main sequence. The primary does not fill its Roche lobe, and is hotter, brighter, and smaller than the secondary (Chen et al., 2006). These binaries generally have an orbital period of a few days, generally falling in the range of 0.5 d $\lesssim P \lesssim$ 30 d, with primary masses 0.1 $M_\odot \lesssim M_1 \lesssim$ 20 $M_\odot$ and secondary masses 0.1 $M_\odot \lesssim M_2 \lesssim$ 5 $M_\odot$ (Giuricin & Mardirossian, 1981). Compared to typical Algol-type binaries, detached Algol-type binaries do not share a common envelope. The light curves typically have one optical minima from eclipse, which is well-defined.

***β* Lyrae binaries (EB)** are a class of eclipsing variable somewhat similar to the EA class, however they exhibit mass transfer between the two stars (Harmanec & Scholz, 1993). As such an accretion disc is present. The masses of the component stars are also generally more massive than with EA binaries, and the optical minima at eclipses are smoother and are typically unequal depths (Jayasinghe et al., 2018).

W Ursae Majoris binaries (EW) are the final class of eclipsing binary considered in this work. The properties listed here are summarised from Terrell et al. (2012). EW binaries have very tight orbits and as such share a convective common envelope, resulting in mass exchange, and are typically of the cooler F, G, and K spectral types. Due to their tight orbits, orbital periods are typically of the order $0.2 \text{ d} \lesssim P \lesssim 0.5 \text{ d}$, and there is a period-colour relation which suggests that these systems evolve towards longer periods and redder colours. The light curves of these systems have smooth eclipses of similar depths.

RR Lyrae variables (RRC and RRAB) are pulsating giant stars (mass $\sim 0.7\ M_\odot$) with periods in the range of 0.2-1.2 d and A-F spectral types (Dambis et al., 2013). They sit on the horizontal branch and exhibit helium core burning, showing a strong period-metallicity-luminosity relationship. In addition, the age ($\gtrsim 10$ Gyr) of these stars means they typically are found in older galaxies and globular clusters and so are commonly used as standard candles for distance estimates. In this work we consider two sub-classes, the asymmetric (RRAB) and first overtone (RRC) RR Lyrae. The first displays the characteristic asymmetric light curve with a steep rising slope and longer tail, while the latter displays more sinusoidal modulation with typically shorter orbital periods (Smith, 2004; Bingham et al., 1984).

Spotted variables with rotation modulation (ROT) are a general class of variable which display rotational modulation, including the RS Canum Venaticorum- and BY Draconic-type systems (Jayasinghe et al., 2018). The latter of these, for example, exhibits star spots which result in luminosity modulation at the star's spin period. As star spots are not constant features, the periods and modulation amplitudes of these stars are not constant over long timescales (Hall & Henry, 1992). Additionally, these stars may appear in binaries and so display additional modulation, and even exhibit stellar flares (Boden & Lane, 2001).

4.2 Data Acquisition

4.2.1 Simulation of GOTO data

To test the functionality of the periodicity search algorithms, we simulated test fields of sources as they would be observed with GOTO. This simulation is based on the simulation algorithm described in section 2.4.2. The GOTO field is comprised of some fraction of periodic sources, with the remainder sources of constant magnitude. Initially, we used a template for the periodic sources based on our database of known redback and black widow pulsar systems. This includes the range of periods

$P \sim 1.5 - 27$ hr, mean R-band magnitudes $R \sim 17 - 22$, fractional amplitudes of modulation $k \sim 0.01 - 0.2$, and colour information which is used to define the magnitude of other bands from a given R-band magnitude.

Of course, this template can be relaxed in order to simulate a more general population of sources, for example with a much larger range of periods. We considered the magnitudes in terms of the corresponding signal-to-noise ratio (SNR) when observed with GOTO. These are calculated using a modified version of the ULTRACAM sensitivity calculator, calibrated to the technical specifications of GOTO. We chose to use SNRs over magnitudes or fluxes as this also takes into account the observing conditions; for example a mag 17 source will be observed with a greater SNR on a dark night compared to one with a full moon. Importantly, we also simulate the observing schedule of GOTO. This includes defining the number of nights in a given time frame on which observations take place, as well as the observation schedule on a given night and other factors such as the average cadence of the observations. This last factor is crucial as aliasing can be a large source of error in determining periodicity using our method. As such, we can compare the quality of the recovery of periods from a strongly aliased case with observations exactly 24 hours apart to observations performed at different times in subsequent nights. Aliasing occurs when periodicity in the observation times is convolved with periodicity in the signal, resulting in aliased peaks in a periodogram of the data. The results of this analysis are discussed in section 4.3.3. The simulation parameters are summarised in table 4.2.

With the system parameters and observation schedule defined, we use a simple model to generate the light curves. Comprised of sinusoidal modulation plus some offset, it is of the form $f(t) = -A(k\cos(2\pi t P^{-1} + \phi) + 1)$, where A is the mean SNR, k is the fraction amplitude of modulation, P is the orbital period, and t is the time since the first observation. Throughout this work we use P and t in units of days. $\phi \in [0, 1]$ is a constant phase offset, chosen from a uniform distribution and added to each of the sources. When creating the dataset, we set a fraction, $f_{\rm var}$, of the sources to be variable, with the rest constant. We shall label datasets produced in this way as Type A, with sinusoidal light curves over the period range 0.1 d $< P <$ 2 d.

We also use the harmonic decomposition described in section 4.4.1 to create non-sinusoidal synthetic light curves. We first decompose light curves from the ASAS-SN variable star catalogue into their Fourier components. Then, by using a similar method to the Type A datasets, we synthesise new light curves from these components using the observing schedule and sensitivity of GOTO. These will be labeled as Type B datasets, with an example shown in figure 4.2. The periods of the

TABLE 4.2: Key GOTO light curve simulation parameters and descriptions.

Parameter	Default	Description
`nsources`	1000	Number of sources to simulate.
`bands`	3	Number of photometric bands to simulate.
`filters`	`'RGB'`	Which bands to simulate. If `bands`= 2, can set as `'RB'` or `'RG'` for example.
`samples`	60	Number of observations per light curve.
`days`	365	Number of days over which the observations are spaced.
`minm`	18	Brightest magnitude of the brightest band.
`maxm`	24	Faintest magnitude of the brightest band.
`meanm`	21	Mean magnitude of the brightest band
`mina`	0.01	Minimum modulation amplitude (fractional).
`maxa`	0.5	Maximum modulation amplitude (fractional).
`jitter`	0.2	Variance of random scatter to add to observation times.
`exp`	180	Exposure time in seconds.
`cadence`	360	Minimum time between exposures in seconds.
`varfrac`	0.5	Fraction of sources to be variable.
`declim`	21.5	Magnitude corresponding to detection limit of simulated telescope.
`offsets`	`True`	Include phase offsets; if set to `False` then all synthetic light curves will be in phase.

sources in this dataset can either be set as the same as the original ASAS-SN sources, or customised when resampling. These data are used to populate a GOTO FoV with sources based on those in the ASAS-SN catalogue to test both the periodicity search algorithm and the machine learning classifier.

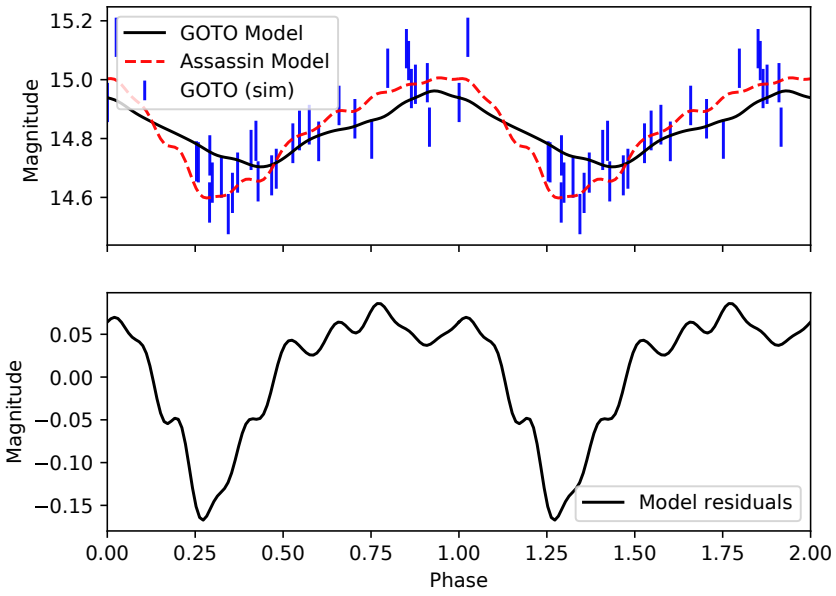

FIGURE 4.2: *Top:* The blue data points are a synthetic light curve drawn from the Fourier decomposition of an ASAS-SN light curve, shown by a red dashed line. The black line shows the Fourier decomposition of the GOTO light curve. *Bottom:* Residuals between the two Fourier decompositions in the top panel. These show a reconstruction to an acceptable degree, with the majority of reconstructed light curves showing residuals of order ~ 0.1 mag.

4.2.2 Obtaining and reducing photometry

An archive of past GOTO photometry is offered through a PostgreSQL database hosted by Warwick University. For this project, we use the Phase 4 dataset, which contains observations from MJD 58536 (2019-02-22) to present. The database is accessed using PSQL commands executed using the `psycopg2` Python package. We implemented two methods of querying the database. The first is optimised for the intended use of the pipeline; for a list of positions and error radii, such as those obtained by FERMI localisations, the database is queried for stars in each ellipse and the data obtained is stored separately for each FERMI ellipse. The second is used for obtaining data on a large number of sources for use with the classifier; this method makes multiple small queries and caches the results which are then saved together as one large dataset. During this step in the pipeline, the GAIA catalogue is also queried using the same position and radius (Gaia Collaboration et al., 2018). We obtain the following observations from the database: the (RA, Dec) position, the G-band magnitude, and the $B_P - R_P$ colour term. Note that the GOTO light curves consist of L-band, which cannot be directly compared to the narrower G-band of GAIA, for instance. The relationship between the two is not well defined

(in the same way that the SDSS filters are, for example), however it is approximately directly proportional. Figure 4.3 illustrates that this relationship does not hold for a significant number of sources, indicating that for situations where a precise source magnitude is required, the GAIA G-band should be used. This discrepancy arises from two main causes. First, the wide-band L filter does not capture colour the same way the narrower-band GAIA G filter does; if a source is significantly brighter in the red or blue wavelengths, it will appear brighter in the L-band than the G band. This effect will be compounded for variable stars which exhibit colour changes. Secondly, since the magnitude of variable sources is not constant, these independently-sampled magnitude measurements will contribute significantly to the discrepancy. We note that figure 4.3 shows a clear bias, such that for a source of a given GAIA G-band magnitude the GOTO L-band magnitude is in most cases fainter.

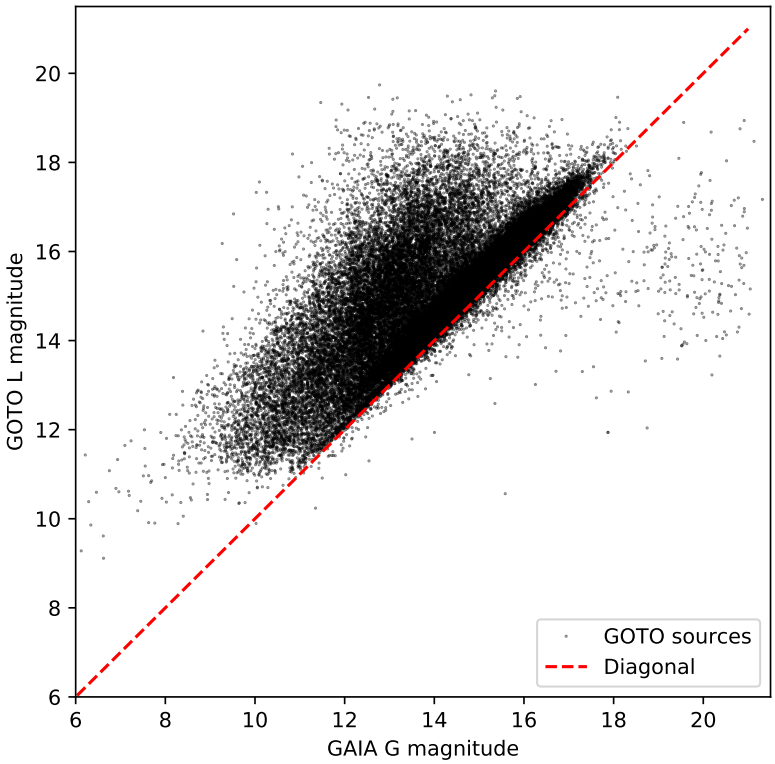

FIGURE 4.3: Comparison of GAIA G-band and GOTO L-band magnitudes of 40,000 variables sources. The red dashed line does not represent a fitted relationship, rather illustrating the diagonal for comparison. Note that most of the sources are clustered along this line, with a broad cloud of points spreading away from it.

The data obtained from the GOTO database are not grouped by source; instead, individual observations are separate. However, since the positions of each source are well-localised, a unique ID for each source is available. We therefore group the observations by source ID to produce light curves. Next, any light curves with less

than 10 observations are removed. While the GOTO pipeline performs a robust reduction and photometry, during testing of the methods in this research we noted the frequent occurrence of unacceptable data points, such as single observations many sigma from the mean of non-varying sources. We found no correlation between the deviation from the mean of these outliers and the FWHM of the source PSF in the corresponding images, suggesting that the photometry is not significantly affected by changes in seeing.

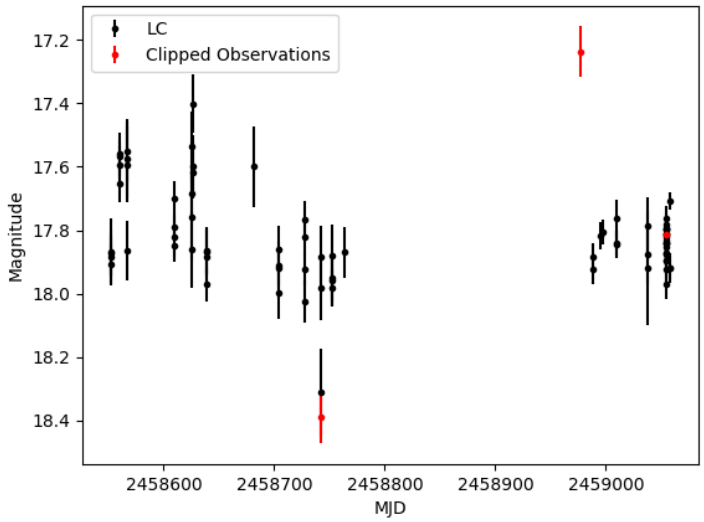

FIGURE 4.4: Light curve with observations flagged by the ensemble clipping shown in red. Note that one clipped observation appears to not be an outlier; this point would have been flagged as an outlier in a significant number of other light curves.

A fraction are likely to be due to cosmic rays not detected in the pipeline as on-source cosmic ray detection is generally less effective than off-source detection. While these extreme outliers can be removed by sigma clipping, other erroneous data points remained. These erroneous points were first noticed in the light curves of sources with little or no intrinsic variability; for example an otherwise flat light curve with some expected scatter would have a number of observations several standard deviations from the mean. We therefore implemented a form of sigma-clipping using the ensemble of sources in a given field. This method identifies frames where a substantial fraction of the observations within deviate significantly from the expected magnitudes. The frames are differentiated by their MJD timestamps. For a field of sources in some FERMI error ellipse, we calculate the means and standard deviation of each source with more than 12 observations. Then, since some sources may reside in the area where two or more UT fields of view overlap, we separate the observations in each light curve by UT number. For example, a given light curve of 20 observations

may have 10 in UT1 and 10 in UT2: the observations comprising one single light curve may be split over several UTs. The deviation from the mean is then calculated for each observation within each UT. The fraction of sources with deviations greater than 3-sigma in each unique observation is calculated. If this fraction is greater than a threshold value then the corresponding UT frame is discarded, such that a frame with more than 10% of the observations lying more than 3-sigma from the mean will be discarded. During this investigation we used a threshold of 0.1, determined through visual inspection. Note that this threshold may not be the optimum value in all cases, however it is an acceptable middle ground which can be adjusted during data acquisition if necessary. A threshold much smaller than this risks incorrectly eliminating a large number of frames, while for a larger threshold frames containing a significant number of outliers were rarely discarded. Figure 4.4 shows a light curve after this stage of clipping.

In addition to the ensemble clipping process, we also perform a standard sigma-clip using the `astropy sigma_clip` package in order to remove outliers which are not correlated within GOTO frames, such as cosmic rays. We use an asymmetric set of thresholds, with 6-sigma and 3-sigma for fainter and brighter points respectively. This way we are sensitive to cosmic rays, but less likely to erroneously remove data points.

4.3 Periodicity Searching

Given a field of GOTO light curves, the first step towards classification is to identify periodic sources. This section describes the process we use and discusses the performance of the algorithms. Before the periodicity search proper, we eliminate non-variable sources from the field using the method described in section 4.3.1. After this, the remaining sources are tested for periodicity. The choice of periodicity search algorithm is discussed in section 4.3.2, and their performance in section 4.3.3. We make use of the synthetic datasets introduced in the previous section for testing.

4.3.1 First-pass elimination of non-varying sources

To dramatically reduce computing time, a first-pass analysis of sources in the catalogue is performed. This is a simple fit of the source's light curve to a linear model centred on the source's mean magnitude, $\bar{m}$. A non-varying source is expected to have measurements consistent with $\bar{m}$, within measurement uncertainty, while we expect a variable source to be poorly fit by this linear model. The reduced chi-square statistic, χ^2_{red} is calculated and compared to a cutoff value, if it exceeds this then

the source is flagged as a variable source. This is equivalent to a goodness-of-fit test, with the null hypothesis being a linear model with slope 0 and intercept equal to $\bar{m}$. Using a Type A dataset containing 1000 sources with a variability fraction of $f_{\text{var}} = 0.5$ (such that 500 sources were variable and 500 were constant), we performed the fit with a range of χ^2_{cut} cutoff values, spaced logarithmically from 0.5 to 10. The default values in table 4.2 were used. At each χ^2_{cut} value the accuracy (fraction of sources identified correctly), number of false positives, and number of false negatives were recorded and can be seen in figure 4.5.

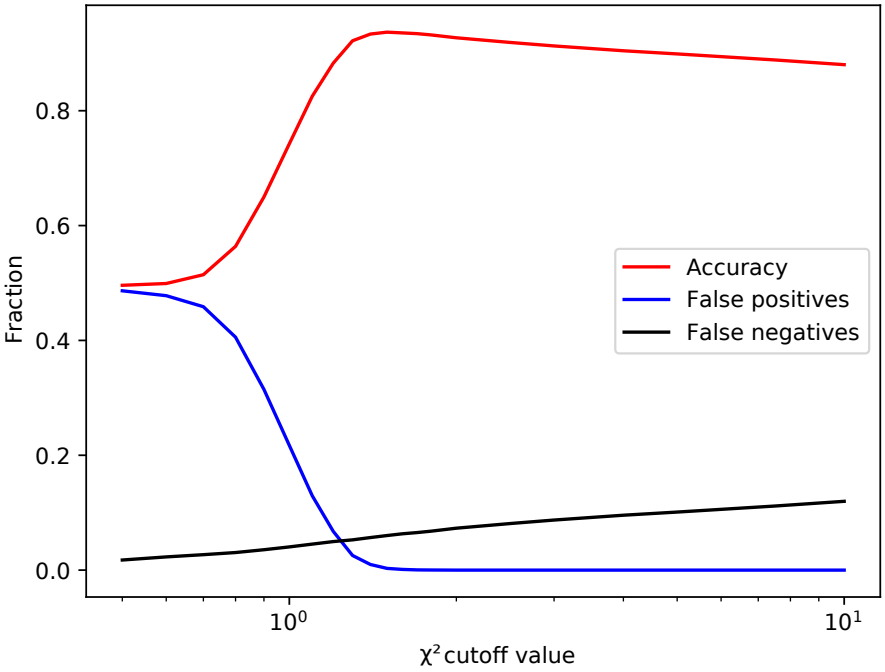

FIGURE 4.5: Performance of the first pass elimination with varying χ^2 cutoff value. With low values, the accuracy converges to the fraction of variable sources in the testing data ($f_{\text{var}} = 0.5$) as almost all light curves are classified as variable. With larger χ^2_{cut} values, the false positive rate quickly drops off while the false negative rate slowly increases.

At the lowest values of χ^2_{cut}, practically all sources are classified as variables as next to no scatter is allowed for. As the cutoff value increases, the number of false positives drops off rapidly to nearly zero while the false negative rate gradually increases. The optimum value for this dataset, with an accuracy (fraction of sources correctly classified) of 0.92, is $\chi^2_{\text{cut}} = 1.5$. At this value the false negative rate is approximately 7%. It is worth noting that this method is not sensitive to variable sources with light curves of low signal-to-noise ratios or small variability amplitudes. In both these cases the amplitude of variability is comparable to the Gaussian scatter of observations and so would likely be a false negative. However, discarding these

sources is sensible as it is unlikely any periodicity would be uncovered due to this fact, and so classification of the source would not be successful. In fact, running our periodicity search on the pool of false negatives from this test recovered the correct period in only 17% of the sources. This first-pass elimination makes the survey less sensitive to both dim and weakly variable sources, but as a result the classification of variable sources will be more robust, and considerably faster as there are fewer light curves to process. Figure 4.6 shows one such light curve, illustrating how the weak variable signal is lost in the noise.

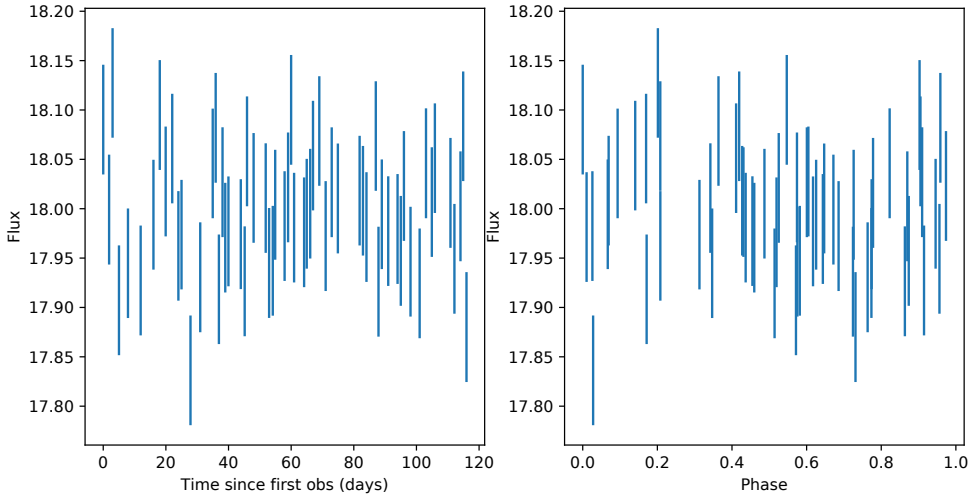

FIGURE 4.6: Left: unphased GOTO light curve, with flux in arbitrary units. Right: light curve phased at period. The sinusoidal modulation is just visible, though the scatter is of similar amplitude

We also note that the threshold determined using this method will not necessarily extend to light curves with a different number of observations (i.e. degrees of freedom) due to the nature of the χ^2 distribution's dependence on the number of degrees of freedom. While we have demonstrated the feasibility of this method for identifying variable sources, the relationship between the number of degrees of freedom and χ^2_{thresh} should be determined before deployment of this code.

4.3.2 Identifying periodic sources

Algorithm choices

The variable sources flagged with the first-pass analysis are then run through a periodicity search algorithm. There are a plethora of methods for period searching, such as those based upon discrete Fourier transforms (Deeming, 1975) or least squares approximations (notably the Lomb-Scargle (LS) periodogram; Lomb (1976)

and Scargle (1982)) and more recent methods such as those based on the analysis of variance (AoV) (Schwarzenberg-Czerny, 1996a). A thorough review and comparison of many different methods can be found in Graham et al. (2013b), in which methods using analysis of variance or conditional entropy (Graham et al., 2013a) were found to be the most effective. Conditional entropy, while effective does not natively account for uncertainties, though this can somewhat be mitigated by using more computationally expensive techniques. Since the majority of spider systems have a magnitude close to or beyond the sensitivity limit of GOTO, we expect the data of these sources will frequently have considerable uncertainties and therefore this method will not be sufficient. However, analysis of variance methods do explicitly take into account errors and perform similarly well. Therefore we proceed using the AoV method described in Mondrik et al. (2015), which describes a multi-band implementation of the multi-harmonic analysis of variance (MHAOV) method. Cutting edge methods, such as the template periodogram of Hoffman et al. (2021) show new improvements in the method, however this particular method relies on a pre-defined light curve template.

Additionally, we use a multi-band, multi-harmonic, extension of the generalised Lomb-Scargle (GLS) periodogram (VanderPlas & Ivezić, 2015) as a comparison to the MHAOV method. We note that while both of the methods implemented here are multi-band extensions, our investigations only consider the single band case as the images from GOTO are only available in single band at the time of writing.

Generalised Lomb-Scargle periodogram The classical periodogram,

$$P_X(\omega) = \frac{1}{N}\left|\sum_{k=1}^{N} y_k e^{i\omega t_k}\right|^2 , \tag{4.1}$$

is a function of angular frequency and is defined for some time series y_k observed at evenly-spaced times $t_k = t_0 + k\Delta t$ for $k = 1, ..., N$ (Schuster, 1898). This periodogram is defined such that if the time series contains a periodic component of frequency ω_0, the factors y_k and $e^{i\omega t_k}$ will be in phase as ω approaches ω_0 and so contribute significantly to the function at that frequency. This corresponds to the appearance of a peak in the frequency spectrum. However, this formulation is only valid for equally spaced observations, which rarely occur in practice, and is also exceptionally prone to noise.

Introduced by Lomb (1976) and Scargle (1982), the so-called normalised periodogram, $P_N(\omega)$, is defined for angular frequency ω and data y_k observed at times

t_k as

$$P_N(\omega) = \frac{1}{2V_y}\left[\frac{\left(\sum_k (y_k - \bar{y})\cos\omega\,(t_k - \tau)\right)^2}{\sum_k \cos^2\omega\,(t_k - \tau)} + \frac{\left(\sum_k (y_k - \bar{y})\sin\omega\,(t_k - \tau)\right)^2}{\sum_k \sin^2\omega\,(t_k - \tau)}\right], \tag{4.2}$$

where $\bar{y}$ is the mean of the data, V_y is the variance, and τ is a time offset which serves two purposes; orthogonalising the model and allowing for a periodic signal of arbitrary phase. Note that this reduces to the classical periodogram for evenly spaced data. P_N in this formulation is equivalent to harmonic least squares analysis; the frequency which maximises P_N is the same which minimises the least-squares statistic when fitting a sine wave to the data (Scargle, 1982). As such, equation 4.2 can be derived from the χ^2 statistic of a single-term sinusoidal model.

We use the multi-band periodogram described in VanderPlas (2018), in which a generalisation of the periodogram is presented which allows for the fitting of a non-zero mean (the 'floating-mean' periodogram) or a multi-term sinusoidal model. This periodogram is implemented using the `astropy.timeseries` class `LombScargle`. In this periodogram, the periodic model is fit with a floating mean (the so-called generalised periodogram, Zechmeister & Kürster (2009); Cumming et al. (1999)). This reduces the likelihood of poorly fitting the model and determining an incorrect period in cases where the sampling is particularly unevenly spaced or sparse, both of which are possible with the data we are processing.

Due to the typical nightly observing schedule of GOTO we expect to see strong aliasing of any true period peaks with an observational period of one day, or for weakly periodic sources the power of the window function peaks may be orders of magnitude larger that of the true period peaks. In order to mitigate this we first calculate the window function for the set of observations by producing a periodogram for a 'constant' light curve sampled at the same times as the data. The frequencies and powers of the 10 best peaks are saved as the window function.

For a source light curve of magnitudes m_i with uncertainties σ_{m_i}, sampled at times t_i, these quantities are read into the search function. A number of bootstrap repeats, n, is chosen. Typically we use $n \sim 20 - 100$ in the periodicity search pipeline, as values in this range provide a good compromise between speed and accuracy. We found that values of $n < 10$ tend to result in large uncertainties, especially for low-power periods, while values of $n > 100$ provide diminishing returns while the computational cost increases. For each repeat, each of the light curve points is re-drawn from a normal distribution centred on m_i with standard deviation σ_{m_i}, such that we create a new light curve with magnitudes $m_i' \sim \mathcal{N}(m_i, \sigma_{m_i})$. The uncertainties and sample times are carried over from the original light curve. This new light curve is then used to create a Lomb-Scargle periodogram using the

`astropy.timeseries.LombScargle` class.

The 10 strongest periodogram peaks are saved as period candidates. Here we base the procedure on that detailed in Shappee et al. (2014): the frequencies of these candidates are compared to the window function, and any that lie within some tolerance of the window function peaks are discarded. We set the tolerance at a fractional difference of 1%. Following the elimination of aliased peaks, we select the strongest remaining peak as the best estimate. We check the power of the peak against the null hypothesis periodogram to estimate the peak significance, using the `false_alarm_level` method. For a false alarm probability (FAP), the probability of a peak of the observed height or higher under the null hypothesis, this method returns the corresponding peak power which is compared to the observed peak power. We use a FAP of 0.001, roughly the 3-sigma p-value. If the peak is significant, this period is saved and we proceed to the next iteration. If not, or if no period peaks were found, then we do not save the period and proceed to the next iteration. After all n iterations, we calculate the mean and standard deviation of the saved periods. The mean period, P is our best period estimate and the standard deviation, σ_P, is an estimate of the uncertainty. If the bootstrap method returns a period, then this period is used to fold the unphased data to produce a phased light curve.

It should be noted that this method is not necessarily sensitive to periods at $2P$, $0.5P$, $3P$, etc., and indeed assumes that the candidate peak with the highest power is the true period. To attempt to mitigate this issue we implemented a crude template fitting algorithm; fitting sinusoids at $\frac{1}{2}$ and $2\times$ the best period to the light curve and comparing the reduced χ^2 statistic in each case. However, this method failed to improve over the base case.

Multi-harmonic analysis of variance (MHAOV) This method involves the application of analysis of variance (ANOVA) to a harmonic sine series such that the maximum likelihood occurs at the true period, if present. As such, it is sensitive to non-sinusoidal light curves if the number of terms in the harmonic series is greater than one. The model used in this method is given by the expression

$$m_{bi} = \beta_{b0} + \sum_{k=1}^{K} a_{bk} \sin\left(k\omega t_{bi} + \phi_{bk}\right) + \epsilon_{bi}, \tag{4.3}$$

which describes a harmonic sine series of K terms (Mondrik et al., 2015). The multi-band data which correspond to this model are composed of n_b observations in each band b, with an observation time t_{bi}, a magnitude m_{bi}, and magnitude uncertainty σ_{bi} for each observation i, for a total of $\sum_b^B n_b$ observations over B bands. It is assumed that the frequency, ω, is constant between bands. β_{b0} define the magnitude

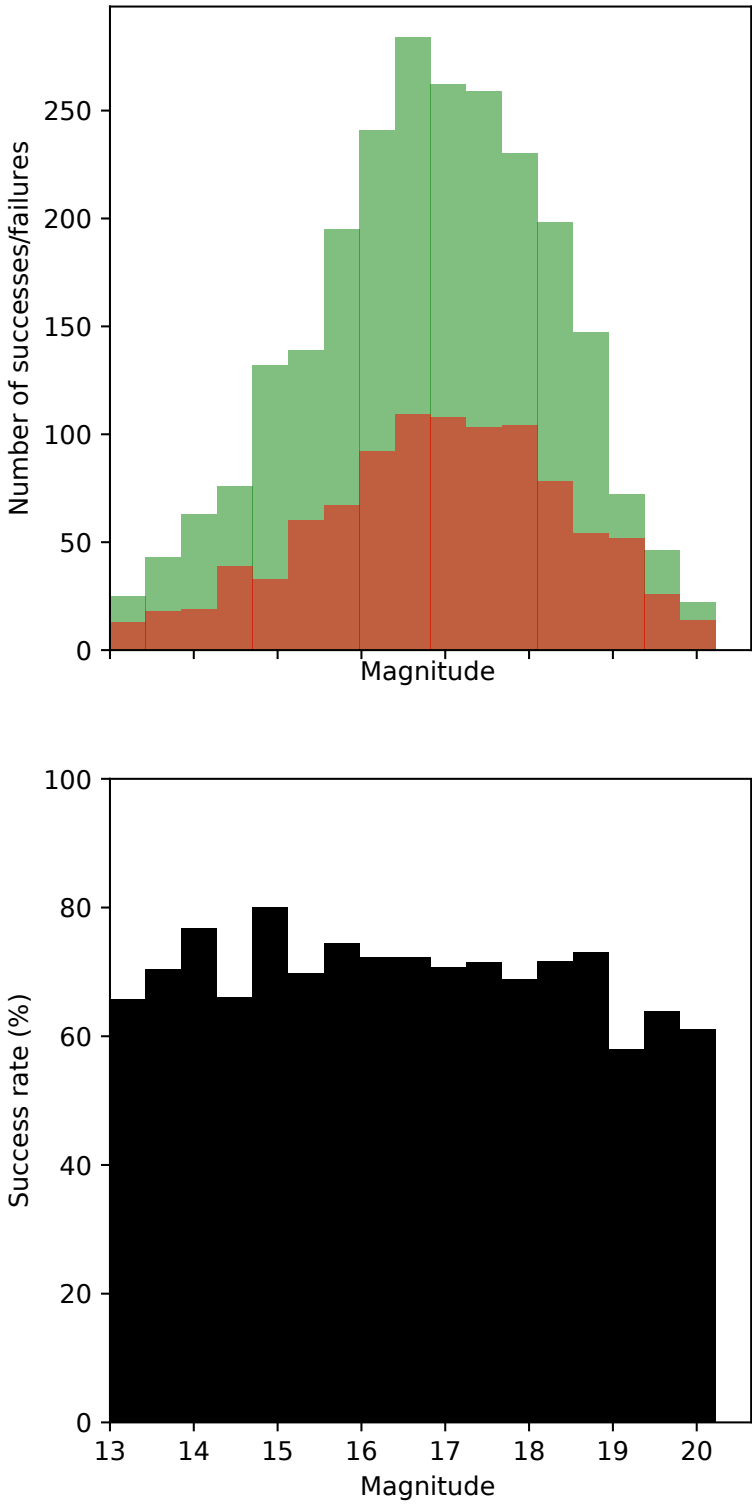

FIGURE 4.7: *Top:* Histogram showing the number of successfully (green) or unsuccessfully (red) recovered periods for a simulated GOTO field with the GLS algorithm. The red bars are overlaid on the green bars, not stacked. *Bottom:* Success rate for the GLS algorithm, showing relatively consistent performance across the range of magnitudes with a slight drop-off near the magnitude limit.

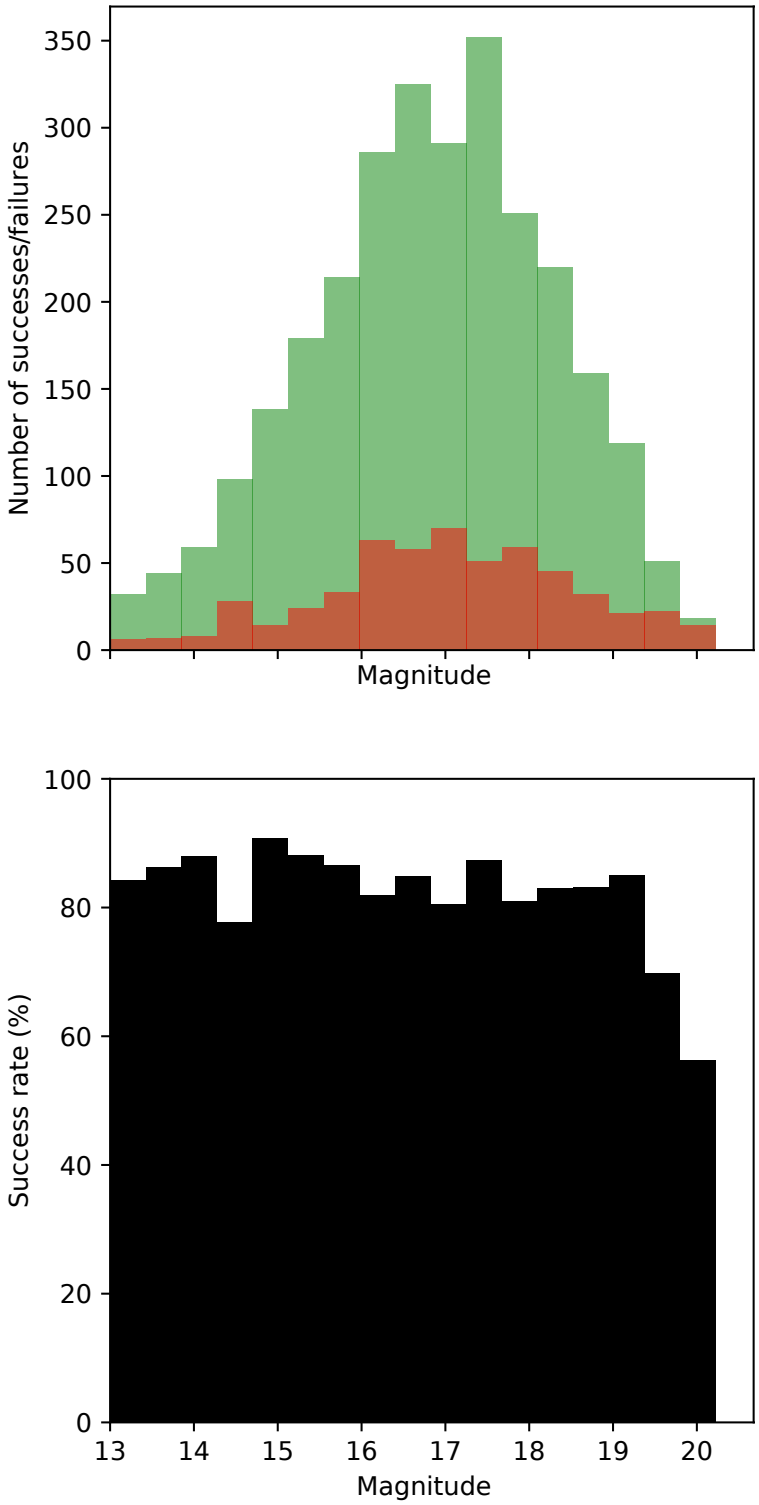

FIGURE 4.8: *Top:* Histogram showing the number of successfully (green) or unsuccessfully (red) recovered periods for a simulated GOTO field with the MHAOV algorithm. The red bars are overlaid on the green bars, not stacked. *Bottom:* Success rate for the GLS algorithm, showing consistent performance across the range of magnitudes with a drop-off near the magnitude limit.

offsets (means) of each band, a_{bk} are constants, and $\epsilon \sim N(0, \sigma_{bi}^2)$ are independent Gaussian noise terms.

As with the GLS periodogram, finding the frequency which minimises the χ^2 statistic is equivalent to finding the maximum likelihood. Note that in the single-band case, this periodogram is equivalent to the case presented in Schwarzenberg-Czerny (1996b). The ANOVA statistic is defined by the expression

$$F = \frac{\mathrm{Var}(F)}{\mathrm{Var}(E)}, \tag{4.4}$$

for observations $X_k = F_k + E_k$, where F_k is the signal, E_k is the noise. Here, the statistic represents the ratio of the variance (Var; the mean sum of squares) of these two components, effectively separating the variability due to the signal from the variability due to the noise.

We use a Python implementation of this periodogram, 'P4J' (Huijse et al., 2018), and a moving-block bootstrap method (see e.g. Bühlmann (2002); Suveges (2012)) to determine the periodogram peak significance, also included in the Python package. Note that again we use the single-band case, as GOTO photometry is single-band at the time of writing, though extension to the multi-band case should in fact yield more favourable results (Mondrik et al., 2015).

Initial performance To establish which of the two algorithms was more suitable, we tested both algorithms on a set of synthetic GOTO light curves. We used a Type A dataset of 800 synthetic sinusoidal light curves, as described in section 4.2.1. Aside from the number generated and a variability fraction of $f_{\mathrm{var}} = 1.0$, the default simulation parameter values of table 4.2 were used. Note that this results in a Gaussian distribution of the mean magnitudes of the light curves; this was chosen to reflect the approximate distribution seen in Spider light curves. Defining a 'success' as a period recovered to within 1% of the true value, we found an average success rate of 66% for the GLS periodogram. With a processing time of 0.5 s per source, this method trades off accuracy for speed. A histogram showing the performance as a function of source magnitude is shown in figure 4.7.

We performed the same test with the MHAOV algorithm, with the results shown in figure 4.8. For this comparison we set the number of harmonic terms to 1 for two reasons. First, to be a direct comparison of algorithms with the GLS periodogram, and secondly because this test dataset uses only single-term sinusoidal light curves. The mean success rate was 84% with a processing time of 1.3 s per source. Since we will not be performing these searches in real-time, the speed is less of a factor in algorithm choice than the accuracy. Note that in this initial testing sample, the

GLS algorithm shows marginally better performance with faint sources. However, not only is this sample size small, but it also does not take into account performance with non-sinusoidal light curves - this is discussed in section 4.3.3. As such, we proceed using the MHAOV method.

Periodogram peak significance

While the MHAOV method of (Huijse et al., 2018) produces robust periodograms, we required a method with which we may determine the statistical significance of a periodogram peak of a given power. Due to noise, aliasing, random fluctuations, outliers, and other erroneous data points, it cannot be assumed that every periodogram peak represents a true period. As such, we used a bootstrap method to determine the statistical significance of periodogram peaks. In this case the null hypotheses, H_0, is that the light curve contains no periodic signal and is solely white noise: the observations are normally distributed around a mean brightness. We may assume this as measurement errors are statistically independent, as opposed to data which contain correlated (red) noise.[1] The alternative hypothesis, H_1, is that there is a periodic signal in the data. The alternative hypothesis is accepted if the probability of observing a periodogram peak at the observed (or greater) power under H_0 is greater than some threshold, the so-called false alarm probability (FAP). However, the distribution of the periodogram maxima is not well defined, and cannot be assumed to be Gaussian. Instead, we follow the procedure outlined in Suveges (2012).

First, periodograms are created for a pool of m light curves generated using random resampling and k periodogram maxima are recorded. A moving block bootstrap method is employed instead of a i.i.d (independent and identically distributed) bootstrap, where the data are randomly resampled with replacement. i.i.d bootstrapping also removes any temporal correlation, such as from the observation window function, which results in a misrepresentation of the probability distribution. Instead, the moving block bootstrap method achieves resampling by dividing the light curve into blocks of some length, a number of days T_B, and shuffling these. We use a block length of $T_B = 10.0$ days, such that the window function, dominated by the daily observations of GOTO, is preserved while any periodic signal is destroyed. Note that the ideal block length depends on several factors, such as the observing schedule, and so this value of T_B should not be taken as the optimum. We generate $m = 100$ synthetic light curves using this bootstrap method and record $k = 20$ maxima from each.

[1] For example, we may expect the noise component of photometry of exoplanetary transients to be red (Baluev, 2013)

An extreme value probability density function (PDF) is fit to these maxima. The extreme value PDF describes the probability distribution of maxima of a random variable and is used to estimate the frequency of occurrence of extreme events of a given magnitude. A typical example given is in predicting the probability of an extreme weather event occurring, such as a catastrophic flood, in geological engineering (e.g. Fowler & Kilsby (2003), Sousa et al. (2011)). In this case, it is used to estimate the probability that a prominent periodogram peak is real and not expected from random noise. The generalised extreme value (GEV) distribution,

$$G(z) = \exp\left[-\left(1+\xi\frac{z-\mu}{\sigma}\right)^{-\frac{1}{\xi}}\right], \tag{4.5}$$

describes the distribution of these maxima, once rescaled. Here, $G(z)$ is the distribution of N independent random variables $Z_1, .., Z_N$, ξ is the shape parameter, μ is the location parameter, and $\sigma > 0$ is the scale parameter. ξ and μ can take any real value, and specific configurations of these result in specific extreme value distributions, such as the Gumbel distribution for $\xi = 0$. In general, ξ defines the shape of the tail of the distribution, with larger values corresponding to longer decays and negative values resulting in a terminating distribution. These parameters are fit to the distribution of periodogram maxima using the `genextreme` class in the `scipy.stats` Python package. From here, the expected periodogram power of a peak with significance p is determined from the associated PDF confidence level, taken to be 3σ or a p-value of 0.0027. If a periodogram peak has a power greater than the corresponding power from the distribution, then the null hypothesis is rejected and it is assumed that this peak is significant.

However, this approach can also indicate that true periods are not significant, for example where the periodogram is substantially contaminated with aliased peaks. This is illustrated in figures 4.9 and 4.10 which show an example periodogram alongside the GEV distribution and light curve respectively. This MHAOV periodogram was created using a light curve which showed clear periodicity chosen from a test GOTO field centred around the candidate FERMI source 4FGL J1818.6+1316. Despite the clear periodicity in the light curve, the periodogram is heavily contaminated by aliased peaks, significantly increasing the power required for a significant period detection. This problem was encountered throughout the investigation, revealing the extent to which aliasing is an issue with these light curves.

4.3.3 Algorithm performance

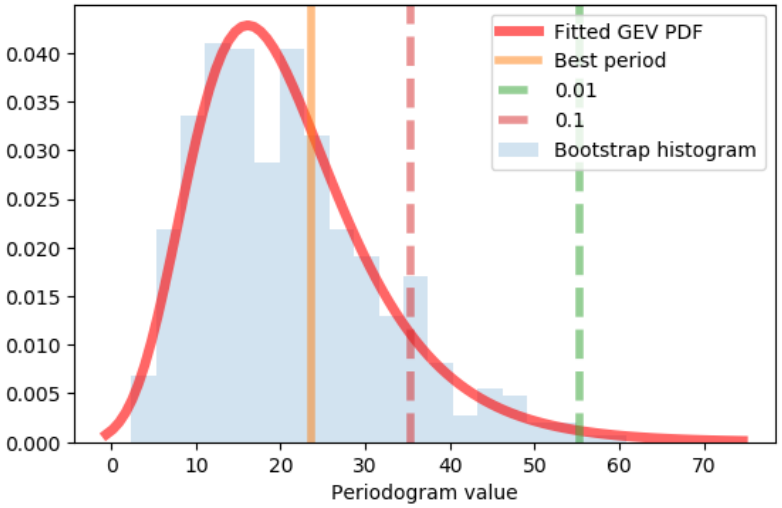

FIGURE 4.9: Fitted GEV distribution for a light curve observed by GOTO. The powers corresponding to the 0.1 and 0.01 p-values are shown by red and green dashed lines respectively, while the power of the best period (as shown in figure 4.10 is in orange. The blue histogram shows the distribution of maxima of the bootstrapped periodograms, with the fitted GEV in red.

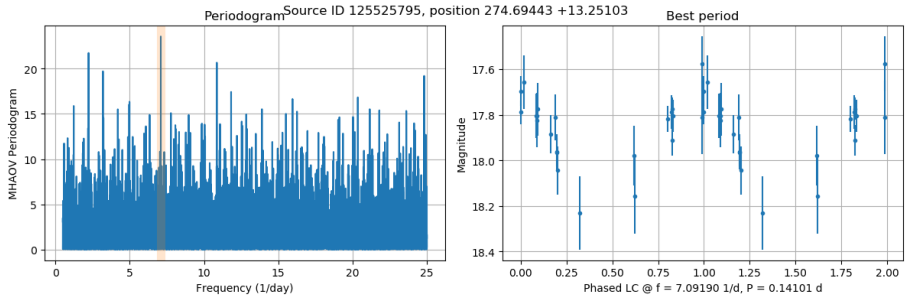

FIGURE 4.10: MHAOV periodogram (left) and GOTO light curve folded at the best period (right) for a test GOTO source. The best period is highlighted in yellow on the periodogram.

Aliasing and jitter

It is well known that periodic structure in the observation times of a signal leads to aliasing in periodograms (Deeming, 1975). This is distinct from the aliasing that occurs when the observation frequency is below the Nyquist limit, and is instead caused by structured observing windows. For ground-based observatories like GOTO, this typically results in strong window function peaks at a period of 1 day (caused by nightly observations), with related aliases. This structure therefore affects the periodogram, which is a convolution of the true spectrum and the power of the window function. If the true periodogram power is weaker than the aliased peaks, such as when the periodic signal is not strongly modulated or the SNR is low, then the aliased peaks will dominate the power spectrum. However, this can

also occur even with strong periodic signals if the observation window function is strongly periodic, for example if all observations are regularly spaced. As can be seen in figure 4.11, this is an issue which GOTO suffers from. This window function was created by producing a periodogram from a time series $y_i(t_i)$, where y_i is unity for each t_i and t_i are the observation times of a sample GOTO light curve. The number of t_i samples used was 43, the median number of light curve observations in a random sample of 10,000 GOTO sources. In addition to the peak at a frequency of 1 d^{-1} (and aliases spaced every $n = 1, 2, 3...$ d^{-1}), substantial structure can be seen surrounding the main peak. This suggests that there is additional periodicity in the observing schedule. Indeed, GOTO frequently observes the same field several times in a given night: this window function illustrates that these observations tend to have a periodic schedule.

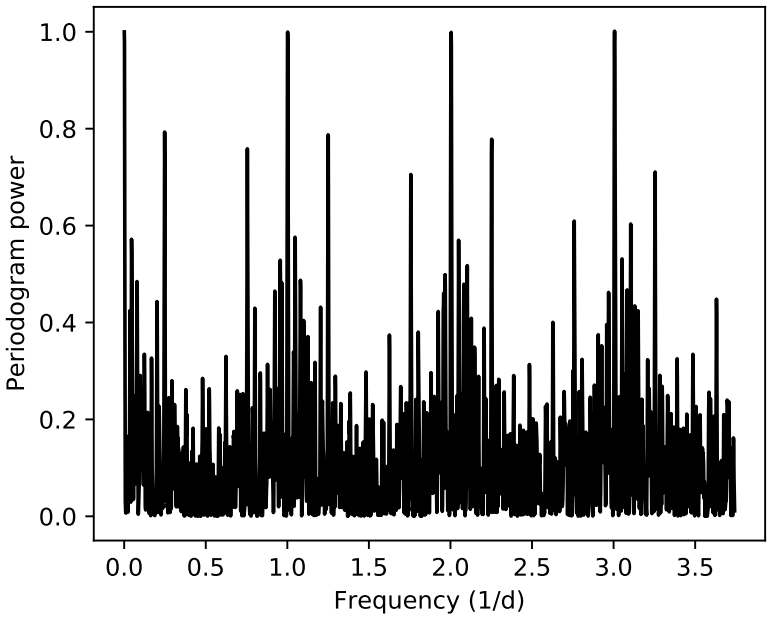

FIGURE 4.11: Window function for a sample GOTO light curve with the median number of samples, 43. The clear peak at 1.0 d^{-1} can be seen. Note that the structure repeats every 1 d^{-1}; these are aliases of the peak at 1.0 d^{-1}.

For comparison with figure 4.11, figure 4.12 shows the window function of a typical ASAS-SN light curve. Again, the number of samples used is the median number of light curve observations in a random sample of 10,000 ASAS-SN sources. The window function shows one dominant peak at 1 day with aliased peaks every n days, whose amplitude decays with $1/n$. As described in Jayasinghe et al. (2019), it is significantly simpler to mitigate the effects of aliasing with window functions like these. Aliased periodogram peaks can be easily identified and eliminated as the

simple structure and well-separated peaks of the window function mean that the periodogram is much less contaminated.

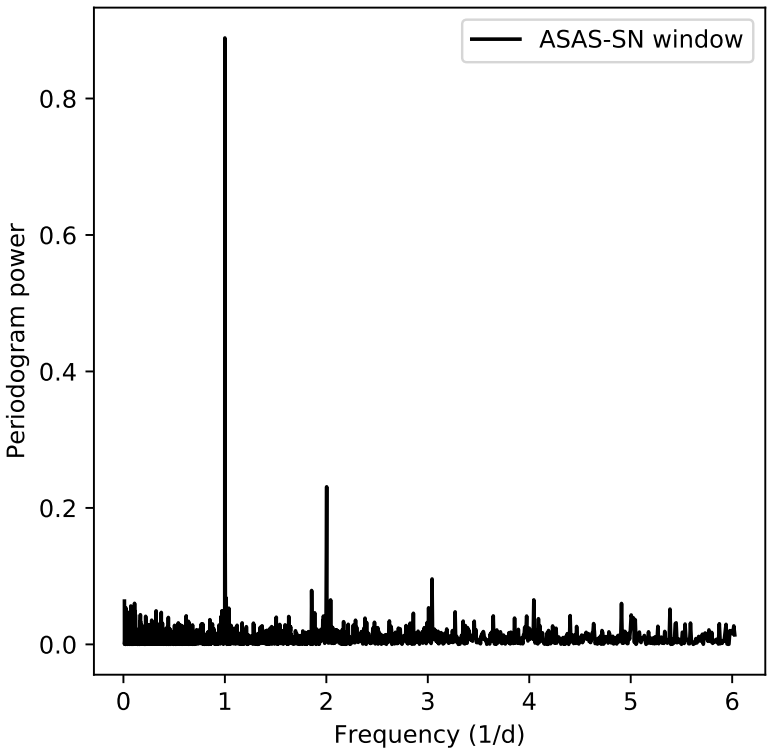

FIGURE 4.12: Window function of an ASAS-SN light curve with the median number of samples, 201. The main peak is at 1 d^{-1}, as with the GOTO window function, however there is significantly less additional structure.

When testing the periodicity search algorithms with real GOTO data, this aliasing has a significant effect. We ran the MHAOV algorithm on 86 GOTO light curves obtained from a pilot field centred around the black widow pulsar, PSR B1957+20 (Fruchter et al., 1988); one of these is shown next to its periodogram in figure 4.13. This figure illustrates one of the main problems encountered: that the low SNR of the data and structured observing window cause significant aliasing, obscuring true the period. The consequences of this in the pipeline are clear; missed periodic sources, and incorrect periods. The former means that many candidates will be thrown away, and the latter means that the light curves will be incorrectly folded, rendering them useless.

Therefore, we investigated if small changes to the observing schedule could sufficiently alter the window function and mitigate the problem. Returning to the simulation, we synthesised groups of 1000 Type A light curves (as described in section 4.2.1, these are sinusoidal light curves with 0.1 d $< P <$ 2 d), each with slightly different sample times. Each group uses the same 60 observations per 120 days

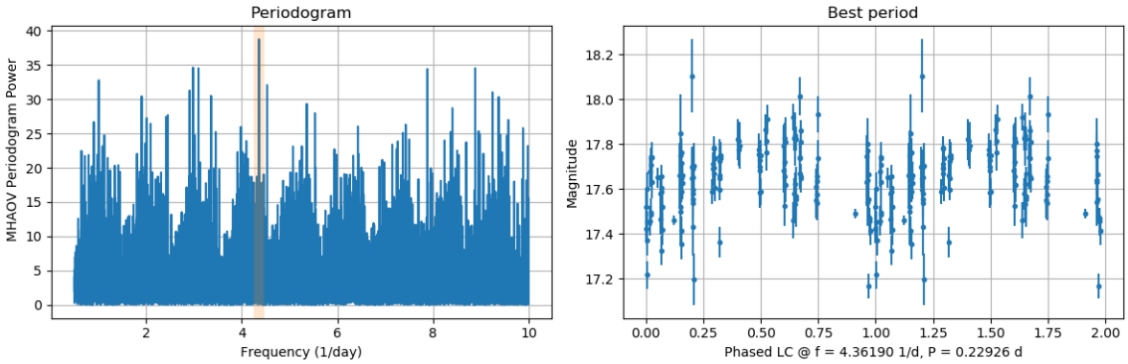

FIGURE 4.13: Periodogram and folded GOTO light curve of a variable source. The aliased peaks, separated by $\Delta p = 1$ day are visible in the periodogram.

schedule, however different amounts of jitter are added to the observation times each night. Here, jitter refers to scatter added to the observation times at night, drawn from a uniform distribution with some width j. For the null hypothesis case, we use a jitter of $j = 0$ days, such that the observations are at the same time on a given night, and increase the jitter for subsequent groups. We use roughly logarithmically spaced values of $j = [0.001, 0.01, 0.05, 0.1, 0.2, 0.3, 0.5]$ days. We compare the performance of the MHAOV algorithm on each set of light curves by scoring each by the percentage of periods recovered to within 1% of the true value. The results are shown in figure 4.14. We see that even with a small amount of jitter applied, the accuracy rapidly increases to a plateau of around 84%, approximately the previously stated performance of the MHAOV algorithm. The optimum scatter is around 0.1 to 0.2 days, or $\pm 2 - 5$ hours.

Testing data

While we were able to recover the true period more than 80% of the time when testing on synthetic sinusoidal light curves, applying these methods to real data is more complicated. As such, we used 1600 synthetic Type B light curves (described in section 4.2.1) to test the MHAOV and GLS algorithms, 200 of each star type. We used an exposure time of 180 s, with 60 samples over 120 days, and allowed for jitter of 0.2 hours.

Factors influencing period recovery

Light curve shape As figure 4.15 illustrates, the shape of the light curve can have a substantial effect on the algorithm performance. The algorithm scores are summarised in table 4.3. The MHAOV algorithm shows good performance on light curves with clearly defined variability, like RR Lyrae-type variables with scores of

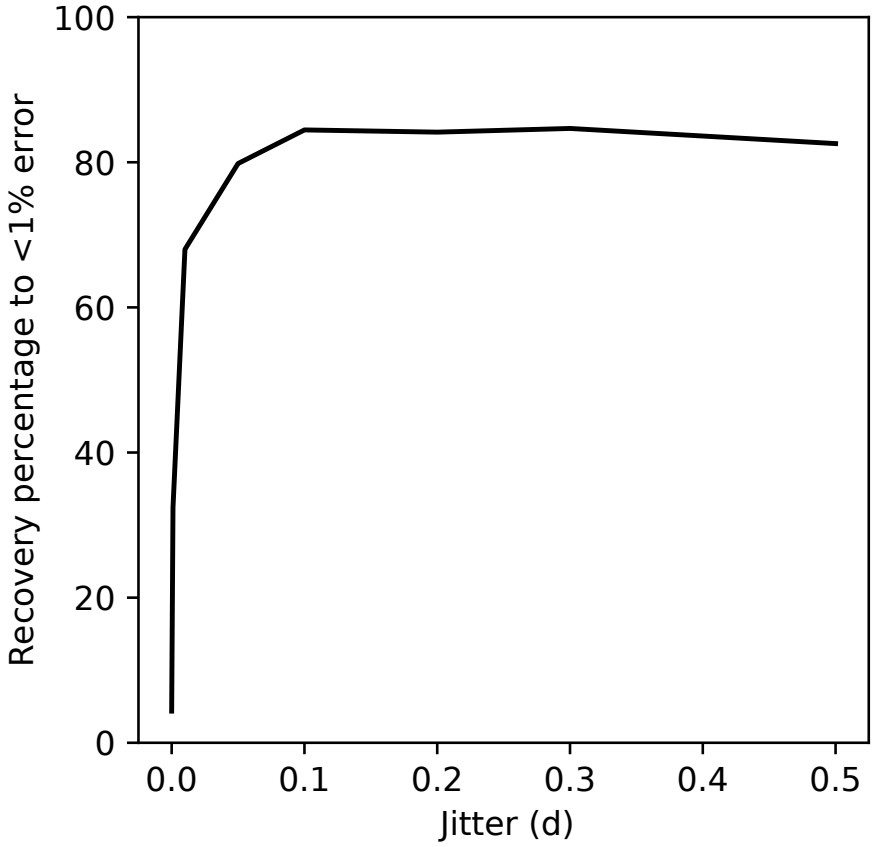

FIGURE 4.14: Relationship between jitter and period recovery for synthetic light curves. The accuracy rapidly increases as jitter is applied

TABLE 4.3: Percentage of sources with period recovered to $< 1\%$ of true period. Sources with asterisks have doubled periods.

Star Type	MHAOV score (%)	GLS score (%)
SR	79	52
M	84	67
RRC	98	94
RRAB	92	84
ROT	75	56
EA	41	13
EB	63	7
EW	46	1
EA*	53	74
EB*	32	84
EW*	50	97
EA+EA*	94	87
EB+EB*	95	91
EW+EW*	96	98

98% and 92% in the RRC and RRAB classes respectively. The SR (semi-regular variables) class contains stars with irregular periods, which often display either several cycles of variation in one ASAS-SN period or very noisy light curves, so the slightly poorer score of 79% is not unexpected. Similarly the spotted binaries (ROT) have regular periods but often irregular brightnesses, hence the middling performance of 75%.

The eclipsing binary classes, EA, EB, and EW show some of the poorest scores, of 41%, 63%, and 46% respectively. At first this appears to be an unusual result as these classes typically host some of the most well-defined periodic light curves. However, they typically contain two distinct eclipses, often of comparable depth. If the light curve does not contain enough information to differentiate between the two eclipses then the periodogram peak at half the true period may (incorrectly) contain more power than the peak at the true period. This can be shown by multiplying the recovered periods in these classes by two and comparing to the true periods. For these doubled periods, we obtain scores of 53%, 32%, and 50% respectively. We see that the 2nd harmonic is dominant in these light curves and that the periodicity search algorithm detected half the true period for a substantial number of sources. Combining these two cases, we find that a total of 94%, 95%, and 96% of eclipsing sources have their periods correctly determined at either the fundamental or second harmonic; a score more in line with the other classes.

Indeed, this is a long-standing issue in periodicity searching; without additional data (such as RV curves) it is non-trivial to determine which of the periods (P or $2P$) is correct [2]. This problem becomes even more difficult to solve if the light curves are sparsely sampled or are of low sensitivity, as the features which distinguish the two cases are often subtle. The simplest method is visual inspection of the light curves folded at P and $2P$, however this is not scalable to the size of the datasets used here. Beyond this, the SuperSmoother algorithm (Reimann, 1994; VanderPlas & Ivezić, 2015) can be used to fit the observed light curve over the range of candidate periods to determine which is the true fundamental period. This method uses a non-parametric running linear regression on the data, and so does not require a-priori knowledge of the light curve shape. However, it is computationally expensive so cannot be realistically used as a standalone period search method.

The performance of the GLS algorithm is consistently several percentage points poorer than the MHAOV algorithm; notably, the score for the eclipsing variables is significantly worse. However, by using the same method and doubling the recovered periods, we see a huge improvement which indicates that the GLS algorithm

[2] This notation assumes that the true period is either P or $2P$; in reality it could equally be P or $\frac{P}{2}$.

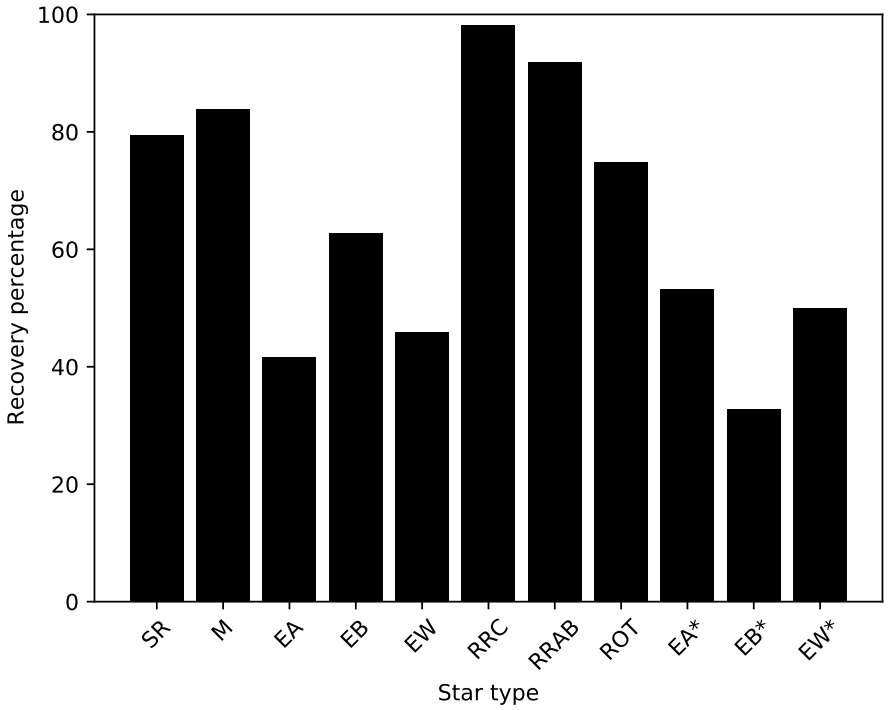

FIGURE 4.15: Percentage of periods recovered using the MHAOV algorithm to within 1% of the true period for 1600 sources, separated by star type. The classes on the right with asterisks show the percentage recovered when the period is multiplied by 2.

identified the incorrect harmonic most of the time. This difference in performance is expected from the two periodograms, as the MHAOV method fits a three-term harmonic series as opposed to the single-term fitting done using the GLS algorithm. This difference is shown in figure 4.16, where the difference been the light curve folded at the true period and at half the true period can be seen. We also performed the same investigation using the MHAOV algorithm with one harmonic term, as a direct comparison to the GLS method. This resulted in comparable performance, though the mean MHAOV was still several percentage points higher than the GLS score. In general we may surmise that the more higher-order harmonic content in a light curve, the poorer the period recovery, especially for algorithms which only fit a single sinusoidal term.

4.4 Classification of periodic sources

Given the enormous number of light curves ($\gtrsim 10^5$ per pointing) we obtain with GOTO, it is clear that manual classification is not feasible. Instead, we employ machine learning (ML) techniques to attempt to identify the light curves which fit our desired description. We use a Random Forest classifier, introduced in section

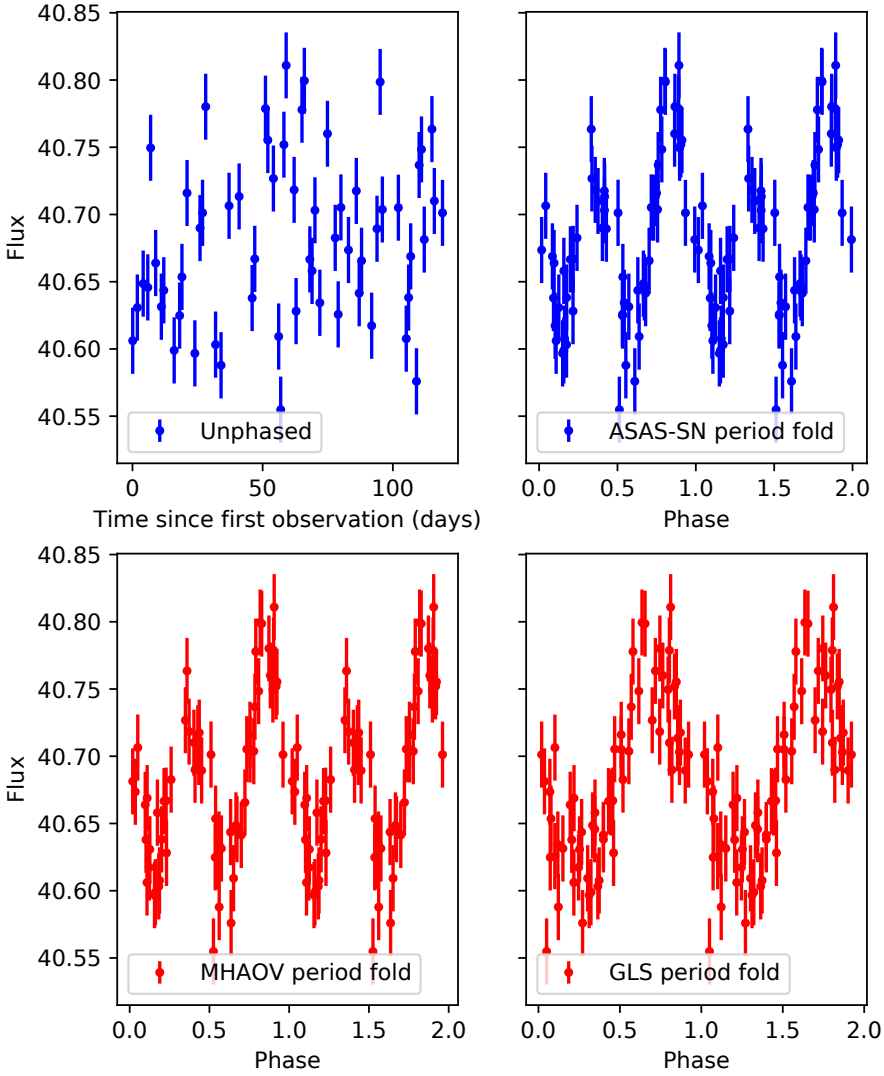

FIGURE 4.16: EW type binary, from top left clockwise: unphased light curve, phased at ASAS-SN period, phased at GLS period, phased at MHAOV period. The flux is in arbitrary units. The MHAOV period is clearly a good match to the true ASAS-SN period, while the GLS period is at double the true period. Data repeated for clarity.

1.6.3, to achieve this. The classifier is trained on a large dataset of variable star light curves (the ASAS-SN catalogue of variable stars) using intrinsic properties like absolute magnitude, period, and colour information as features, as well as the shape of the light curves. The light curve shape is parameterised using a Fourier decomposition, described in section 4.4.1[3]. This builds on the work of an earlier Master's research project by Matt Wheeler and Isaac Richardson under the supervision of R. P. Breton. This project developed a method of parametrising the shape of periodic light curves from the ASAS-SN catalogue in terms of a finite number of Fourier

[3]Note that this method is also used in the creation of Type B datasets.

harmonic amplitudes and phases, to be used as features for a Random Forest classifier. This classifier was trained and validated on the same dataset of ASAS-SN light curves used in this work, and as such can be considered as a proof-of-concept for the viability of such a classifier trained on features derived from the light curve shape. In this work we will refer to this research as the 'precursor study'. In section 4.4.2 the training and testing datasets used with this classifier are discussed in detail.

4.4.1 Light curve modelling

Using a least squares estimation (LSE) process we model the phased light curves using a Fourier series,

$$m(\phi) = A_0 + \sum_{i=1}^{12} \left(a_i \cos(2\pi i\phi) + b_i \sin(2\pi i\phi)\right), \tag{4.6}$$

where $m(\phi)$ is the source brightness at phase ϕ and A_0 is equal to the mean brightness of the source. Note that this model can be applied to light curves of both flux and magnitude. We truncated the series at 12 terms, a limit determined during the precursor study using Bayesian information criterion (BIC) (Schwarz, 1978) tests using Catalina Sky Survey (CSS) data. BIC is a model selection technique that is used here to determine the optimum number of model parameters by introducing a penalty term which scales with the number of parameters. This 12-term limit allows us to model light curves of a very broad range of shapes without over-fitting the data. Using the coefficients a_i and b_i as parameter terms, $\mathbf{a}_i$, we define a hypothesis function

$$h(\mathbf{x}) = \sum_{i=0}^{N} \mathbf{a_i} \mathbf{X}_i(\mathbf{x}), \tag{4.7}$$

where $\mathbf{X}_i$ are the basis functions of the predictor variables, $\mathbf{x}$. The hypothesis function defines the model used for the fitting, with N terms. Note that A_0 is absorbed into a_i. With this we define the cost function, C, as

$$C = \sum_{k=1}^{M} \left(\sum_{i=0}^{N} \frac{a_i X_i(x_k) - y_k}{\sigma_k} \right)^2 + \lambda \sum_{j=0}^{N} W_j(a_j)^2, \tag{4.8}$$

$$C = \chi^2 + \lambda \sum_{j=0}^{N} W_j(a_j)^2, \tag{4.9}$$

where the χ^2 statistic, summed over M datapoints $k = 1, ..., M$, is joined by a weighting term. This weighting term arises from the L2 regularisation we use to constrain the magnitude of higher-order harmonic terms and reduce the impact

of over-fitting. Without this regularisation these higher order terms can become large when fitting sparsely sampled light curves. λ is the Lagrange multiplier for this regularisation, and W is a diagonal weighting matrix. To minimise the cost function we differentiate with respect to a and express it in matrix notation,

$$\mathbf{a} = \left(\mathbf{X}^T\mathbf{X} + \lambda\mathbf{W}\right)^{-1}\mathbf{X}^T\mathbf{b}, \tag{4.10}$$

where we define the design matrix

$$X_{kj} = \frac{X_j(x_k)}{\sigma_k} \tag{4.11}$$

and the vector $b_k = \frac{y_k}{\sigma_k}$. To parameterise the weighting matrix we define a 'critical harmonic',

$$H_c = \frac{1}{2G_m}, \tag{4.12}$$

where G_m is the maximum gap in phase between two consecutive data points. We use a different weighting scheme dependent on the critical harmonic. The weighting scheme is based on the Fermi-Dirac distribution: harmonics $H_i \leq H_c$ are unaffected, while harmonics $H_i > H_c$ are increasingly suppressed to zero. For light curves with a critical harmonic of $H_c < 9$, we use the scheme

$$W_i = 1 + \frac{100}{1 + 99^{\frac{H_{\max}+H_c-2H_i}{H_{\max}-H_c}}}, \tag{4.13}$$

and for light curves with $H_c \geq 9$,

$$W_i = 1 + \frac{100}{\left(1 + 99^{\frac{2}{3}(H_{\max}-H_c)-1}\right)}. \tag{4.14}$$

In these equations, $H_{\max} = 12$, the maximum harmonic, and H_i is the fitted harmonic value. Note that H_i does not appear in equation 4.14, as this weighting scheme is constant for all harmonics. We use k-fold cross validation to estimate the optimum value of λ for each light curve. This is a technique commonly used for evaluation of ML classifier performance: the data are divided into k subsets, $k-1$ of which are used to train the model and one to validate. Here, the light curve is divided such that $\frac{k-1}{k}$ of the data points are used to train the model and $\frac{1}{k}$ are reserved for testing. This is repeated k times with each of the k subsets used as the testing set once. Figure 4.2 shows an example result of the decomposition process. The light curve in this figure has a large gap in phase coverage, approximately phase 0.05–0.25, which will result in the suppression of higher-order harmonics.

4.4.2 The Random Forest classifier

Training and testing datasets

While there are several novel methods of light curve classification using deep learning, Random Forest classifiers continue to be among the most robust and successful (Mahabal et al., 2017). By design, Random Forest classifiers are robust to noisy datasets and outliers. We use the Random Forest classifier `RandomForestClassifier` implemented in the `scikit-learn` package (Pedregosa et al., 2011). As the GOTO light curves are unlabelled, we make use of two other published datasets when training and testing the classifier; the ASAS-SN catalogue of variable stars (Shappee et al., 2014; Jayasinghe et al., 2018), and data from the GAIA telescope (Gaia Collaboration et al., 2016). From the GAIA catalogue (Gaia Collaboration et al., 2018), we extract the G-band magnitude, the parallax and parallax error, RA and Dec positions, and the $B_p - R_p$ excess magnitudes. First, we will summarise the composition of the training and testing datasets. After this, we discuss the performance of the classifier.

ASAS-SN training data and classifier features Our training set comprises labelled light curves from the ASAS-SN catalogue. The testing set uses GOTO light curves, cross referenced with the GAIA catalogue for additional information like colour, and with the ASAS-SN catalogue for labels. The cross-referencing is performed by matching the positions and magnitudes of the sources; this process will be detailed shortly. The ASAS-SN catalogue of variable stars contains multi-band photometry of over 660,000 sources, the majority of which have defined periods and labelled types. The classifier is trained on a subset of this catalogue; we only select sources with one of the 8 star types in table 4.1, and ensure that the same number of sources are present for each. This is done to exclude unusual star classes such as those with very long or irregular periods and to combat class imbalance. This is also done to simplify the problem; by reducing the number of possible classes to a key few, the initial performance of the classifier can be assessed before moving onto a more complicated problem. In total, this data set is composed of 433,294 light curves. Training features are extracted from the light curves using the above harmonic decomposition method: the amplitudes and phases first 12 harmonics and the fundamental amplitude are used. To normalise these features, each of the higher order harmonics are given as a fraction of the fundamental amplitude. In addition to these harmonics and the fundamental, the features used are the absolute G-band magnitude, $B_P - R_P$ colour excess, period, and phase offset for a total of 27 features. The classifier described in Jayasinghe et al. (2019) uses several additional

colour terms, present in the ASAS-SN dataset. However, these were not available in the GOTO or GAIA data and so were not used as training features in the final classifier. In section 4.4.3 we discuss the potential impact this may have on our results.

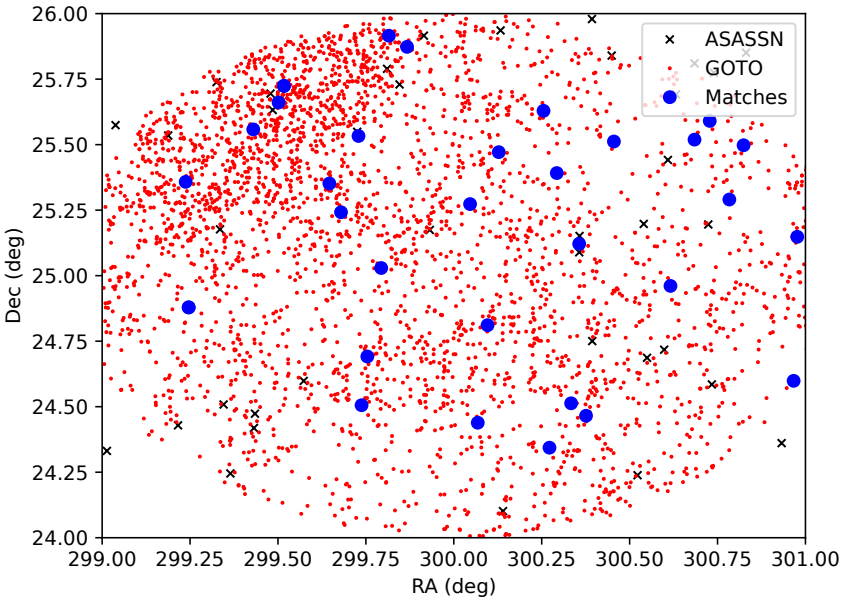

FIGURE 4.17: Positions of ASAS-SN (black cross) and GOTO (red dot) sources in a field of radius 1 degree. The GOTO sources in this field are only those with a significant detected period and with more than 10 observations in the light curve. The ASAS-SN sources are those in one of the 8 star types used in the classifier. Positive matches are indicated with a blue circle

GOTO testing data and source matching Since the sensitivity of GOTO (L$\sim$ 20.5 mag) is greater than that of ASAS-SN (V$\sim$ 17 mag), for a given field observed by each telescope there will be significantly fewer ASAS-SN sources than GOTO sources. This is exacerbated by the fact that we limit the number of stars to 8 types of variables. Figure 4.17 shows that, for a large field containing several thousand GOTO sources, there are only of order $\sim$ 50 ASAS-SN sources in the same region. This presents a major problem as matches between GOTO and ASAS-SN sources are required to label our light curves.

As such, to create the testing dataset, we search around each ASAS-SN source for a matching GOTO source. The process is as follows. For a given ASAS-SN source with some (RA, Dec) coordinates, we query the GOTO database. We use a search radius of 30 arcsec, which approximately corresponds to the 3-sigma radius of the ASAS-SN source localisation, based on the FWHM quoted in Jayasinghe et al.

(2018). We also query the GAIA database with the same parameters. Suppose that this search returns N candidate GOTO sources. For each of these we run the process detailed in section 4.2.2; we construct light curves from the raw data and remove erroneous observations. In order to be able to compare magnitudes, we cross-reference these N GOTO sources with the GAIA sources using the `match_coordinates_sky` function in the `astropy.coordinates` package.. The maximum allowed separation between the sources from the two catalogues is defined as the FWHM of the GOTO source. We obtain a successful match $\sim 99\%$ of the time, assessed using a one square degree GOTO test field. As some GAIA sources are missing magnitude or colour measurements, we fill any such missing values in the GAIA data with the respective median value once all sources have been matched. This is done as the `scikit-learn` classifier is not robust enough to deal with missing features (Breiman, 2001).

At this stage, several candidate sources are typically eliminated due to poor or no matches, such that we are left with $M < N$ candidates. These M candidates are then compared to the ASAS-SN source using two methods; first by the sky separation, using the `SkyCoord.separation` method in the `astropy.coordinates` package, and second by the G-band magnitudes. The ideal candidate will have both the smallest sky separation and smallest magnitude difference to the ASAS-SN source. If the same candidate fulfils both of these conditions, we assume that this is a successful match. If not, then we compare the fractional magnitude difference to a pre-defined threshold of 0.1%. The GAIA and ASAS-SN G-band magnitudes are well-calibrated, such that this threshold is small enough that few true matches will be eliminated. If neither of the conditions are met, then it is assumed there are no matching sources in the candidate pool and the ASAS-SN source is discarded.

Due to the difference in sensitivity between ASAS-SN and GOTO, the use of this method means that GOTO sources fainter than the magnitude limit of the ASAS-SN survey are not included in the testing set. This is the most immediate limitation of this method of generating the testing set, since it is no longer representative of the GOTO survey as a whole. This difference is illustrated in figure 4.18, which compares the magnitude distribution of periodic GOTO light curves in the pilot 1° field to those in the training and testing data sets. However, since these eliminated light curves are of relatively faint sources close to the magnitude limit of GOTO, it is unlikely that a significant fraction of them would have accurate period detections due to the low SNR. However, since some classes such as the eclipsing binary tend to be fainter sources, we may be selecting against sources of those classes.

The last step in this process is to phase-fold the GOTO light curves at their periods. Due to the poor performance of the periodicity search with real GOTO light curves, we instead use the periods of the associated ASAS-SN sources. To

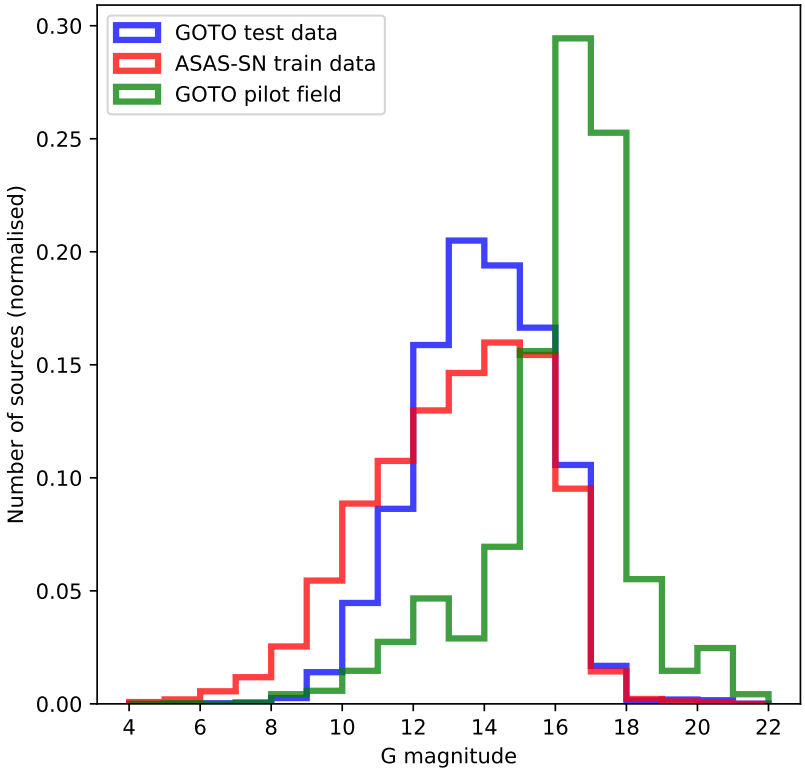

FIGURE 4.18: Normalised histograms illustrating magnitude distributions of three data sets; the training set of ASAS-SN light curves in red, the testing set of GOTO light curves in blue, and light curves from the GOTO pilot 1° field in green. We matched the GOTO sources to the GAIA catalogue to obtain the GAIA G-band magnitudes.

ensure that these periods are the 'true' periods of the GOTO sources, we perform a similar bootstrapping procedure to that described in section 4.3.2. To form the null hypothesis case, we create 50 periodograms using the moving-block bootstrap method and determine the periodogram power corresponding to a significance of 3σ as before. Instead of running a full periodicity search with the GOTO light curve, we calculate the periodogram power at the ASAS-SN period and compare to the 3σ power; if greater than or equal we add the source to the testing set. It is at this step that we discard the largest number of sources, again highlighting the issues of using the GOTO light curves in their current state for periodicity searching.

Indeed, at each stage of the process the number of sources in the testing data shrinks; we are left with 42834. However, the number of light curves of each type is capped at the least populous class (in this case, RRC) so that the testing set is balanced. As such, the final testing set is composed of 7128 light curves, 891 of each

type. The median number of samples per light curve is 43, with a standard deviation of 43. The minimum number of samples per light curve is 10 and the maximum is 685.

4.4.3 Classifier analysis

Feature importance

An appealing feature of the Random Forest classifier is that the relative predictive power, or importance, of each classification feature is easily determined. For each tree, the feature chosen at each node is that which maximises the decrease in prediction error compared to the previous node. As described in Breiman (2001), the features which decrease this error the most when averaged across all trees are those with the greatest importance. Figure 4.19 shows the feature importances of the final version of the classifier. The feature importances sum to unity. To examine the feature importances, we trained the classifier on our dataset of ASAS-SN variable stars. Validation was performed using k-fold cross-validation with $k = 10$. The dominance of the source period is immediately obvious, and is due to the fact that the period of variable stars is often a strong indicator of their type. For example, the SR class will typically contain the sources with the longest periods, and all greater than 30 days, while the eclipsing binary classes will typically contain some of the shortest period variables. After this, the (apparent) G magnitude and $B_P - R_P$ colour excess follow; the colour information and brightness are also typically strong predictors of star type. Following this is the percentage fluctuation, and then the amplitudes and phases of Fourier components. The predictive power of the Fourier components are individually relatively weak, however it should be noted that the predictive power is spread across the components.

Figure 4.20 shows the feature importances of a model trained on the same dataset but including the additional colour terms that are not present in the GOTO data; W1-W2, H-K, J-K, and J-H (Wright et al., 2010; Skrutskie et al., 2006). Here, W1 and W2 are derived from the Wesenheit magnitudes (Madore, 1982; Lebzelter et al., 2018; Jayasinghe et al., 2019). Since the sum of individual feature importances sum to one, we see that the lack of colour information in the final model is compensated for by a higher reliance on the period. The distribution of power is otherwise very similar. This difference between the two models indicates the importance of colour information in the classification of light curves. When the predictive power is concentrated in one feature, like here with the period, model performance and reliability can suffer. This is an especially prominent issue within this project because of the unreliable performance of the periodicity search algorithm on GOTO light curves.

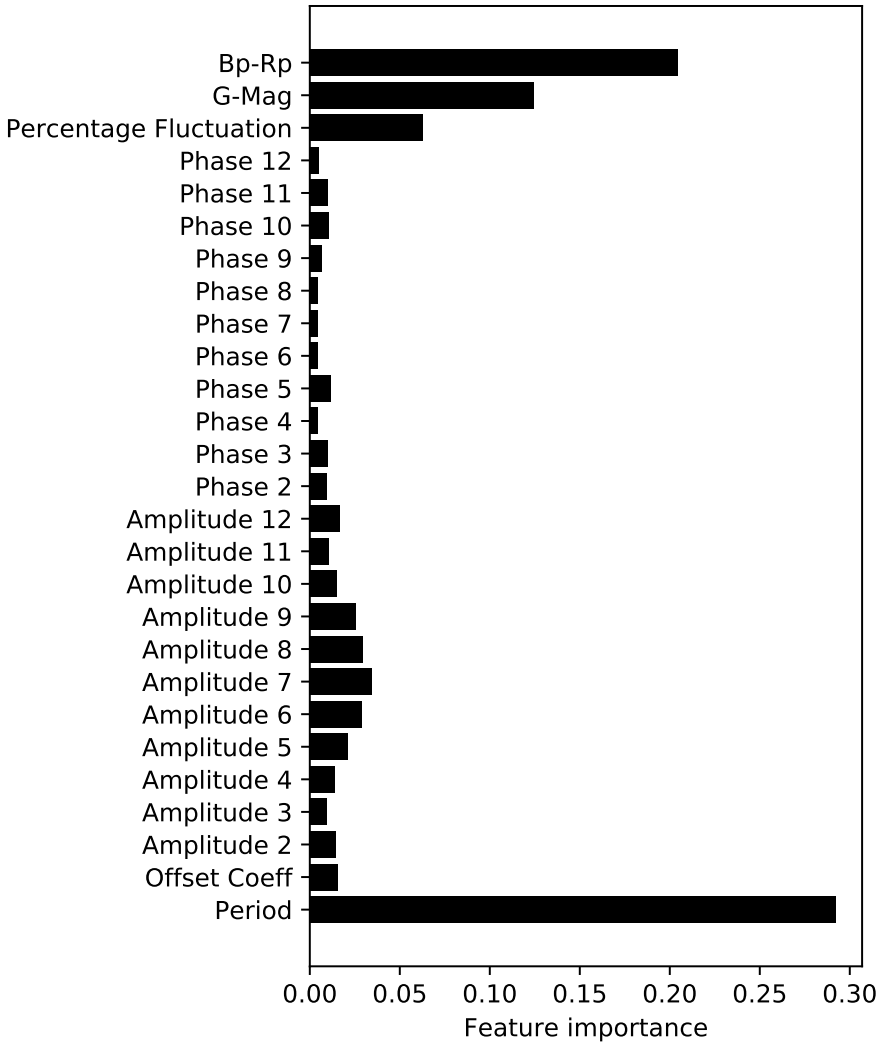

FIGURE 4.19: Feature importances of the final classifier. Note the dominant predictive power of the period.

The performance of this model, tested on a set of ASAS-SN light curves, is discussed in section 4.4.3 as a comparison to the model used with GOTO light curves.

Classifier performance

In this section we consider the performance of the classifier model excluding the additional colour terms. The confusion matrix for this classifier is shown in figure 4.21, which straight away shows that the classifier is not working as intended. As described in section 1.6.3, the ideal confusion matrix is diagonal; here we see significant off-diagonal terms. Clearly, this is not a promising result for the project.

It can immediately be seen that the performance of the classifier is extremely class-dependent, to the extent where some classes (M, RRC) are never predicted.

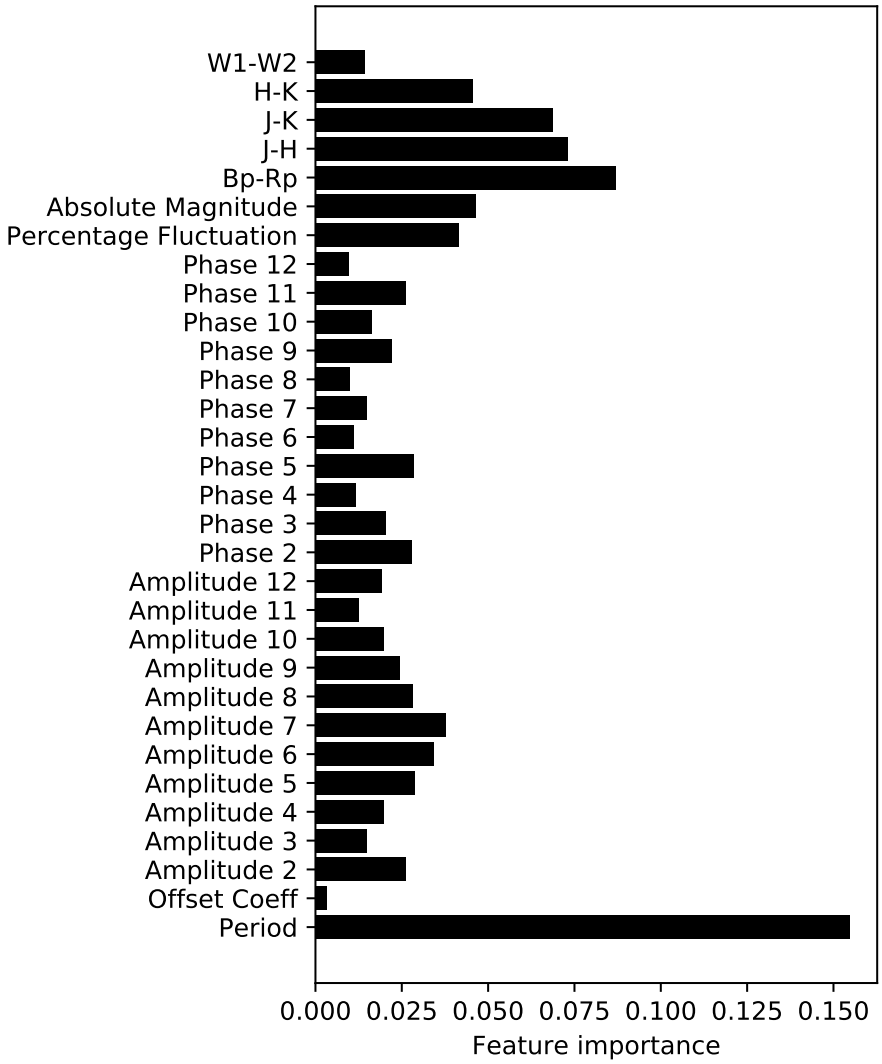

FIGURE 4.20: Feature importances of the original classifier, including all colour terms. W1 and W2 are Note the concentration of predicted power in the colour terms, reducing the reliance on the period alone.

Others (EW, SR) contain the bulk of false predictions. Table 4.4 illustrates this dependence by comparing several machine learning performance metrics (introduced in section 1.6.3), separated by star class. The overall accuracy is on average acceptable, with a mean value of 0.823, although this metric varies; from good score of 0.876 for the EB class to 0.569 for the EW class. However, the accuracy alone is a poor descriptor of the performance of the algorithm; this can clearly be seen from the sensitivity of the EB class of 0.056, despite the good accuracy. This is because the accuracy also incorporates factors like the specificity, which is generally acceptable for this classifier, so this score can exaggerate the performance. For example, consider the sensitivity (or true positive rate; the fraction of positives which are correctly identified). The average score of this metric is just 0.290. In fact, the

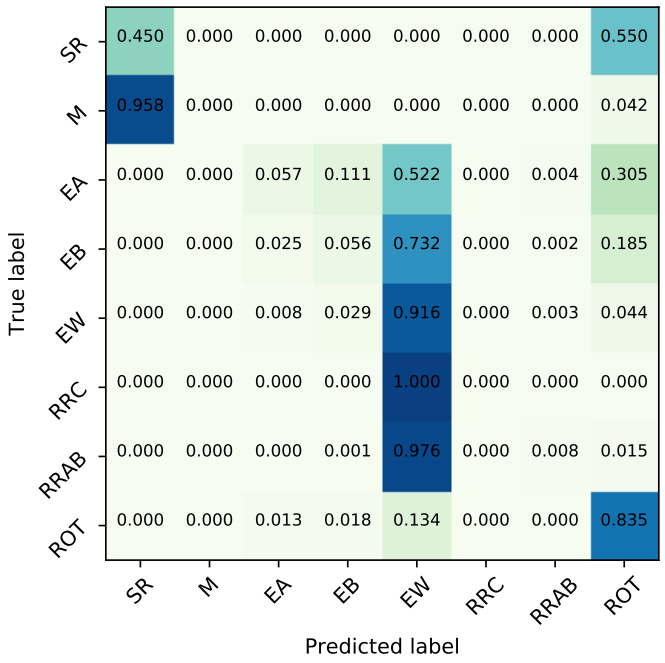

FIGURE 4.21: Confusion matrix for the Random Forest classifier. The matrix is normalised by true class, in rows.

sensitivity rarely crosses the threshold expected from random guessing (0.125). As such, this table and confusion matrix can instead give us significant insights into the shortcomings of the data and the classifier, which we will discuss in this section. We also constructed period-magnitude, period-colour, and colour-magnitude plots to illustrate the separation of classes in various parameter spaces, which we will introduce over the next few paragraphs. Since the classifier is dominated by the G band magnitude, $B_P - R_P$ colour, and period features, these plots are very useful for showing how class confusion occurs.

We will focus on the 3 star types which compose the vast majority (95.7%) of predictions; SR (17.6%), EW (53.4%), and ROT (24.7%).

Semi-regular variables With a sensitivity of 0.450, the classifier predictions for the SR class are split between correct classifications and misclassifications as ROTs. Conversely, the bulk (95.8%) of M class variables were classified as SRs, with the remaining 4.2% classified as ROTs. This results in a sensitivity of 0.0 and a false negative rate (FNR) of 1.0; exceptionally poor predictive power. Considering the plots shown in figures 4.22, 4.23, 4.24, we see a consistent and considerable overlap between the M-class and SR-class light curves. This indicates that, given these features alone, the classifier would be largely unable to differentiate the light curves as we see.

TABLE 4.4: Summary of results from classifier. Missing values are caused by some columns or rows summing to zero. All values have been rounded to 3 decimal places for display. The column headers are defined in section 1.6.3

Star Type	Sensitivity	Specificity	Precision	NPV	FPR	FNR	False discovery rate	Accuracy
SR	0.450	0.863	0.320	0.917	0.137	0.550	0.680	0.811
M	0.000	1.000	-	0.875	0.000	1.000	-	0.875
EA	0.057	0.993	0.554	0.881	0.007	0.943	0.447	0.876
EB	0.056	0.977	0.260	0.879	0.023	0.944	0.740	0.862
EW	0.915	0.519	0.214	0.977	0.481	0.084	0.786	0.569
RRC	0.000	1.000	-	0.875	0.000	1.000	-	0.875
RRAB	0.008	0.998	0.438	0.877	0.001	0.992	0.563	0.875
ROT	0.835	0.837	0.423	0.973	0.163	0.165	0.577	0.837
Mean	0.290	0.899	0.368	0.906	0.101	0.710	0.632	0.823

A similar observation can be made about the overlap of ROT light curves and SR light curves, although in these figures the distribution of ROT light curves is less coincident, most clearly so in the period-colour and colour-magnitude plots. Note however, that in all three figures the ROT and SR classes are generally well-separated from the other star types as they are in general brighter and have longer periods; this is reflected in the classifier as there is no confusion between these longer-period sources and the shorter-period eclipsing binaries and RR-Lyre-type objects. Still, these overlapping regions illustrate how reliance on such a small number of features leads to poor classification. We can surmise that the Fourier term features have little effect at resolving these issues.

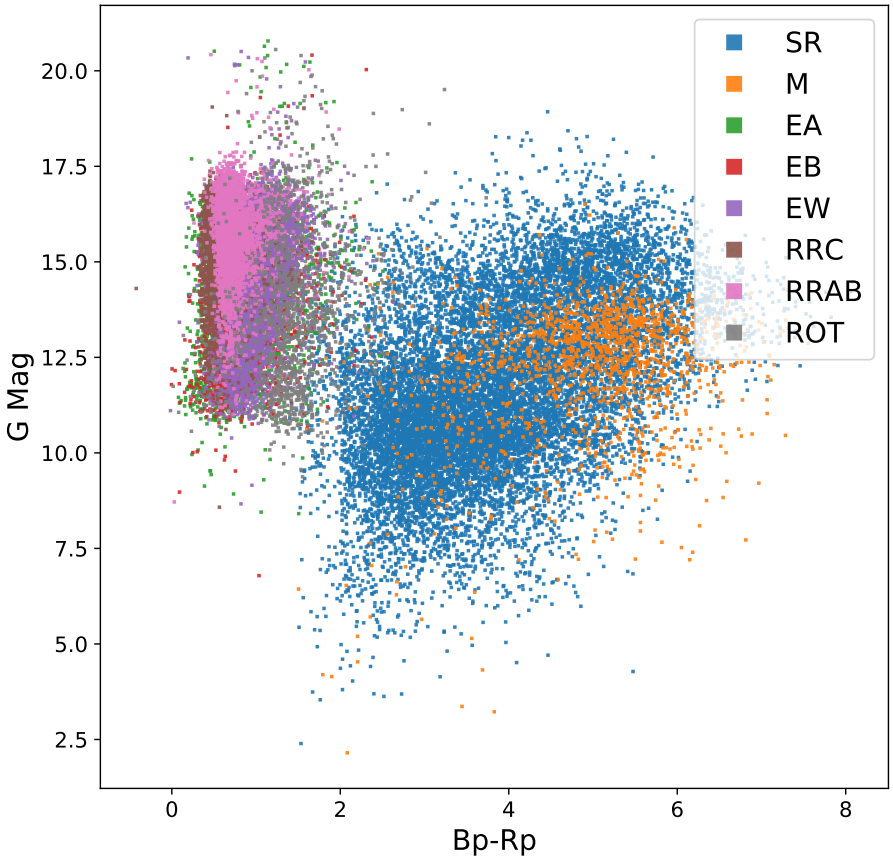

FIGURE 4.22: Colour-magnitude plot showing the locations of sources in the training set.

With regards to breaking this degeneracy, the introduction of additional colour terms can greatly help. Jayasinghe et al. (2019) show that the Wesenheit W_{JK} magnitude can help discriminate between ROT and SR. While the same cannot be done for Mira-type variables, this class generally has a higher modulation amplitude than SR which can help discriminate between the two. However, the issues we uncovered with GOTO light curves, uneven photometric calibration and frequent outliers, mean that this feature is less useful when classifying GOTO light curves.

Additionally, Mira variables can display period variations (Percy & Colivas, 1999) which may also contribute to this confusion.

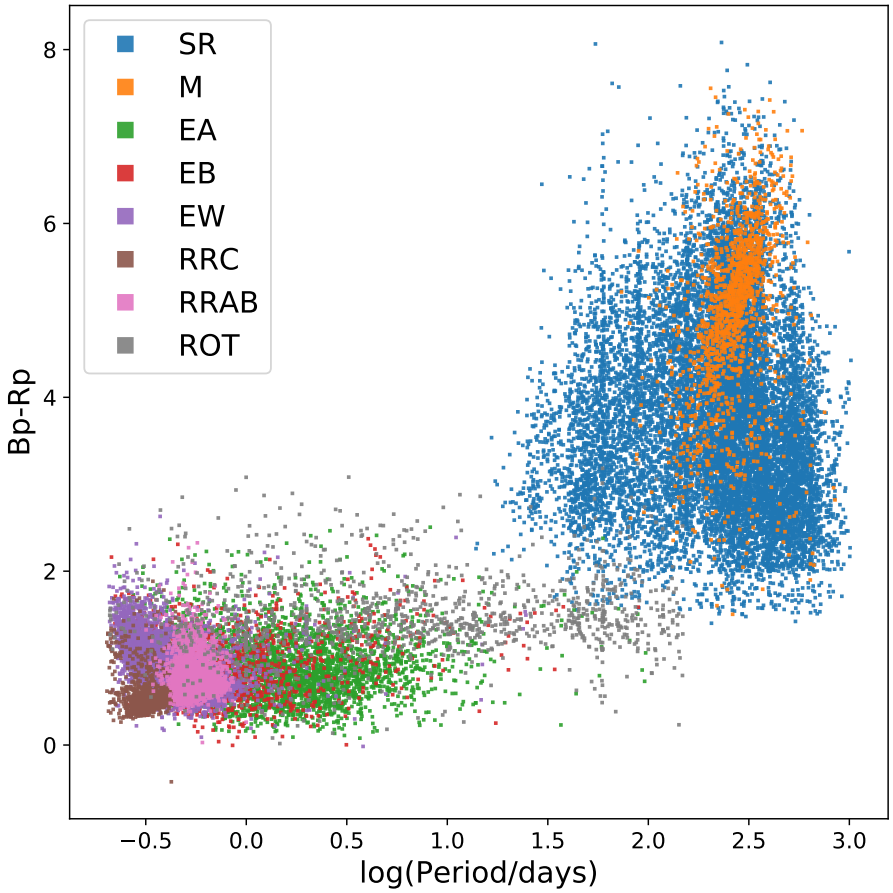

FIGURE 4.23: Colour-period plot showing the locations of sources in the training set.

W Ursae Majoris-type binaries Similarly to the Semi-regular variables, The EA, EB, EW, RRC, and RRAB classes all show significant overlap of the three major features. This is reflected in the classifier performance, where these classes are consistently misclassified as EW sources (though this also results in the majority of EW light curves being classified correctly). We obtain a precision of 0.214 for the EW class accordingly; the worst of all the classes. Figure 4.24 shows this especially well, where the RRC and RRAB light curves are completely coincident with the region occupied by the EW light curves. Again, Jayasinghe et al. (2019) makes it clear that the additional colour information afforded by the Wesenheit magnitudes breaks this degeneracy. What is not clear however, is why there is not more confusion between these classes instead of the dominance of the EW class as we see. This may simply be due to the greater extent of the EW light curves in each of the three parameter spaces, again a sign that the dominance of these three class features is a

serious problem, or it could be a further symptom of the GOTO light curves if the shape of the GOTO light curves is such that most resemble EW variables.

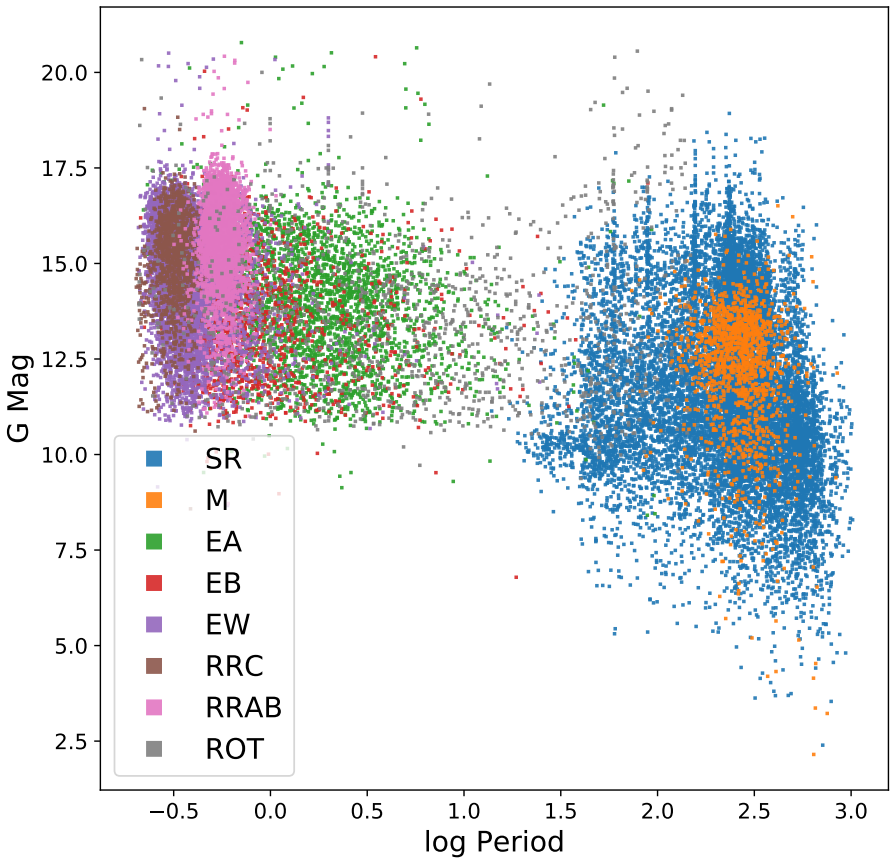

FIGURE 4.24: Period-magnitude plot showing the locations of sources in the training set.

Spotted variables This third case has the widest range of periods, magnitudes, and $B_P - R_P$ magnitudes, as shown in each of the three figures. These light curves overlap both the long-period region occupied by the SR and M variables, as well as the short-period region occupied by the other variables. As such, the fact that other classes are frequently classified as ROT variables is expected, resulting in the false positive rate of 0.568.

Absolute magnitude

Faced with this poor performance, we introduced the absolute magnitude as a further feature. The absolute magnitude is derived from the GAIA G-band magnitude and the GAIA parallax. We calculate an estimate of the source distance, d, in parsecs, from the parallax, p, in arcseconds, with the expression $d = 1/p$. Out of the 42834 sources with GAIA matches, we eliminate 5792 sources with negative or missing

parallax measurements. From the remaining 37042, we use the distance estimate to calculate the absolute magnitude, M, using the expression $M = m - 5\log_{10}(d) + 5$, where m is apparent G-band magnitude. Our new balanced testing set is composed of 6000 light curves (750 of each type).

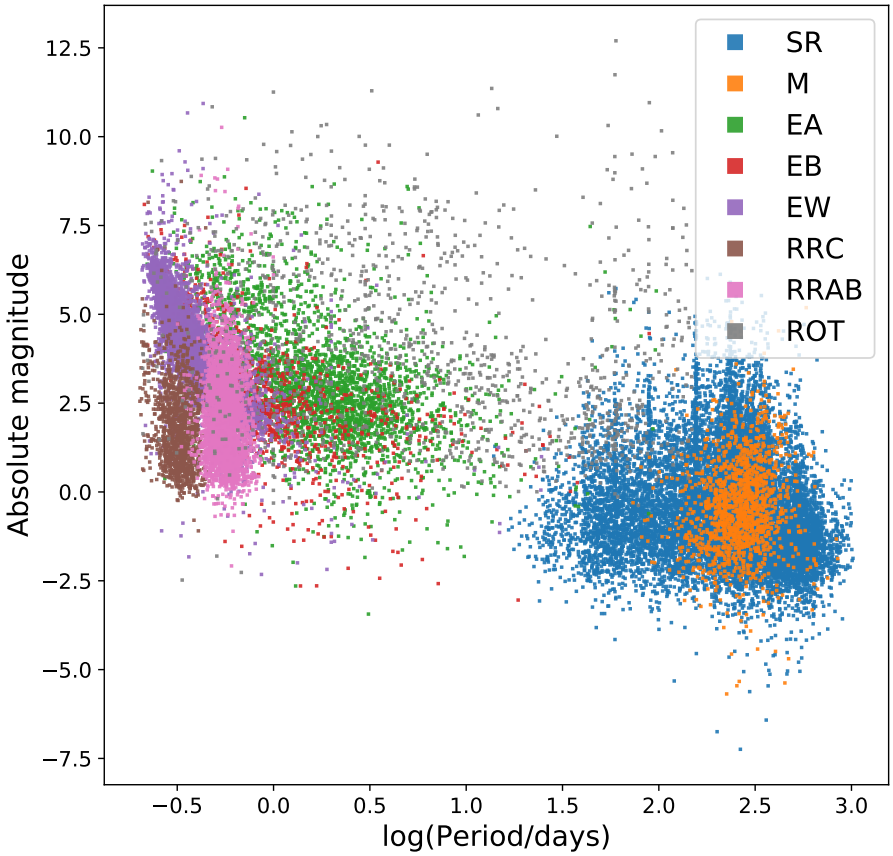

FIGURE 4.25: Period-absolute magnitude plot showing the locations of sources in the training set.

Figure 4.25 shows the period-luminosity diagram with the absolute magnitudes of the sources, illustrating that this feature results in a slight increase of the distinction between the short-period variables, however SR and Mira variables, as well as EA and EB variables, are still indistinguishable. Correspondingly, we found roughly identical performance with this additional feature, with minimal improvements (<5%) in differentiation between EW RRC, and RRAB classes. This suggests that the absolute magnitude it is not sufficient to break the degeneracy and indeed more colour terms are required with this dataset.

Removal of light curves

Visual inspection of the GOTO light curves had previously revealed significant scatter and anomalous data points. As such, we trained the classifier using only the

ancillary features - absolute magnitude, period, GAIA G magnitude, and GAIA $B_P - R_P$ colour - to determine the effect of the light curves on the classifier performance. We suspected that the poor quality of the light curves would have a large impact on the Fourier components; for example, excess noise may result in overestimation higher-order harmonics.

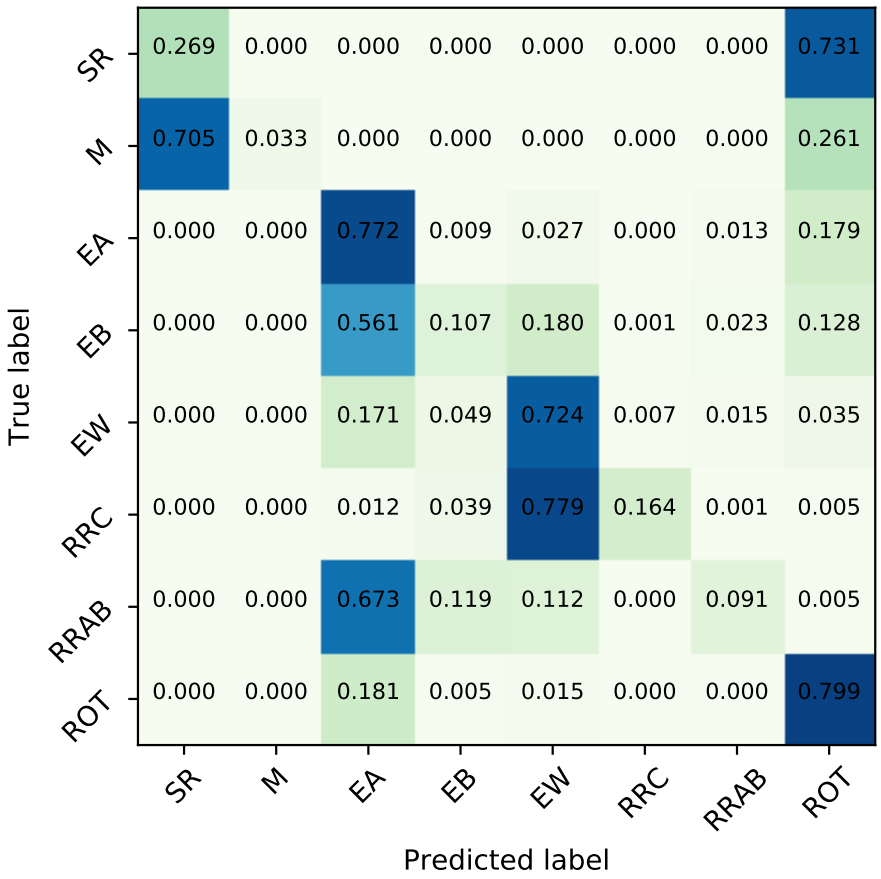

FIGURE 4.26: Normalised confusion matrix for classifier without light curves harmonics as features. Note that the diagonal terms are more dominant than the previous case, though the same confusion patterns (such as between SR and ROT) are still present

Figure 4.26 shows the confusion matrix of this classifier, tested on the second data set of 7500 sources. Note that the diagonal terms are more significant than with the original case, notably with the EA class. Additionally, fewer sources are classified as EW variables and indeed there is a wider distribution of predictions. This suggests that this problem was a symptom of the light curves

Some features remain from the original case, such as the confusion between ROT and SR, and SR and Mira variables, and frequent misclassification of short-period variables as EW variables. However, additional issues are introduced, such as the misclassification of RRAB and EB-type variables as EA. These misclassifications

are explained by the overlap of the populations in the parameter spaces shown in figures in this section.

However, that the removal of light curve harmonics as features does not have a significant detrimental effect on the classifier performance indicates that the quality of the GOTO light curves is inadequate for these purposes. This suggests that the outliers discussed in section 4.2.2 are still a considerable issue.

Classification of ASAS-SN light curves

To further illustrate the importance of colour information and high quality light curves, we also tested the classifier on a sample of ASAS-SN light curves. The additional colour information in the broader ASAS-SN database ($W_1 - W_2$, $H - K$, $J - K$, and $J - H$) were also included. Other than these additional colour terms and the the use of ASAS-SN rather than GOTO light curves, the classifier remained unchanged. The dataset used to train the classifier thus-far was shuffled and split into training and testing sets, with 33% of the 330372 light curves retained for testing.

The confusion matrix of this classifier is shown in figure 4.27; the mean sensitivity is 86.6%. Immediately the dominance of the diagonal terms can be seen, showing that the classifier performance is significantly improved over the instances using GOTO light curves. Six of the eight classes show sensitivities of over 90%, with the best performance in the SR class with 99.8% of SR variables classified correctly. The sensitivities of the remaining two classes, EB with 41.6% and ROT with 81.3%, indicate that some of the problems present in the other classifiers are still present. However, the overall accuracy of 98.5% clearly illustrates the vast improvements obtained by using full colour information and less noisy light curves.

The EB variables are frequently misclassified as EA and EW variables; similarly for ROT variables. Comparing the EB light curve in figure 4.29 with the EA and EW light curves in figures 4.28 and 4.30 respectively, the similarities can be seen. Notably, the presence of at least one sharp eclipse indicates that higher-order harmonics will contribute significantly in the Fourier decomposition. While this may differentiate these short-period variables from other classes, the similarity between them that results contributes to the misclassification that is observed. While the examples shown here indicate that the EB and EW classes are the most similar, the depth of the second eclipse in our ASAS-SN EB light curves varies significantly such that in many cases they appear more similar to EA light curves. As such we see that EB sources are frequently misclassified as both EA and EW variables. It should be noted that this confusion occurs much less frequently in reverse, implying that the populations of EA and EW sources are each relatively homogeneous when compared

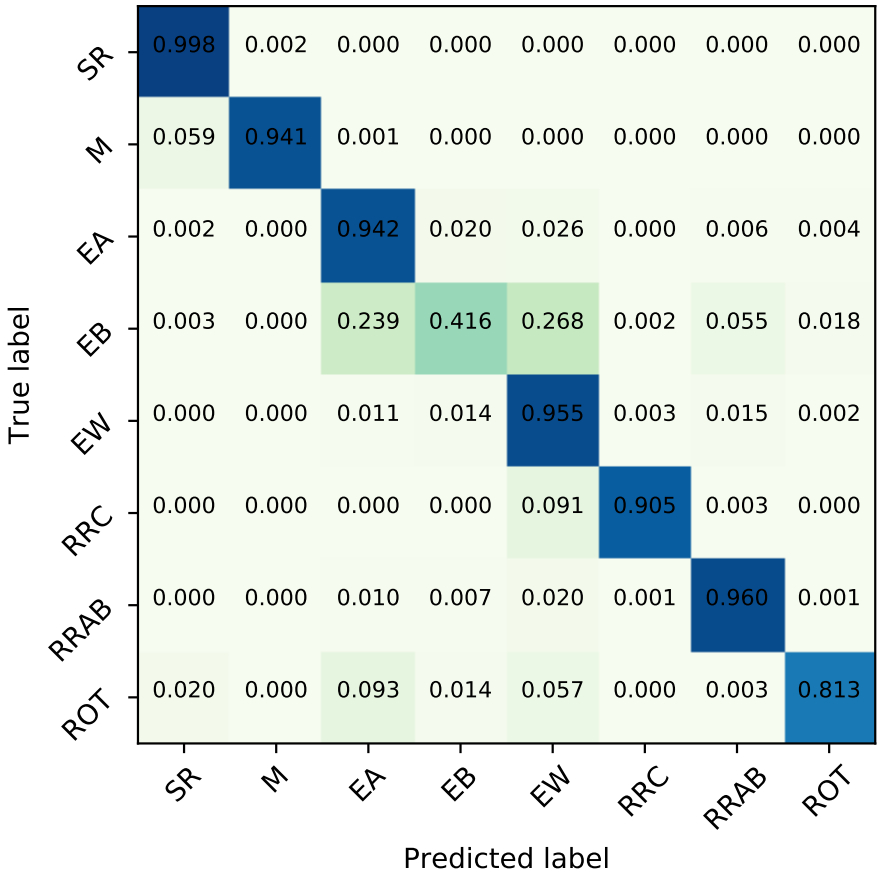

FIGURE 4.27: Normalised confusion matrix for classifier trained and tested on the ASAS-SN dataset, including full colour information and original ASAS-SN light curves. The diagonal terms dominate.

to the EB variables. With the ROT variables, the light curves population is again quite inconsistent. The light curve shown in figure 4.31 displays two minima, similar to the EA, EB, and EW variables, however most do not. As such, the majority are classified correctly with only a minority misclassified as eclipsing variables.

From this analysis we see the advantage that high quality light curves and auxiliary information provide. However, it should be noted that the ASAS-SN light curves, with a median of 201 and a range of $7 - 1122$ samples, generally consist of many more samples than GOTO light curves. Naturally, the period of a well-sampled light curve will be easier to recover, and this logic extends to the Fourier decomposition we perform. With a more accurate Fourier decomposition, the small differences between similar light curves are more prominent. Conversely, the Fourier decomposition of poorly-sampled light curves will produce components with significant uncertainties; in this case the predictive power of these components is drastically reduced. Indeed, removing the Fourier components of the GOTO light curves as classification features entirely does not drastically reduce the performance of the

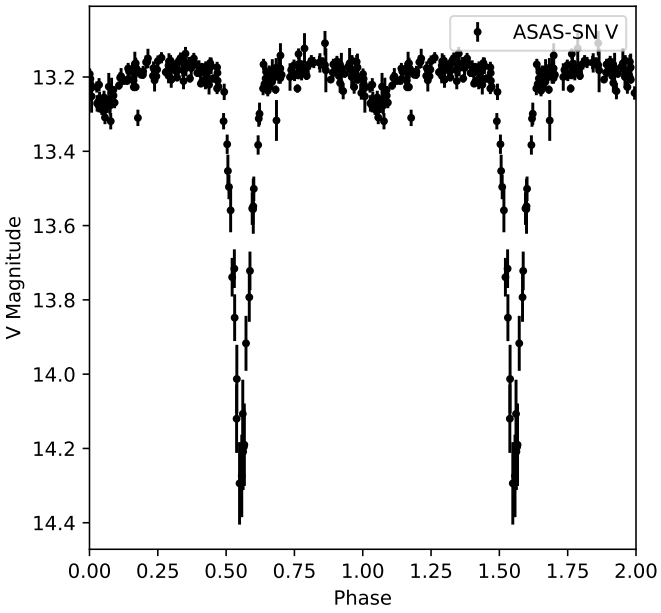

FIGURE 4.28: ASAS-SN V-band light curve of a randomly selected EA (Algol-type) variable. The characteristic single eclipse is located around phase 0.5-0.6, with a smaller secondary eclipse around phase 0.05.

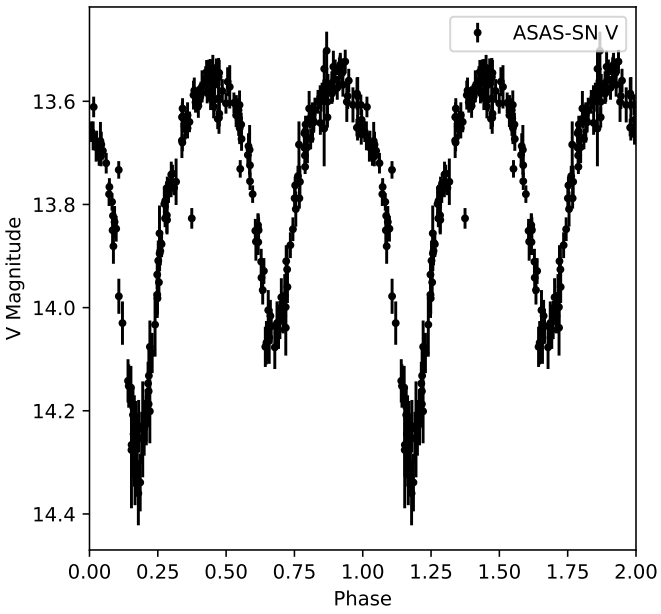

FIGURE 4.29: ASAS-SN V-band light curve of a randomly selected EB (β Lyrae-type) variable. The light curves in this sample display one main eclipse and a secondary whose amplitude varies significantly between sources.

classifier.

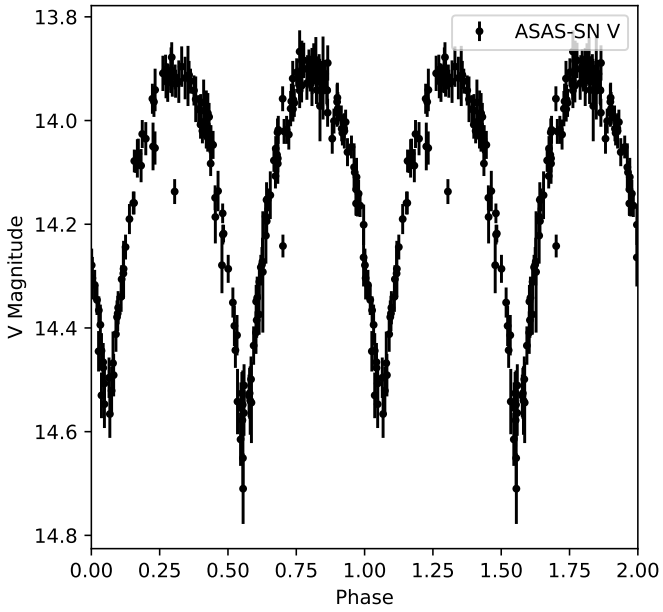

FIGURE 4.30: ASAS-SN V-band light curve of a randomly selected EW (W Ursae Majoris-type) variable. These light curves typically display two equally-spaced minima of similar amplitude.

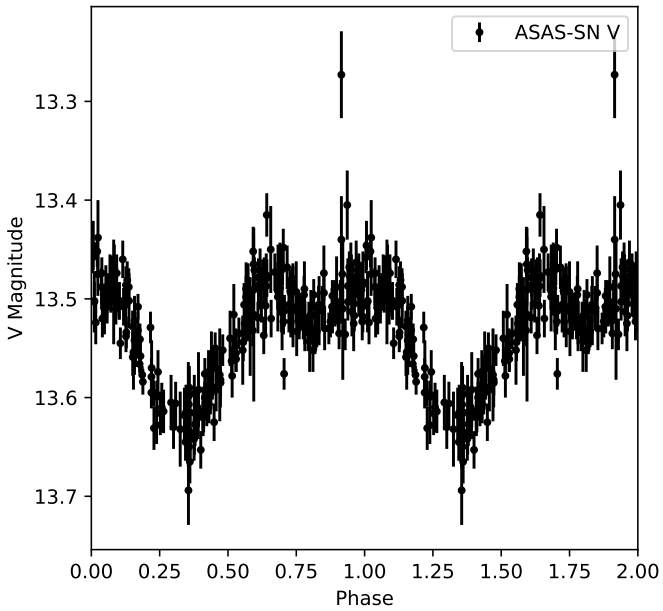

FIGURE 4.31: ASAS-SN V-band light curve of a randomly selected ROT (spotted variable) variable. The shape of these light curves varies significantly, due to the asymmetric natures of these sources. The light curve shown here, with two clear minima displays some similarities to the light curves of eclipsing variables.

Chapter 5

Discussion and conclusions

5.1 Discussion

5.1.1 tMSPs

Lessons for modelling

The photometry of PSR J1227–4853 (J1227) and PSR J1023+0038 was in both cases exceptionally high quality and precisely calibrated. Despite this, chapter 3 has shown that the modelling process is far from a straight shot to well-constrained system parameters. Instead, this work can give useful insights about the ultra-precise modelling required with such high quality data; especially important as data of this quality is becoming increasingly more common with the usage of instruments like HiPERCAM and ULTRACAM. The asymmetric nature of the optical light curves can no longer be hidden under observational noise or sparse sampling. In this section we will discuss the impact some of the small changes, such as fixing or constraining certain parameters, has on the model fit, as well as what the next steps for modelling these systems might be.

Intricacies of Icarus

Distance and extinction The importance of distance estimates in the modelling of spider binaries is clear. Without a measure of the distance the mass cannot be easily constrained (see, e.g., Thorstensen & Armstrong (2005a)). However, frequently overlooked is the extinction parameter, A_V, which incorporates the effect of interstellar reddening. Since this affects the colour of the object, i.e. the relative flux of each band, it can have a pronounced effect on both the temperature and the system distance. With J1023, where the distance was systematically underestimated, we hypothesised that an incorrect prior on the A_V extinction was responsible. However,

we found consistent results when fitting with two possible A_V priors. Further to this, we ensured that the extinction was correctly calibrated to each band to avoid any doubt about our distance estimates

We performed our fitting of J1023 with an especially small band calibration uncertainty of 0.005 mag as we calibrated the photometry with extreme care. This uncertainty accommodates for imprecisely calibrated magnitudes of each band of photometry by allowing for a small shift when fitting and can therefore introduce systematic errors into the distance estimates. However, we found consistent results even with larger band calibration uncertainties, suggesting that the calibration is optimal.

We also considered the potential effect of a third light, which may act to push the distance further away. However, this is not the only effect it would have on the system. For J1023 the ellipsoidal modulation is the most significant, which is modelled by i, q and f. The inclination, i, affects the geometry of the system, such that a more edge on system will display great modulation amplitude. Likewise, q and f define the size of the companion; a larger star will show ellipsoidal modulation of greater amplitude. We derive q from the companion radial velocity, K_2, which is well-constrained for J1023, leaving the free parameters i and f. When adding a third light to the system, any parameter changes can be absorbed by the inclination and filling factor, as well as the temperature. For instance, the increased flux could be represented by a larger star (i.e. f increases) if the irradiation temperature or inclination also adjust to keep the same ratio of irradiation to ellipsoidal modulation; adding a 3rd light dilutes the fractional variability of the total light. This means that to reproduce the same total fractional variability, a model including a 3rd light would need the intrinsic companion-only LC to have a larger fraction ellipsoidal amplitude. Thus, if q is fixed and the distance does not change, this means i moves closer to edge-on, or the filling factor increases.

In the end, this was roughly what we saw. By removing the expected flux from a third light, the filling factor decreased by around 5%, suggesting that this change was absorbed by changing the size of the star.

Temperatures We have shown that the result of the fitting is highly model-dependent, and that there is no one-size-fits all solution at this stage. Beyond just the goodness-of-fit, the choice of model can have a significant effect on the best-fit parameters, such as the pulsar mass estimates of J1227 or the filling factor and distance of J1023. This is due to the considerable degeneracy between parameters in the `Icarus` model; this degeneracy is normally broken by determining the temperature of the companion through modelling. By introducing model extensions which

alter the temperature distribution of the companion, either by adding flux in the form of a hot spot or by redistribution, we cause the degeneracy to be broken in a different way. This also extends to the different heat redistribution models, such as whether the effects of diffusion are included along with convection. It should be noted that these changes are small but significant. However, it should be noted that the black-body temperatures of the day and night side of the companion are comparable between models despite often significant changes in individual temperature parameters.

For instance, for J1023 we obtain a significantly larger irradiation temperature for the convection and diffusion model than with the hot spot model; $\sim$ 7350 K compared to $\sim$ 4700 K. With convection and diffusion, the heat is more efficiently redistributed over the star, lowering the actual day-side temperature. Conversely, the spot model has the extra contribution from the hot spot on the day side of the star, reducing the required irradiation for the same modulation amplitude. We see this with the broadly similar blackbody temperatures, however there are some consequences for the model. The heat diffusion allows for a larger ellipsoidal modulation with a smaller filling factor, $f \sim 0.81$ compared to $f \sim 0.94$. This is then compensated by bringing the system closer such that the correct flux is received. Indeed, the heat redistribution model with diffusion and convection returned a significantly smaller distance than any other model, and certainly significantly below the prior value. This may suggest that despite the physical motivation, the heat redistribution model is not an accurate description of the mechanism behind the asymmetry in this system.

We note a further consequence of this difference in irradiation temperature. While the temperature distribution of the companion star may not be significantly changed in this case, a larger irradiation temperature implies that the companion intercepts more energy from the pulsar. This could be caused by either a larger spin-down luminosity, $\dot{E}$, or an increase in the irradiation efficiency. However, because of the model dependency of the irradiation temperature we do not claim that these changes are indeed observed.

Spot constraints Initially when implementing the hot spot model we made no assumptions about the parameters (spot temperature, $T_{\rm spot}$, spot radius, $R_{\rm spot}$, and polar coordinates (θ, ϕ)) other than physical constrains like $T_{\rm spot}/K \geq 0$, $R_{\rm spot}/\deg \geq 0$. Our initial modelling revealed a strong preference for very small, very hot spots, with $T_{\rm spot} > 2000$ K and $R_{\rm spot} < 5^\circ$. We interpreted this as a sign of overfitting. While we might expect star spots to be small compared to the stellar radius, combined with the very high temperature this suggests that this combination is not

physical. For instance, it is unlikely that the heat would not dissipate within the convective layer of the star and instead remain contained in a tight spot without a significant magnetic field. We note that Kandel et al. (2020) describe hot spots at the magnetic pole of the companion, though these spots have a radius of around $30-45°$, suggesting that even under the influence of magnetic fields the spots are not that small. However, we had no physical constraints to apply to these parameters.

As such, we introduced the spot intensity, $I_{\rm spot} = T^4_{\rm spot} R^2_{\rm spot}$ with units $\mathrm{K}^4\ \mathrm{deg}^2$, as a further parameter which is analogous to the luminosity. This parameter is constrained by a Gaussian prior centred around zero with a large ($\sim 10^{10} - 10^{15}$) width; note that a change in this prior by a factor of 10^4 results in a change in temperature by a factor of 10. With this parameter, very hot spots (with necessarily small radii) are selected against. Despite this, a small, hot spot for J1227 remains, with the best-fitting spot temperature of 2100 K and radius of 7.8° with a spot prior width of 10^{12}. The reason for this is simply that the penalty for the spot is not sufficient to alter the fit, indicating a strong preference for this configuration.

This highlights one shortfall of modelling, in that the best fitting solution does not always correlate with physical expectations and priors. This is also seen in the consistent underestimation of the system distance of J1023, despite strong priors, or the mass estimates of J1227.

5.1.2 GOTO

Lessons for surveys

While the sky coverage, sensitivity, and resolution of GOTO mean it is excellently placed as a transient laboratory, the results discussed in chapter 4 show that there is some way to go before this performance extends to wider uses. We have seen that the photometry produced by GOTO is not sufficiently tailored for use in periodicity searching or classification of fainter objects, though Mong et al. (2020) present promising results on the classification of supernovae light curves. Naturally, the primary focus of GOTO should remain as a transient laboratory for the detection of gravitational wave counterparts, but how could the program be modified so that we may also obtain photometry useful for our aims?

The importance of reliable photometry is also key for period searching and classification, that is, consistent magnitude measurements and robust detection of cosmic rays. We found that, even after cleaning through two methods of clipping, the light curves displayed significant artefacts. These artefacts have a significant impact on both the periodicity searching and classification performance. For the periodicity search they increase the χ^2 statistic of harmonic series fits even at the true period,

reducing the power of true peaks in the periodogram compared to the background. For the classifier, the artefacts will remain after phase folding and so have an affect on the Fourier components when fit. It is likely that this is a contributing factor to the inability of the classifier to distinguish between classes with similar periods, magnitudes, and $B_P - R_P$ colours.

The key finding of this research, however, is is how large an effect the observing cadence has on the performance of periodicity search algorithms. We have shown that small changes to the observation schedule can result in a significant increase in periodicity search performance through the addition of random noise to observation times, without compromising the daily observations required of GOTO. Furthermore, the periodicity search performance is an important precursor to the classifier performance, due to both the phase folding of light curves as well as the period estimate itself being the dominant feature. Without concrete period estimates, we must rely on matching to existing surveys (such as ASAS-SN, as used in this research) for ephemerides. In this case, the fainter limiting of magnitude of ASAS-SN resulted in a significant fraction of GOTO light curves without matching ASAS-SN periods, nullifying the advantage of such a deep survey. It is at this point that the importance of ancillary information becomes clear.

The importance of ancillary information

In this research, we make use of two ancillary data sources; the GAIA database and the ASAS-SN catalogue of variable stars. The former provides the essential calibrated G-band magnitudes, distances, and $B_P - R_P$ colour information, while the latter provides periods, as well as our classifier training data set and labels for the testing data. In our case, the significance of the G-band magnitude and $B_P - R_P$ magnitude is clear given the importance of these classification features. Without these the classifier would be relying solely on the GOTO light curves and magnitudes, and the periods derived from these. Given the results we have obtained, including the challenges faced searching for periods and the inconsistent magnitude calibration of the GOTO L-band observations, this cross matching is essential. However, the process of cross-matching with other data sets can introduce a new set of problems, like the elimination of fainter GOTO sources without ASAS-SN matches.

A clear solution would be to use the multi-band imaging that GOTO is capable of. Not only will this provide the needed colour information for precise classification of variables, but will also allow the use of multi-band light curves for period classification. It is well-known that multi-band periodograms can outperform single-band periodograms (VanderPlas, 2018), especially when the signal is not strong in individual bands. Since GOTO is not capable of simultaneous multi-band imaging the

same way that ULTRACAM is, for example, the necessary staggering of observations in each band may also help alleviate the problems with aliasing.

Other surveys

We will now discuss some of the approaches taken by contemporaries of GOTO; other all sky or transient surveys. The most immediate comparison can be made with the BlackGEM project, as this telescope shares the aim of detecting gravitational wave counterpart transients. With the first phase installed at the La Silla observatory, Chile, in 2019, BlackGEM consists of three wide-field telescopes with 0.65 m primary mirrors, individually mounted (Bloemen et al., 2015). Each will have a field of view of 2.7 square degrees with a CCD resolution of 10k by 10k pixels for a pixel scale of 0.56 in/px (Roelfsema et al., 2016) . The resolution is seeing limited at 1 arcsec, and the sensitivity is expected to extend to 23rd magnitude in SDSS g band with a 5-minute exposure (Groot et al., 2019). Similar to GOTO, the primary function as a transient laboratory will be augmented by operating as a deep sky survey (though only in the southern sky) in SDSS ugriz filters, on similar timescales to GOTO of hours to days. The MeerLICHT telescope has identical specifications to BlackGEM, but instead shadows the MeerKAT radio telescope in Sutherland, South Africa (Bloemen et al., 2016).

A precursor to these projects is the SkyMapper telescope, and the associated Southern Sky Survey (Keller et al., 2007). SkyMapper is a single, larger, 1.3 m telescope with a 5.7 square degree field of view, with a similarly high CCD resolution of 16k by 16k pixels. This results in a pixel scale of 0.5 arcsec such that the telescope is also seeing-limited. This larger telescope is capable of photometry accuracy of better than 3% down to 23rd magnitude in SDSS g bands. Like BlackGEM it has interchangeable SDSS ugriz filters, with an additional v-band filter. SkyMapper is optimised for a range of stellar astrophysics science goals, such that it used multi-epoch sampling; from 4 hours through 1 day, week, month, and year. This range of epochs may help to alleviate aliasing issues if this survey is used for periodicity searching, though it is possible that the periodogram would be contaminated with even more aliased peaks. Akhter et al. (2013) successfully recover 90% known RR Lyrae variables in the globular cluster NGC 3201, though periodicity searching techniques were not used, instead by calculating the variability amplitude and colour information to identify RR Lyraes. However, this illustrates the sensitivity possible with SkyMapper.

Andreoni et al. (2020) present results from the Dark Energy Survey's Dark Energy Camera (DECam), which used continuous 20 s exposures with a magnitude limit of $g > 23$. This instrument is considerably more specialised than the others

discussed here; the project does not aim for full sky coverage, and the extremely short cadence is more suited to fast transient detection than detection of most variable stars.

The Vera Rubin Observatory (shorthand 'Rubin', formerly the Large Synoptic Survey Telescope), conducting the Legacy Survey of Space and Time (LSST) has perhaps the widest range of science goals yet; from mapping objects in the solar system, through high redshift galaxies and supernovae, to the distribution of galaxies and cosmological constraints. This large telescope has an effective aperture of 6.7 meters and a field of view of 9.6 square degrees with a total survey area of over 20,000 square degrees. Each pointing will be imaged 2000 times over the course of the 10 year program, with fifteen second exposures in six broad bands modelled on the SDSS ugriz(y) system, to a maximum depth of $r \sim 27.5$ (LSST Science Collaboration et al., 2009). From the real-time classification pipeline of Narayan et al. (2018), to promising recovery of simulated RR Lyrae light curve periods (Oluseyi et al., 2012), the breadth of this project and the volume of data have already produced interesting research which highlights the potential of such a system for the detection and classification of periodic variables. Though naturally, the size (and indeed budget) of the LSST or Rubin cannot be fairly compared to GOTO.

Lastly, the Zwicky Transient Facility (ZTF), the successor of the Palomar Transient Factory (PTF). The ZTF uses a 47 square degree wide-field camera on the Palomar 48 inch Schmidt telescope, representing more than an order of magnitude improvement in survey speed over the PTF (Bellm et al., 2018). It uses g, r, and i-band filters. The ZTF aims to discover new optical transients and variables through all-sky surveys, and as a result uses relative photometry as well as point source function (PSF) photometry for more accurate light curves.

While the science aims are often different, from the gravitational wave transient detection of BlackGEM and GOTO to the broad goals of the LSST, one of the key similarities between these surveys is the use of multi-band imaging. The advantages for classification of detections is clear, as colour information and its variation over time is one of the most important measures in optical astronomy. Another facet of surveys like the ZTF is the application of relative photometry correction (see Masci et al. (2019), ensuring more consistent calibration of light curves, which means identification of these light curves is more reliable.

GOTO sits among the smallest telescopes, smaller even than BlackGEM, however the field of view is one of the largest. The smaller collecting area means that GOTO is not as sensitive as its contemporaries; considering the faintness of black widow pulsars, especially at their optical minima, a limiting magnitude of ~ 20.5 is not

sufficient for reliable detections. In this regard, deeper surveys like the LSST may be better-suited for discovering the optical counterparts to spider pulsars.

Surveys like the ZTF and the LSST show that a primary aim of transient detection can be performed alongside standard photometry. While each takes a different approach, this shows that small adjustments to the GOTO program will provide bountiful returns for periodicity searching, classification, and other research with well-calibrated photometry.

5.2 Conclusions

The work in this thesis concerns two extremes of photometric data, from the precisely-calibrated, high time resolution photometry from ULTRACAM to the sparse, noisy data from GOTO. However, in both cases the work revealed significant intricacies.

For the modelling of ULTRACAM photometry in chapter 3 we see that the modelling was unable to fully account for the asymmetric shape of the light curve. While this asymmetry provides a substantial challenge for modelling, it also provides a window to new physics. With lower-quality data it is likely that some of this fine structure would be obscured by noise, but in this case it reveals that our current models are not a complete description of the system. However, we were able to obtain coarse estimates of the mass and inclination of both systems, as well as finding a system temperature that is consistent with other measurements. Notably, we find slightly under-filled Roche lobes in both systems, which may give insights into the nature of the transition mechanism; we expect the Roche lobes of tMSPs to be over-full in the AP state, suggesting that the star shrinks during the transition to the RP state. While these results are not as complete as one may hope, they encourage further study of these, and other, tMSP systems to further unveil their role in the evolution of spider pulsars.

While the periodicity search pipeline and ML classifier discussed in chapters 2 and 4 did not produce satisfactory results when implemented on GOTO data, they did highlight several key shortcomings of this photometry. This work is instead useful as a means of illustrating the desirable properties of photometric surveys - namely well-calibrated photometry, multi-band imaging, low aliasing, and good phase coverage - and showing the limitations which occur when some or all of these are missing. The importance of ancillary information has also been demonstrated. It should be noted that this work is focused on the study of short-period binaries, which are generally fainter than other transients like supernovae. Indeed, Mong et al. (2020) and Killestein et al. (2021) have shown very promising results in the detection and classification of brighter transients. However, if we wish to expand

these successes to the field of spider pulsars, then adjustments must be made in the scheduling of observations to reduce aliasing and additional colour information must be available. As well as this, we have shown that the quality of photometry must be improved when studying these faint sources, as the shape of the optical light curve is a key feature for differentiating between otherwise similar variables.

To summarise, this work typifies some of the key issues in modern astronomy; that the quality of the model and the data are rarely in tandem. However, we have shown that with some perseverance, interesting results can be obtained and where this is not possible provided a discussion on how best to proceed for future investigations.

Bibliography

Abdollahi S., et al., 2020, Astrophysical Journal, Supplement, 247, 33

Akhter S., Da Costa G. S., Keller S. C., Schmidt B. P., Bessell M. S., Tisserand P., 2013, Publications of the Astron. Soc. of Australia, 30, e051

Alpar M. A., Cheng A. F., Ruderman M. A., Shaham J., 1982, Nature, 300, 728

Andreoni I., et al., 2020, Monthly Notices of the RAS, 491, 5852

Applegate J. H., 1992, The Astrophysical Journal, 385, 621

Applegate J. H., Shaham J., 1994, The Astrophysical Journal, 436, 312

Archibald A. M., et al., 2009, Science, 324, 1411

Archibald A. M., Kaspi V. M., Hessels J. W. T., Stappers B., Janssen G., Lyne A., 2013, arXiv e-prints, p. arXiv:1311.5161

Arzoumanian Z., Fruchter A. S., Taylor J. H., 1994, The Astrophysical Journal, 426, L85

Asseo E., 1993, Monthly Notices of the Royal Astronomical Society, 264, 940

Baglio M. C., D'Avanzo P., Campana S., Zelati F. C., Covino S., Russell D. M., 2016, Astronomy & Astrophysics, 101, 1

Baluev R. V., 2013, Monthly Notices of the RAS, 429, 2052

Bassa C. G., et al., 2014a, Monthly Notices of the RAS, 441, 1825

Bassa C. G., et al., 2014b, Monthly Notices of the Royal Astronomical Society, 441, 1825

Becker A., 2015, HOTPANTS: High Order Transform of PSF ANd Template Subtraction (ascl:1504.004)

Bedding T. R., Zijlstra A. A., 1998, The Astrophysical Journal, 506, L47–L50

Bellm E. C., et al., 2018, Publications of the Astronomical Society of the Pacific, 131, 018002

Benvenuto O. G., De Vito M. A., Horvath J. E., 2014, The Astrophysical Journal Letters, 786

Bertin E., Arnouts S., 1996, Astronomy and Astrophysics, Supplement, 117, 393

Bhattacharya D., 2002, Journal of Astrophysics and Astronomy, 23, 67

Bhattacharya D., van den Heuvel E. P., 1991, Physics Reports, 203, 1

Bhattacharyya B., et al., 2013, The Astrophysical Journal Letters, 773, L12

Bingham E. A., Cacciari C., Dickens R. J., Pecci F. F., 1984, Monthly Notices of the RAS, 209, 765

Bloemen S., Groot P., Nelemans G., Klein-Wolt M., 2015, in Rucinski S. M., Torres G., Zejda M., eds, Astronomical Society of the Pacific Conference Series Vol. 496, Living Together: Planets, Host Stars and Binaries. p. 254

Bloemen S., et al., 2016, in Hall H. J., Gilmozzi R., Marshall H. K., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9906, Ground-based and Airborne Telescopes VI. p. 990664, doi:10.1117/12.2232522

Boden A. F., Lane B. F., 2001, Astrophysical Journal, 547, 1071

Bogdanov S., Patruno A., Archibald A. M., Bassa C., Hessels J. W. T., Janssen G. H., Stappers B. W., 2014, Astrophysical Journal, 789, 40

Bond H. E., White R. L., Becker R. H., O'Brien M. S., 2002, Publications of the ASP, 114, 1359

Bonnet-Bidaud J. M., de Martino D., Mouchet M., Falanga M., Belloni T., Masetti N., Mukai K., Matt G., 2012, Mem. Societa Astronomica Italiana, 83, 742

Breiman L., 2001, Machine Learning, 45, 5

Breton R. P., et al., 2008, Science, 321, 104

Breton R. P., Rappaport S. A., van Kerkwijk M. H., Carter J. A., 2012, The Astrophysical Journal, 748, 115

Breton R. P., et al., 2013, Astrophysical Journal, 769, 108

Britt C. T., Strader J., Chomiuk L., Tremou E., Peacock M., Halpern J., Salinas R., 2017, Astrophysical Journal, 849, 21

Buchner, J. et al., 2014, A&A, 564, A125

Bühlmann P., 2002, Statist. Sci., 17, 52

Cardelli J. A., Clayton G. C., Mathis J. S., 1989, Astrophysical Journal, 345, 245

Castelli F., Kurucz R. L., 2003, in Piskunov N., Weiss W. W., Gray D. F., eds, Vol. 210, Modelling of Stellar Atmospheres. p. A20 (`arXiv:astro-ph/0405087`)

Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2014, Physics of Plasmas, 21

Chakrabarty D., Morgan E. H., 1998, Nature, 348, 346

Chen A. Y., Beloborodov A. M., 2014, The Astrophysical Journal Letters, 795, L22

Chen K., Ruderman M., 1993, Astrophysical Journal, 402, 264

Chen W.-C., Li X.-D., Qian S.-B., 2006, Astrophysical Journal, 649, 973

Chen H.-L., Chen X., Tauris T. M., Han Z., 2013, The Astrophysical Journal, 775, 27

Cheng K. S., Ho C., Ruderman M., 1986, The Astrophysical Journal, 300, 500

Chinarova L. L., Andronov I. L., 2000, Odessa Astronomical Publications, 13, 116

Clark C. J., et al., 2021, Monthly Notices of the Royal Astronomical Society, 502, 915

Contopoulos I., Kazanas D., Fendt C., 1999, The Astrophysical Journal, 511, 351

Cordes J. M., 1978, The Astrophysical Journal, 222, 1006

Coti Zelati F., et al., 2019, Astronomy and Astrophysics, 622, A211

Cumming A., Marcy G. W., Butler R. P., 1999, The Astrophysical Journal, 526, 890

Cumming A., Zweibel E., Bildsten L., 2001, Astrophysical Journal, 557, 958

Dambis A. K., Berdnikov L. N., Kniazev A. Y., Kravtsov V. V., Rastorguev A. S., Sefako R., Vozyakova O. V., 2013, Monthly Notices of the Royal Astronomical Society, 435, 3206

De Falco V., Kuiper L., Bozzo E., Ferrigno C., Poutanen J., Stella L., Falanga M., 2017, Astronomy and Astrophysics, 603, A16

Deeming T. J., 1975, Astrophysics and Space Science, 36, 137

Deller A. T., et al., 2012, Astrophysical Journal, Letters, 756, L25

Deller A. T., Moldon J., Patruno A., Hessels J. W. T., Archibald A. M., Paragi Z., Heald G., Vilchez N., 2015, The Astrophysical Journal, 809, 13

Demory B.-O., et al., 2013, Astrophysical Journal, Letters, 776, L25

Dhillon V. S., et al., 2007, Monthly Notices of the RAS, 378, 825

Dhillon V. S., et al., 2016, Proc.SPIE 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 9908

Dhillon V., et al., 2018, in Evans C. J., Simard L., Takami H., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 10702, Ground-based and Airborne Instrumentation for Astronomy VII. p. 107020L (arXiv:1807.00557), doi:10.1117/12.2312041

Dyer M. J., et al., 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 114457G (arXiv:2012.02685), doi:10.1117/12.2561008

Eggleton P. P., 1983, Astrophysical Journal, 268, 368

Feast M. W., Glass I. S., Whitelock P. A., Catchpole R. M., 1989, Monthly Notices of the RAS, 241, 375

Feroz F., Hobson M. P., Bridges M., 2009, Monthly Notices of the Royal Astronomical Society, 398, 1601

Ferraro F. R., Possenti A., D'Amico N., Sabbi E., 2001, Astrophysical Journal, 561, L93

Fowler H. J., Kilsby C. G., 2003, International Journal of Climatology, 23, 1313

Fruchter A. S., Stinebring D. R., Taylor J. H., 1988, Nature, 333, 237

Fruchter A. S., et al., 1990, The Astrophysical Journal, 351, 642

Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, Astronomical Journal, 111, 1748

Gaia Collaboration et al., 2016, Astronomy and Astrophysics, 595, A1

Gaia Collaboration et al., 2018, Astronomy and Astrophysics, 616, A1

Giuricin G., Mardirossian F., 1981, Astrophysical Journal, Supplement, 46, 1

Glendenning N. K., 1992, Physical Review D, 46, 4161

Gold T., 1968, Nature, 218, 731

Goldreich P., Julian W. E. L., 1969, Astrophysical Journal, 157, 869

Gompertz B. P., et al., 2020, Monthly Notices of the RAS, 497, 726

Graham M. J., Drake A. J., Djorgovski S. G., Mahabal A. A., Donalek C., 2013a, Monthly Notices of the Royal Astronomical Society, 434, 2629

Graham M. J., Drake A. J., Djorgovski S. G., Mahabal A. A., Donalek C., Duan V., Maker A., 2013b, Monthly Notices of the Royal Astronomical Society, 434, 3423

Green G. M., et al., 2018, Monthly Notices of the RAS, 478, 651

Groot P., Bloemen S., Jonker P., 2019, in The La Silla Observatory - From the Inauguration to the Future. p. 33, doi:10.5281/zenodo.3471366

Gunn J. E., et al., 2006, The Astronomical Journal, 131, 2332

Hall D., Henry G., 1992, The Astronomical Journal, 104, 1936

Harding A. K., 1986, The Astrophysical Journal, 309, 362

Harding A. K., Muslimov A. G., 2002, The Astrophysical Journal, 568, 862

Harmanec P., Scholz G., 1993, Astronomy and Astrophysics, 279, 131

Helfand D. J., Gotthelf E. V., Halpern J. P., 2001, The Astrophysical Journal, 556, 380

Hernández Santisteban J. V., et al., 2016, Nature, 533, 366

Hessels J. W. T., Ransom S. M., Stairs I. H., Freire P. C. C., Kaspi V. M., Camilo F., 2006, Science, 311, 1901

Hester J. J., et al., 2002, The Astrophysical Journal, 577, L49

Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, Nature, 217, 709

Hinderer T., Lackey B. D., Lang R. N., Read J. S., 2010, Physical Review D, 81, 123016

Hobbs G., et al., 2010, Classical and Quantum Gravity, 27, 084013

Hoffman J., Vanderplas J., Hartman J., Bakos G., 2021, arXiv e-prints, p. arXiv:2101.12348

Honeycutt R. K., 1992, Publications of the ASP, 104, 435

Huijse P., Estévez P. A., Förster F., Daniel S. F., Connolly A. J., Protopapas P., Carrasco R., Príncipe J. C., 2018, The Astrophysical Journal Supplement Series, 236, 12

Hulse R. A., Taylor J. H., 1975, The Astrophysical Journal, 195, L51

Illarionov A. F., Sunyaev R. A., 1975, Astronomy and Astrophysics, 39, 185

Ireland M. J., Scholz M., Tuthill P. G., Wood P. R., 2004, Monthly Notices of the Royal Astronomical Society, 355, 444

Jackson J. D., 1962, Classical Electrodynamics. Wiley, New York

Jackson B., Adams E., Sandidge W., Kreyche S., Briggs J., 2019, Astronomical Journal, 157, 239

Jaodand A., Archibald A. M., Hessels J. W. T., Bogdanov S., D'Angelo C. R., Patruno A. r., Bassa C., Deller A. T., 2016, Astrophysical Journal, 830, 122

Jaodand A. D., et al., 2021, arXiv e-prints, p. arXiv:2102.13145

Jayasinghe T., et al., 2018, Monthly Notices of the Royal Astronomical Society, 477, 3145

Jayasinghe T., et al., 2019, Monthly Notices of the RAS, 486, 1907

Johnson T. J., et al., 2015, Astrophysical Journal, 806, 91

Kalapotharakos C., Contopoulos I., 2009, Astronomy & Astrophysics, 496, 495

Kandel D., Romani R. W., 2020, Astrophysical Journal, 892, 101

Kandel D., Romani R. W., Filippenko A. V., Brink T. G., Zheng W., 2020, arXiv e-prints, p. arXiv:2009.04105

Kaplan D. L., Bhalerao V. B., Van Kerkwijk M. H., Koester D., Kulkarni S. R., Stovall K., 2013, Astrophysical Journal, 765

Kaspi V. M., Helfand D. J., 2002, ASP Conference Series, 271, 3

Keane E. F., et al., 2018, Monthly Notices of the Royal Astronomical Society, 473, 116

Keith M. J., et al., 2011, Monthly Notices of the Royal Astronomical Society, 414, 1292

Keller S. C., et al., 2007, Publications of the Astron. Soc. of Australia, 24, 1

Kennedy M. R., Clark C. J., Voisin G., Breton R. P., 2018, Monthly Notices of the RAS, 477, 1120

Kennedy M. R., et al., 2020, Monthly Notices of the RAS, 494, 3912

Khechinashvili D., Melikidze G., Gil J., 2000, The Astrophysical Journal, 541, 335

Killestein T. L., et al., 2021, arXiv e-prints, p. arXiv:2102.09892

King A. R., 1988, Quarterly Journal of the RAS, 29, 1

Kippenhahn R., Weigert A., Weiss A., 2012, Stellar Structure and Evolution, doi:10.1007/978-3-642-30304-3.

Kirk J. G., Lyubarsky Y., Pétri J., 2009, in Becker W., ed., , Neutron Stars and Pulsars - Introduction. Springer, Berlin, Chapt. 16, pp 421–450

Kluzniak W., Ruderman M., Shaham J., Tavani M., 1988, Nature, 334, 225

Komacek T. D., Showman A. P., 2020, Astrophysical Journal, 888, 2

Komesaroff M. M., 1970, Nature, 225, 612

Konar S., Bhattacharya D., 1997, Monthly Notices of the RAS, 284, 311

Kramer M., Lange C., Lorimer D. R., Backer D. C., Xilouris K. M., Jessner A., Wielebinski R., 1999, The Astrophysical Journal, 526, 957

LSST Science Collaboration et al., 2009, arXiv e-prints, p. arXiv:0912.0201

Lang D., Hogg D. W., Mierle K., Blanton M., Roweis S., 2010, Astronomical Journal, 139, 1782

Lara F. E., Rieutord M., 2011, Astronomy & Astrophysics, 43, 1

Lattimer J. M., Prakash M., 2001, Astrophysical Journal, 550, 426

Lattimer J. M., Prakash M., 2004, Science, 304, 536

Lattimer J. M., Prakash M., 2007, Physics Reports, 442, 109

Lebzelter T., Kerschbaum F., Hron J., 1995, Astronomy and Astrophysics, 298, 159

Lebzelter T., Mowlavi N., Marigo P., Pastorelli G., Trabucchi M., Wood P. R., Lecoeur-Taïbi I., 2018, Astronomy and Astrophysics, 616, L13

Lesch H., Jessner A., Kramer M., Kunzl T., 1998, Astronomy & Astrophysics, 24, 21

Levin L., et al., 2013, Monthly Notices of the RAS, 434, 1387

Li M., Halpern J. P., Thorstensen J. R., 2014a, Astrophysical Journal, 795, 1

Li M., Halpern J. P., Thorstensen J. R., 2014b, Astrophysical Journal, 795, 115

Linares M., 2014, Astrophysical Journal, 795, 72

Linares M., 2018, Monthly Notices of the Royal Astronomical Society: Letters, 473, L50

Linares M., 2020, in Multifrequency Behaviour of High Energy Cosmic Sources - XIII. 3-8 June 2019. Palermo. p. 23 (arXiv:1910.09572)

Lomb N. R., 1976, Astrophysics and Space Science, 39, 447

Lorimer D. R., 2008, Living Reviews in Relativity, 11, 8

Lorimer D. R., Kramer M., 2004, Handbook of Pulsar Astronomy. Cambridge University Press, Cambridge

Lothringer J., Casewell S., 2020, in American Astronomical Society Meeting Abstracts #235. p. 409.07

Lucy L. B., 1967, Zeitschrift fuer Astrophysik, 65, 89

Lucy L. B., 1968, apj, 153, 877

Luri X., et al., 2018, Astronomy and Astrophysics, 616, A9

Lyapin A. R., Arkhangelskaja I. V., Larin D. S., 2017, J Phys. Conf. Ser., 798

Lyne A. G., et al., 1990, Nature, 347, 650

Ma B., Li X., 2009, Astrophysical Journal, 5, 1611

Machabeli G. Z., Usov V. V., 1979, Soviet Astronomy Letters, 5, 238

Madore B. F., 1982, Astrophysical Journal, 253, 575

Mahabal A., Sheth K., Gieseke F., Pai A., Djorgovski S. G., Drake A., Graham M., the CSS/CRTS/PTF Collaboration 2017, arXiv e-prints, p. arXiv:1709.06257

Manchester R. N., 2014, The Astronomical Journal, 129, 1993

Martin T. J., Davey S. C., 1995, Monthly Notices of the RAS, 275, 31

Masci F. J., et al., 2019, Publications of the ASP, 131, 018003

Masetti N., et al., 2006, Astronomy and Astrophysics, 459, 21

McConnell O., Callanan P. J., Kennedy M., Hurley D., Garnavich P., Menzies J., 2015, Monthly Notices of the RAS, 451, 3468

Melrose D. B., 1992, Coherent Radio Emission from Pulsars, 341, 105

Mestel L., 1972, Nature, 233, 149

Mestel L., Pryce M. H. L., 1992, Monthly Notices of the Royal Astronomical Society, 254, 355

Milne E. A., 1921, Monthly Notices of the Royal Astronomical Society, 81, 361

Mondrik N., Long J. P., Marshall J. L., 2015, Astrophysical Journal, Letters, 811, L34

Mong Y. L., et al., 2020, Monthly Notices of the RAS, 499, 6009

Narayan G., et al., 2018, Astrophysical Journal, Supplement, 236, 9

Nieder L., et al., 2020, Astrophysical Journal, Letters, 902, L46

Oluseyi H. M., et al., 2012, Astronomical Journal, 144, 9

Pacini F., 1967, Nature, 216, 567

Papitto A., de Martino D., 2020, arXiv e-prints, p. arXiv:2010.09060

Papitto A., et al., 2013, Nature, 501, 517

Papitto A., et al., 2018, Astrophysical Journal, Letters, 858, L12

Papitto A., et al., 2019, Astrophysical Journal, 882, 104

Patruno A., et al., 2014, Astrophysical Journal, Letters, 781, L3

Pedregosa F., et al., 2011, Journal of Machine Learning Research, 12, 2825

Percy J. R., Colivas T., 1999, Publications of the ASP, 111, 94

Phinney E. S., Evans C. R., Blandford R. D., Kulkarni S. R., 1988, Nature, 333, 832

Pletsch H. J., Clark C. J., 2015, The Astrophysical Journal, 807, 18

Podsiadlowski P., 1991, Nature, 350, 136

Polzin E. J., Breton R. P., Stappers B. W., Bhattacharyya B., Janssen G. H., Osłowski S., Roberts M. S. E., Sobey C., 2019, Monthly Notices of the RAS, 490, 889

Polzin E. J., Breton R. P., Bhattacharyya B., Scholte D., Sobey C., Stappers B. W., 2020, Monthly Notices of the RAS, 494, 2948

Pretorius M. L., 2009, Monthly Notices of the RAS, 395, 386

Pringle J. E., Rees M. J., 1972, Astronomy and Astrophysics, 21, 1

Pylyser E. H. P., Savonije G. J., 1989, Astronomy and Astrophysics, 208, 52

Radhakrishnan V., Cooke D. J., 1969, Astrophysical Letters, 3, 225

Rappaport S., Verbunt F., Joss P. C., 1983, Astrophysical Journal, 275, 713

Ray P. S., et al., 2012, arXiv e-prints, p. arXiv:1205.3089

Ray P. S., et al., 2014, in American Astronomical Society Meeting Abstracts #223. p. 140.07

Reimann J. D., 1994, PhD thesis, UNIVERSITY OF CALIFORNIA, BERKELEY.

Roberts M. S. E., 2011, in Burgay M., D'Amico N., Esposito P., Pellizzoni A., Possenti A., eds, American Institute of Physics Conference Series Vol. 1357, American Institute of Physics Conference Series. pp 127–130 (`arXiv:1103.0819`), doi:10.1063/1.3615095

Roelfsema R., et al., 2016, in Riva M., ed., Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 10012, Integrated Modeling of Complex Optomechanical Systems II. p. 1001206, doi:10.1117/12.2204767

Romani R. W., 1996, The Astrophysical Journal, 470, 469

Romani R. W., Sanchez N., 2016, Astrophysical Journal, 828, 7

Romani R. W., Graham M. L., Filippenko A. V., Zheng W., 2016, Astrophysical Journal, 833, 138

Romani R. W., Kandel D., Filippenko A. V., Brink T. G., Zheng W., 2021, arXiv e-prints, p. arXiv:2101.09822

Rowe E. T., 1995, Astronomy & Astrophysics, 296, 275

Roy J., et al., 2015a, apjl, 800, L12

Roy J., et al., 2015b, Astrophysical Journal, Letters, 800, L12

Salvetti D., et al., 2017, Monthly Notices of the Royal Astronomical Society, 470, 466

Sanpa-arsa S., 2016, PhD thesis, University of Virginia

Sazonov S. Y., Revnivtsev M. G., 2004, Astronomy and Astrophysics, 423, 469

Scargle J. D., 1982, Astrophysical Journal, 263, 835

Schaffenroth V., et al., 2019, Astronomy and Astrophysics, 630, A80

Schlafly E. F., Finkbeiner D. P., 2011, The Astrophysical Journal, 737, 103

Schroeder J., Halpern J., 2014, The Astrophysical Journal, 793, 78

Schuster A., 1898, Terrestrial Magnetism (Journal of Geophysical Research), 3, 13

Schwarz G., 1978, The Annals of Statistics, 6, 461

Schwarzenberg-Czerny A., 1996a, Astrophysical Journal, Letters, 460, L107

Schwarzenberg-Czerny A., 1996b, Astrophysical Journal, Letters, 460, L107

Seward F. D., Tucker W. H., Fesen R. A., 2006, The Astrophysical Journal, 652, 1277

Shahbaz T., et al., 2015, Monthly Notices of the RAS, 453, 3461

Shahbaz T., Linares M., Breton R. P., 2017, Monthly Notices of the RAS, 472, 4287

Shahbaz T., Dallilar Y., Garner A., Eikenberry S., Veledina A., Gandhi P., 2018, Monthly Notices of the RAS, 477, 566

Shahbaz T., Linares M., Rodríguez-Gil P., Casares J., 2019, Monthly Notices of the RAS, 488, 198

Shappee B. J., et al., 2014, The Astrophysical Journal, 788, 48

Shklovsky I. S., 1967, Astrophysical Journal, 148, 1

Skrutskie M. F., et al., 2006, Astronomical Journal, 131, 1163

Smith H. A., 2004, RR Lyrae Stars

Soberman G. E., Phinney E. S., van den Heuvel E. P. J., 1997, Astronomy and Astrophysics, 635, 620

Soker N., 2015, Astrophysical Journal, 800, 114

Sousa P. M., Trigo R. M., Aizpurua P., Nieto R., Gimeno L., Garcia-Herrera R., 2011, Natural Hazards and Earth System Sciences, 11, 33

Spitkovsky A., Arons J., 2002, ASP Conference Series, 271, 81

Stairs I. H., 2004, Science, 304, 547

Stappers B. W., Bailes M., Lyne A. G., Camilo F., Manchester R. N., Sandhu J. S., Toscano M., Bell J. F., 2001, Monthly Notices of the Royal Astronomical Society, 321, 576

Stappers B. W., et al., 2014a, Astrophysical Journal, 790, 39

Stappers B. W., et al., 2014b, Astrophysical Journal, 790, 39

Steeghs D., 2017, Nature Astronomy, 1, 741

Strader J., Li K.-L., Chomiuk L., Heinke C. O., Udalski A., Peacock M., Shishkovsky L., Tremou E., 2016b, Astrophysical Journal, 831, 89

Strader J., Li K., Chomiuk L., Heinke C. O., Udalski A., Peacock M., Shishkovsky L., Tremou E., 2016a, The Astrophysical Journal, 831, 89

Sturrock P. A., 1971, The Astrophysical Journal, 164, 529

Sutherland P. G., 1975, Astrophysical Journal, 196, 51

Suveges M., 2012, in Starck J.-L., Surace C., eds, Seventh Conference on Astronomical Data Analysis. p. 16

Szkody P., et al., 2003, Astronomical Journal, 126, 1499

Tan X., Showman A. P., 2020, Astrophysical Journal, 902, 27

Tan C. M., et al., 2018, Astrophysical Journal, 866, 54

Tauris T. M., 2015, arXiv e-prints, p. arXiv:1501.03882

Tauris T. M., Savonije G. J., 1999, Astronomy and Astrophysics, 350, 928

Tauris T. M., Savonije G. J., 2001, in Kouveliotou C., Ventura J., van den Heuvel E., eds, Vol. 567, The Neutron Star - Black Hole Connection. p. 337 (arXiv:astro-ph/0001014)

Tauris T. M., van den Heuvel E. P., 2006, in , Compact stellar X-ray sources. Cambridge University Press, Cambridge, pp 623–665

Tauris T. M., Langer N., Kramer M., 2012, Monthly Notices of the RAS, 425, 1601

Taylor J. H., Cordes J. M., 1993, The Astrophysical Journal, 411, 674

Terrell D., Gross J., Cooney W. R., 2012, The Astronomical Journal, 143, 99

Thompson C., Blandford R. D., Evans C. R., Phinney E. S., 1994, Astrophysical Journal, 422, 304

Thorstensen J. R., Armstrong E., 2005a, Astronomical Journal, 130, 759

Thorstensen J. R., Armstrong E., 2005b, Astronomical Journal, 130, 759

Tin Kam Ho 1998, IEEE Transactions on Pattern Analysis and Machine Intelligence, 20, 832

Tonry J. L., et al., 2018, Astrophysical Journal, 867, 105

Truemper J., Pietsch W., Reppin C., Voges W., Staubert R., Kendziorra E., 1978, apjl, 219, L105

VanderPlas J. T., 2018, The Astrophysical Journal Supplement Series, 236, 16

VanderPlas J. T., Ivezić Ž., 2015, Astrophysical Journal, 812, 18

Verbiest J. P. W., et al., 2008, Astrophysical Journal, 679, 675

Verbunt F., Phinney E. S., 1995, Astronomy & Astrophysics, 296, 709

Verbunt F., Zwaan C., 1981, Astronomy and Astrophysics, 100, L7

Voisin G., Kennedy M. R., Breton R. P., Clark C. J., Mata-Sánchez D., 2020, Monthly Notices of the RAS, 499, 1758

Wade R. A., Horne K., 1988, Astrophysical Journal, 324, 411

Wang Z., Archibald A. M., Thorstensen J. R., Kaspi V. M., Lorimer D. R., Stairs I., Ransom S. M., 2009, Astrophysical Journal, 703, 2017

Watson C. A., Steeghs D., Shahbaz T., Dhillon V. S., 2007, Monthly Notices of the RAS, 382, 1105

Weisberg J. M., Nice D. J., Taylor J. H., 2010, The Astrophysical Journal, 722, 1030

Wheaton W. A., Doty J. P., Primini F. A., Cooke B. A., Dobson C. A., 1979, Nature, 282, 240

Witte M. G., Savonije G. J., 1999, Astronomy & Astrophysics, 350, 129

Woudt P. A., Warner B., Pretorius M. L., 2004, Monthly Notices of the RAS, 351, 1015

Wright E. L., et al., 2010, Astronomical Journal, 140, 1868

Young M. D., Manchester R. N., Johnston S., 1999, Nature, 400, 848

Zechmeister M., Kürster M., 2009, A&A, 496

Zilles A., Kotera K., Rohrmann R., Althaus L., 2020, Monthly Notices of the RAS, 492, 1579

de Martino D., et al., 2010a, Astronomy and Astrophysics, 515, A25

de Martino D., et al., 2010b, Astronomy and Astrophysics, 515, A25

de Martino D., et al., 2013, Astronomy and Astrophysics, 550, A89

de Martino D., et al., 2014a, Monthly Notices of the Royal Astronomical Society, 444, 3004

de Martino D., et al., 2014b, Monthly Notices of the Royal Astronomical Society, 444, 3004

de Martino D., et al., 2015, Monthly Notices of the RAS, 454, 2190

de Wit J., Gillon M., Demory B. O., Seager S., 2012, Astronomy and Astrophysics, 548, A128

van Kerkwijk M. H., Breton R. P., Kulkarni S. R., 2011, The Astrophysical Journal, 728, 95

van den Heuvel E. P. J., van Paradijs J., 1988, Nature, 334, 227