Optimising the Detection of Temperate, Terrestrial Planets around Ultra-Cool Dwarfs with SPECULOOS

Catriona Anne Murray

Darwin College

UNIVERSITY OF CAMBRIDGE

2021 November

This thesis is submitted for the degree of Doctor of Philosophy

Declaration

This thesis, in its entirety, is the result of my own work, except where explicitly stated in the text and declared in the Preface. The appropriate references, citations, and acknowledgements for work done in collaboration have been included in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

The work in Chapter 2, *SPECULOOS (The Search for habitable Planets EClipsing ULtra-cOOl Stars)*, and parts of Chapter 3, *Assessing the Photometry and Performance of SPECULOOS-South*, were published in *Photometry and Performance of the SPECULOOS-South* in the Monthly Notices of the Royal Astronomical Society, volume 495 (Murray et al. 2020).

Chapter 4, *Stellar Variability in the Untapped Ultra-Cool Dwarf Regime*, was published in *A Study of Flares in the Ultra-Cool Regime with SPECULOOS-South* in the Monthly Notices of the Royal Astronomical Society, volume 513 (Murray et al. 2022).

The content of Chapter 5 will feature in a publication, currently in preparation, which evaluates the transit detection efficiency of SPECULOOS and constrains the occurrence rates of small, temperate planets orbiting ultra-cool dwarfs.

This dissertation comprises 60 figures (including 5 in appendices) and 39,671 words excluding the table of contents, figures and tables, figure captions, bibliography and acknowledgements. Therefore, this thesis does not exceed the 60,000 word limit set out by the Degree Committee of Physics and Chemistry.

Abstract

Optimising the Detection of Temperate, Terrestrial Planets around Ultra-Cool Dwarfs with SPECULOOS

Catriona Anne Murray

During my PhD, I worked as part of the SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars) team to optimise the detection potential of Earth-sized, habitable-zone planets transiting ultra-cool dwarfs (UCDs) from the ground. These cool, red dwarfs are small, abundant, and predicted to host a wealth of terrestrial planets. Within our current detection capabilities, these planets are also the most favourable candidates to search for traces of life. Our main goal is to provide the most promising candidates for future atmospheric characterisation by the James Webb Space Telescope (JWST). For this thesis, I focused on removing both the astrophysical and non-astrophysical sources that limit the detection of rocky planets and on constraining these planets' frequency.

Firstly, I developed the SSO Pipeline, a photometric pipeline specific for the SPECULOOS Project and the ultra-cool objects that it observes. This automated pipeline carries out image reduction, cross-matches with preexisting stellar catalogues, and performs precise aperture and differential photometry every night. With the SSO Pipeline, I address both the instrumental and atmospheric contamination that can prevent detecting small planets. I tackle the first and second-order effects of the Earth's rapidly-varying atmosphere on photometric observations through specialised photometric treatment. Specifically, this pipeline removes the photometric impact of water vapour absorption by implementing a novel, first-principles correction. I optimise the performance of this pipeline through rigorous quality checking and demonstrate SPECULOOS's unprecedented precision for UCDs. I showed that SPECULOOS is reaching its survey goals in regularly achieving the required precisions to detect temperate, terrestrial planets.

Secondly, due to the enhanced stellar activity of UCDs, accounting for photometric variability is an essential step in detecting rocky planets and in understanding whether these planets can initiate and sustain life. My work assesses whether this activity is beneficial or detrimental

to the habitability of planetary systems around UCDs, and probes their complex magnetic behaviour. In performing a flaring and rotation study of SPECULOOS's targets, I found that the activity of the coolest stars would be too low to initiate abiogenesis or to sustain an Earth-like biosphere on their planets. By studying the activity of these targets, I mitigate the astrophysical sources of contamination in our lightcurves.

Finally, I established bounds on the unexplored planetary populations around UCDs, by quantifying the detection potential of SPECULOOS for planets like TRAPPIST-1b. I developed a transit-search pipeline and performed transit injection-recovery tests to measure its detection efficiency. I concluded that worlds like TRAPPIST-1b, very short-period, rocky planets, are rare for these stellar hosts; otherwise, we would have already detected a small number. Several more years of SPECULOOS observations will be needed to confirm these results, and potentially, to find another TRAPPIST-1 system.

Contents

Declaration iii

Abstract v

Contents vii

List of Figures xi

List of Tables xiii

Acknowledgements xv

1 Exoplanets **1**
1.1 A Discovery of Exoplanets . . . 2
1.1.1 Techniques for Planet Detection . . . 2
1.2 The Transit Detection Method . . . 5
1.2.1 Outline . . . 5
1.2.2 Ancillary Products . . . 13
1.2.3 Searching for Transits from the Ground and the Sky . . . 15
1.2.4 Contamination of Planetary Signals from Non-Astrophysical Origins . . . 23
1.2.5 Contamination of Planetary Signals from Astrophysical Origins . . . 28
1.3 Stellar Classification . . . 29
1.3.1 M-dwarfs . . . 30
1.3.2 Ultra-Cool Dwarfs . . . 31
1.3.3 Brown Dwarfs . . . 31
1.4 Ultra-cool Dwarfs as Planetary Hosts . . . 31
1.4.1 Seizing the "M-dwarf Opportunity" . . . 34
1.4.2 The Planets Around our Coolest Neighbours . . . 35
1.4.3 Life around an Ultra-Cool Sun? . . . 37
1.5 Conclusions . . . 37

2 The Search for habitable Planets EClipsing ULtra-cOOl Stars (SPECULOOS) **39**
2.1 The SPECULOOS Project . . . 40
2.1.1 The TRAPPIST Prototype Survey . . . 40
2.1.2 SPECULOOS Global Network . . . 41
2.1.3 SPECULOOS Target List . . . 43
2.1.4 Observation Strategy . . . 45
2.1.5 Discoveries with SPECULOOS . . . 46

2.2 The SPECULOOS-South Pipeline . . . 48
2.2.1 Pipeline Design . . . 48
2.2.2 The ESO Archive . . . 50
2.2.3 Reduction . . . 50
2.2.4 Astrometry . . . 54
2.2.5 Generating a Stellar Reference Catalogue . . . 55
2.2.6 Aperture Photometry . . . 56
2.2.7 Differential Photometry . . . 58
2.3 Global vs. Nightly Lightcurves . . . 64
2.4 The Precipitable Water Vapour Effect . . . 66
2.4.1 Quantifying the PWV Effect . . . 70
2.4.2 Implementing the PWV Correction into the SSO Pipeline . . . 70
2.4.3 Impact of the PWV Correction . . . 71
2.4.4 Limitations . . . 72
2.5 Conclusions . . . 74

3 Assessing the Photometry and Performance of SPECULOOS-South 77
3.1 Global Observing Statistics from the SSO . . . 77
3.1.1 Targets Observed in the SSO Data Sample . . . 78
3.1.2 Cleaning the Lightcurves . . . 79
3.1.3 Blended Targets and Crowded Fields . . . 82
3.1.4 Quality Checks . . . 82
3.2 Photometric Precision of the SPECULOOS-South Observatory . . . 87
3.2.1 Calculating the Theoretical Noise . . . 87
3.2.2 Measured Photometric Precision . . . 90
3.3 Comparison with TESS . . . 94
3.4 Conclusions . . . 96

4 Stellar Variability in the Untapped Ultra-Cool Dwarf Regime 97
4.1 The Flares and Rotation of Ultra-Cool Dwarfs . . . 98
4.2 A Study of Flares . . . 99
4.2.1 Flare Detection . . . 99
4.2.2 Modelling Flares . . . 104
4.2.3 Extracting Flare Energies . . . 104
4.2.4 Calculating Flare Rates . . . 106
4.2.5 The Completeness of the Flare Sample . . . 106
4.3 Identifying Stellar Rotation Periods . . . 111
4.4 Combined Results from the SSO Flare and Rotation Study . . . 113
4.4.1 The SSO Flare Sample . . . 113
4.4.2 The Relationship between Rotation and Activity . . . 119
4.5 The Consequences for Life on Ultra-Cool Dwarf Planets . . . 123
4.5.1 Flare Frequency Distributions . . . 124
4.5.2 Converting Bolometric Flare Energies to U-band . . . 126
4.5.3 The Abiogenesis Zone . . . 127
4.5.4 Oxygenic Photosynthesis . . . 127
4.5.5 Ozone Depletion . . . 128
4.5.6 Applying the Sensitivity to the Flare Sample . . . 128

4.5.7 Habitability of Earth-sized planets around UCDs 133
4.6 The Evolution of Stellar Activity from Mid-M to Ultra-cool Dwarfs 135
4.7 Conclusions 137

5 The Frequency and Diversity of Temperate, Earth-sized Planets around UCDs 139
5.1 Gaussian Processes as a Tool for Removing Variability 140
5.1.1 What is a Gaussian Process? 140
5.1.2 GPs Applied to Exoplanet Problems 140
5.1.3 GP-Detrending of the SSO Variability 141
5.2 Searching for Transits Hidden in the SSO Lightcurves 142
5.2.1 Box Least Squares 144
5.2.2 Applying BLS to SPECULOOS 144
5.3 Transit Detection Efficiency 152
5.3.1 Injection-Recovery Tests 152
5.3.2 The Probability of Detecting a Planet 155
5.3.3 Results of the Injection-Recovery Tests 158
5.4 The Occurrence Rate of Planets Like TRAPPIST-1b 158
5.4.1 The Probability of Finding TRAPPIST-1b 163
5.4.2 Predictions about Future Detections 164
5.5 In Perspective 166
5.6 Future Directions 170

6 Conclusions and the Future 173
6.1 Summary 173
6.2 Looking Forward 176

A SSO Quality Checks 179

B Blended Targets 183

C SSO Data Sample 185

D Correcting Flare Occurrence Rates for Completeness 195

Funding and Data Acknowledgements 201

Bibliography 203

List of Figures

1.1 Orbital Geometry . . . 6
1.2 A Transit Lightcurve . . . 8
1.3 Brightness Variations during an Orbital Cycle . . . 15
1.4 Ground and Space-based Planet Detections . . . 17
1.5 Advantages of Low-Mass Stars for Transit Detection and Spectroscopy . . . 32

2.1 The SPECULOOS-South Observatory . . . 42
2.2 SSO Telescope . . . 43
2.3 SPECULOOS Transmission Curves . . . 45
2.4 Simplified flowchart of the SPECULOOS-South Pipeline. . . . 49
2.5 Master Calibration Images . . . 51
2.6 Reduction . . . 53
2.7 Differential Photometry . . . 59
2.8 Distance Weighting . . . 61
2.9 Magnitude-Dependent Correction . . . 63
2.10 $I + z'$-Magnitude against Gaia colour, $G - G_{\mathrm{RP}}$ . . . 65
2.11 Atmospheric Transmission Curves with Varying Water Vapour . . . 66
2.12 Precipitable Water Vapour in Paranal . . . 67
2.13 The Effect of PWV on Photometry . . . 68
2.14 Impact of PWV on a Nightly Lightcurve . . . 69
2.15 Predicted Frequency of PWV-induced Photometric Features . . . 71
2.16 Impact of PWV over Many Nights . . . 73

3.1 The SPECUULOOS-South Data Sample . . . 79
3.2 Observation Durations for SSO Targets . . . 80
3.3 Bad Weather Threshold . . . 81
3.4 Clean Lightcurve . . . 83
3.5 Monitoring of the Photometric Transform Between SPECULOOS and Gaia . . . 86
3.6 Theoretical Noise Model . . . 88
3.7 Photometric Precision of the SSO . . . 90
3.8 Photometric Precision of the SSO - Updated . . . 92
3.9 Photometric Quality Impact of PWV Correction . . . 93
3.10 Simultaneous Observations with TESS - A Photometric Comparison . . . 95

4.1 Flare Detection . . . 101
4.2 Recovered Amplitudes and FWHMs from Flare Injection-Recovery Tests . . . 108
4.3 Recovered Energies from Flare Injection-Recovery Tests . . . 109

4.4 Flare Sensitivity Map 110
4.5 Rotation Classes 112
4.6 The SSO Flare Sample 114
4.7 Flare Energies against Spectral Type 115
4.8 Spectral Distribution of Rotating Objects 120
4.9 Flaring on Fast Rotators 121
4.10 Comparison of Rotating Objects with MEarth 122
4.11 Flare Frequency Distributions 125
4.12 Average Flare Recovery Fractions for Spectral Type Bins 129
4.13 MCMC fit to Flare Occurrence Rates 130
4.14 Flare Occurrence Rates Corrected for Sensitivity 132

5.1 GP-Detrending of SSO Variability 143
5.2 Transit Candidate 1 in SSO Data Sample (Entire Lightcurve) 146
5.3 Transit Candidate 1 in SSO Data Sample 147
5.4 Transit Candidate 2 in SSO Data Sample 148
5.5 Transit Candidate 1 in SSO Data Sample 148
5.6 BLS-Recovered Transits from LP 791-18b and c 150
5.7 BLS-Recovered Transits from LHS 3844b 151
5.8 Signal-to-Noise Ratios of Recovered Planets 154
5.9 GP-Detrending and Transit-Search on a Non-Variable Object 156
5.10 GP-Detrending and Transit-Search on a Variable Object 157
5.11 Transit Detection Sensitivity 159
5.12 All Injected Planets 160
5.13 Smoothed Recovery Grid 161
5.14 The Expected Numbers of Planets Detected if *r=100%* 162
5.15 Occurrence Rate of Planets Similar to TRAPPIST-1b 165
5.16 Finding Another TRAPPIST-1b 166
5.17 Mulders (2018) Review of Planet Occurrence Rates 168

A.1 Readout Noise 180
A.2 Dark Current 181

B.1 Blended Target 183
B.2 Lightcurve from a Blended Target 184

D.1 Flare Recovery Fractions for Stars in Different Spectral Type Bins 196
D.2 MCMC fit to Flare Occurrence Rates in Different Spectral Type Bins 196

List of Tables

1.1 Table of the 12 *bona fide* planets around confirmed ultra-cool dwarfs. 35

2.1 SPECULOOS-South Telescope Specifications 44

4.1 Rotation Class Definitions 113
4.2 Overlap of Flaring and Rotating Objects in the TESS (Seli et al. 2021) and SPECULOOS Surveys 118
4.3 Overlap of Rotating Objects in the MEarth (Newton et al. 2018) and SSO Surveys . 123
4.4 Observed Flare Occurrence Rates for Spectral Type and Flare Energy Bins 131
4.5 α for Spectral Type and Flare Energy Bins 132

C.1 The SSO Data Sample 194

Acknowledgements

First and foremost, I would like to thank my supervisor, Didier Queloz, for his guidance and mentorship over the past four years. I thank him for this PhD project, which has provided countless opportunities to present, publish and learn. The encouragement to work independently and creatively has taught me a considerable amount, but most importantly, curiosity.

My thanks also go to Laetitia Delrez. Her patience for my constant interruptions and questions, and her careful attention to detail, are unmatched. She has been a great support and source of advice and remains so even after moving back to Belgium.

I would like to thank the SPECULOOS team. There have been many interesting and thought-provoking discussions, from the weekly teleconferences to the whole consortium meetings. I have felt nothing but genuine warmth, encouragement and collaboration in these meetings. Involvement in the SPECULOOS project has opened many doors and allowed me to travel to the telescopes in Chile a few times, each a wonderful experience. It has been a privilege to work with the SPECULOOS team, and I look forward to reading all the exciting science they will produce in the future.

Next, I thank the Cavendish Astrophysics and Exoplanet groups for broadening my perspective through their weekly coffee talks and seminars (even if I needed to be lured in by free doughnuts). Thanks particularly to Greg Willatt and Dave Green for their continual help throughout this PhD and everyone I have shared an office with over these past four years. A special mention goes to my office mates and friends, Josh and Patrick, with whom I have shared a staggering eight years of Cambridge experiences.

Lastly, I humbly thank my close friends and family. Finishing a PhD during the pandemic has been tough, and so to my housemates, Akash, Katie and Anna, with whom I shared multiple lockdowns, I owe a great thanks. They never failed to inject fun, joy and warmth into our house to make the experience not only bearable but happy. I thank Sarah and Robyn, whose friendships span a decade and, despite us living far apart, I know they will last many more. Finally, I thank my parents, the coolest and kindest people I know. They have provided endless (financial and emotional) support and advice throughout this journey. I will forever be grateful to them for encouraging me to pursue whichever path I wanted.

“There is nothing like looking, if you want to find something. You certainly usually find something, if you look, but it is not always quite the something you were after."

— J.R.R. Tolkien, *The Hobbit*

CHAPTER 1

EXOPLANETS

"This space we declare to be infinite... In it are an infinity of worlds of the same kind as our own."
— Giordano Bruno, *De l'infinito universo et mondi (1584)*

In this chapter, I outline the history and landscape of the exoplanet field, upon which my research is built. Firstly, I provide a brief history, followed by an introduction to the various techniques used to find exoplanets, focusing particularly on the transit detection method. Then, I trace the evolution of transit surveys, on the ground and in space, that has taken us from the first discovery of a transiting extra-solar planet to the thousands of planetary detections we have today. I include a discussion of the many sources of instrumental, astrophysical and atmospheric contamination these surveys face. Finally, I consider the core motivation of exoplanetary science: to find a temperate, Earth-like planet capable of supporting life. For this purpose, I examine the case of ultra-cool dwarfs as planetary hosts.

The numbers of planets quoted throughout this chapter have been taken from the NASA Exoplanet Archive[a]. There are several exoplanet databases, each with its own criteria for the planets they include. The NASA Exoplanet Archive has only non-free-floating planets with masses less than or equal to 30 Jupiter masses. Inclusion in that catalogue also requires sufficient follow-up or vetting to ensure a low chance of false alarms and is restricted to announcements made in peer-reviewed publications. Consequently, the NASA archive often quotes lower numbers of confirmed planets than other similar lists. I chose to use this catalogue due to the

[a]https://exoplanetarchive.ipac.caltech.edu/

rigour of their selection process, but with the acknowledgement that it may represent a lower bound on real exoplanet populations.

1.1 A Discovery of Exoplanets

Although the existence of undiscovered worlds outside of the Solar System (also known as exoplanets) has been predicted for centuries, the study of exoplanets is a deceptively young field. As early as the 16$^{\text{th}}$ Century, with the advent of the Copernican theory that the Earth and other planets orbited the Sun, philosophers began to wonder whether there could be planetary systems like our own around the other stars in the sky.

However, only in the past thirty years have we empirically tested the theory that exoplanets exist. The first exoplanet was discovered in 1992 around a millisecond pulsar (Wolszczan & Frail 1992), followed in 1995 by the detection of the first planet to orbit a main sequence star (Mayor & Queloz 1995). Mayor & Queloz (1995) found that this Jupiter-sized planet orbited the nearby star 51 Pegasi once every four days. With these early discoveries (that looked strikingly different to our own Solar System), a new era of exoplanetary science began.

The core motivation behind the exoplanet field is to provide answers to fundamental questions about our origins and our place in the universe. *Are we alone? How did life originate here on Earth? Is the Solar System unique?* These are not new questions, but it is only very recently that we have been able to address them. There are now over 4500 confirmed exoplanets in over 3300 planetary systems, none of which look exactly like the Solar System. We no longer think that all other planetary systems would mirror our own; instead, they exhibit a huge diversity in physical, orbital and atmospheric properties. In the thirty years since the first extra-solar discovery of a planet, instrumental limitations have meant that the sought-after “Earth-twin” has remained elusive. However, as we herald in a new age of exoplanet surveys, with cutting-edge detectors and unprecedented instrumental precisions, the search for answers has never been more promising.

1.1.1 Techniques for Planet Detection

There are several methods, both direct and indirect, which astronomers have successfully used to detect exoplanets. The two most prolific detection techniques are the transit and the radial velocity methods. However, for completeness, I will also very briefly describe three other methods: direct imaging, astrometry and microlensing. Each detection technique has its own biases, limitations, and risks of false positives. Each one can also uncover a different piece of the planetary puzzle. Individually, these methods can tell us about a planet’s formation, evolution, composition, orbital dynamics, atmospheric physics, or about the architecture of a multi-planetary system. Nevertheless, to gain a full, unified understanding of a planet and its environment, we would need to use a combination of detection techniques. If the reader is interested, Exoplanets (Seager 2010) and the Exoplanet Handbook (Perryman 2018) provide a full description of each of these techniques.

Radial Velocity

When we think about a planetary system, we often think about the gravitational influence exerted by a much larger body, such as a star, onto something much smaller, such as a planet. In reality, both the planet and the star will orbit separately around the barycentre (a common centre-of-mass). Because of this, we can use the periodic "wiggle" of the star moving around the barycentre to detect whether it has any orbiting bodies and how massive they are. Indirectly observing the star moving towards and away from us, through the Doppler shift of its stellar spectrum (the intensity of light as a function of wavelength), is known as the Radial Velocity technique.

The larger the mass ratio between planet and star, the greater the radial velocity of the star (Lovis & Fischer 2010). Therefore, this technique can constrain the planet's mass (if the stellar mass is well-known) and is biased towards detecting heavier planets. Notably, however, only a lower limit for the mass is yielded if we do not know the inclination of the orbital plane with respect to the observer. The closer the planet is to the star, the more frequently the star "wiggles", which is why it has been utilised so successfully to detect so-called "hot Jupiters", gas giant planets at very short orbital distances (e.g., Mayor & Queloz 1995; Bouchy et al. 2005; Quinn et al. 2012). The vast majority of radial velocity surveys employ a targeted approach. Traditional spectrometers can only measure a stellar spectrum from one star at a time. Therefore, while it has not historically been possible to perform a wide-field survey of multiple stars at once (as we can with the transit method), targets do not need to be monitored continuously. Further specifics and a review of the almost 30-year history of the radial velocity detection method are presented in Lovis & Fischer (2010).

Direct Imaging

As implied from the name, direct imaging is the process of capturing light from exoplanets, enabling us to "see" them. The light can either be reflected or re-emitted light from a planet's host star or generated internally through the planet's thermal emission. While straightforward in concept, this method is extremely challenging in practice due to the faintness of planets compared to the influx of stellar light and the relative closeness of planets to their stars (Claudi 2016; Bowler 2016). These difficulties have limited the use of direct imaging to predominantly detect massive planets or planets on extremely wide orbits (Chauvin 2018). Significant instrumental advances have helped make high-resolution direct imaging possible. Adaptive optics systems on ground-based telescopes mitigate difficulties in observing through the Earth's atmosphere, and coronographs can block the light from bright stars to more easily see faint planets. To date, the NASA Exoplanet Archive states that at least 50 planets have been directly imaged (e.g., Kuzuhara et al. 2013; Macintosh et al. 2015; Chauvin et al. 2017). For a more in-depth discussion of the direct imaging method, see the book chapter by Traub & Oppenheimer (2010), or the review paper by Chauvin (2018) for a more recent description of the field.

Astrometry

As with the radial velocity technique, astrometry relies on the fact that stars that host massive bodies will move around the common barycentre. However, in this method we track these motions directly, through measuring the precise positions and movements of observable bodies in the sky. This technique is described in more detail in Sozzetti (2005) and Quirrenbach (2010). In contrast to the radial velocity and transit techniques, astrometry is more sensitive to long period planets, which cause greater stellar displacement (Quirrenbach 2010). Astrometry also measures two components of orbital motion, unlike with radial velocity; therefore, the mass can be determined directly without knowing the orbital inclination. These factors make astrometry an attractive detection technique. Nevertheless, due to the extremely high precisions needed to detect a planet through astrometry (Sozzetti 2005), the planets found with this method do not often hold up to scrutiny, e.g. DENIS-P J082303.1-491201 b (Sahlmann et al. 2013) which exists on the boundary between giant planet and brown dwarf (see Section 1.3.3). The space-based astrometry mission, Gaia (Gaia Collaboration et al. 2016a), launched in 2013, revolutionised this method. Gaia aims to map the positions and kinematics of around a billion light sources with extraordinary accuracy. This accounts for around 1% of the Galactic stellar population.

Microlensing

This method is based on the gravitational lens effect. A large mass will distort spacetime so that light will bend around it. In this way, the gravitational field of an object in the foreground can act as a lens for distant light sources (such as stars) in the background along our line of sight. Gravitational microlensing uses the same principle on a much smaller scale (Liebes 1964; Refsdal & Bondi 1964; Paczynski 1986). Here the mass of the lens object is too small to clearly distort the image, but the lensing effect can be "seen" indirectly by a brightening of the light source due to the relative motion of the foreground and background objects. A planet orbiting the lens star can add an additional lens to the system, imprinting a detectable signature if the alignment happens to be just right (Mao et al. 1991; Gould & Loeb 1992). With this, we can derive the planetary mass and orbital separation.

The precise alignment needed for microlensing to occur is very rare, and the planet's signature is short-lived, meaning that this detection technique is heavily time-intensive. It is also a random, unique event, so the exoplanet cannot be re-detected. However, this method's advantage is that it is sensitive to planets of a few Earth-masses (Beaulieu et al. 2006) at large physical separations from their host star (Gaudi 2010). These planetary systems can be used to constrain planetary formation models, as they exist beyond the "snow line" where volatiles freeze out. There are several excellent reviews of this detection method, such as those by Paczyński (1996), Sackett (1998) and Gaudi (2010).

1.2 The Transit Detection Method

1.2.1 Outline

We will use Kepler's three laws of planetary motion for this section. The first law, the "Law of Ellipses", states that every planet's orbit around its star is an ellipse with the common centre of mass at one focal point (the star moves on a much smaller elliptical path around the barycentre with the same orbital period as the planet). The second law, the "Law of Equal Areas", states that a line drawn from a focal point to the planet's centre will sweep out equal areas in equal time intervals. Finally, the third law, or the "Law of Harmonies", states that all planets a star hosts have the same ratio P^2/a^3, where P is the orbital period, and a is the semi-major axis (half the longest diameter of a planet's orbit).

Orbital Geometry and the Probability of Transit

To derive the orbital geometry of a planetary system, we imagine the scenario of a planet of radius R_p and mass M_p orbiting a star of radius R_* and mass M_*. The distance between the planet and the star at a given point in the orbit, r, is

$$r = \frac{a(1-e^2)}{1+e\cos f}, \tag{1.1}$$

where e is the eccentricity of the orbit (a number between 0 and 1 that describes an orbit's ellipticity, where $e = 0$ is a circular orbit) and f is defined as the true anomaly. The true anomaly is a function of time that depends on the eccentricity and period and can be pictured as the angle between the periapse (shortest planet-star distance) and a line drawn from the star

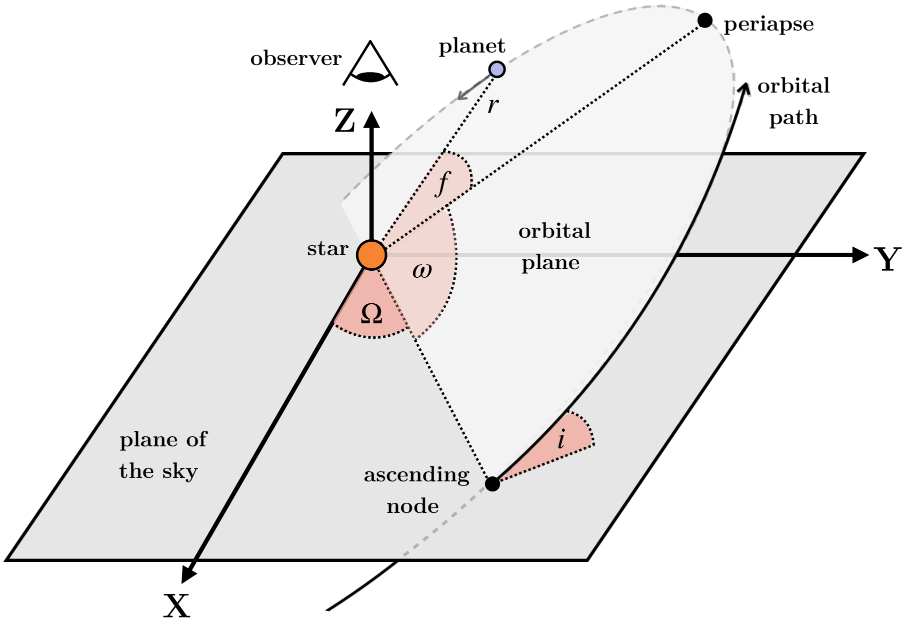

(a) General case orbital geometry.

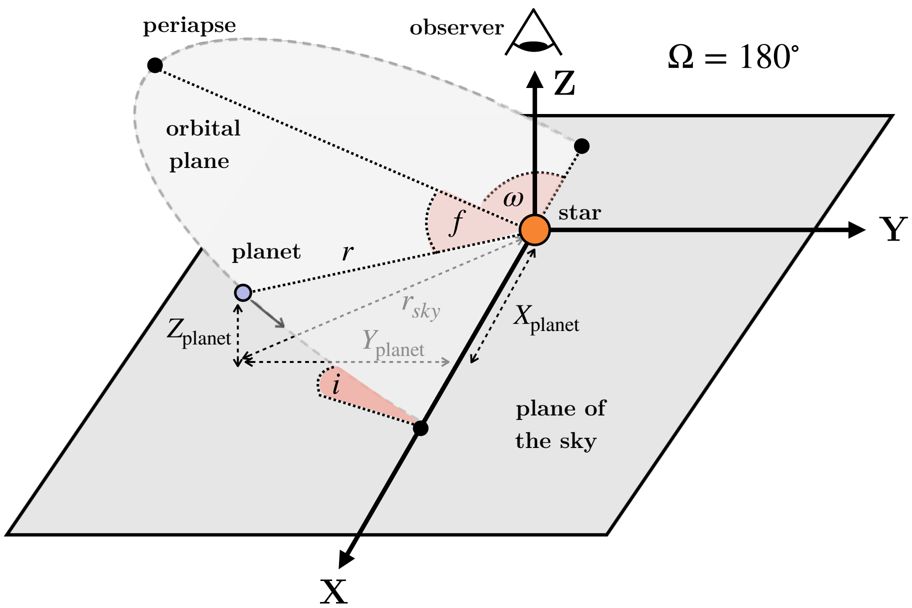

(b) Orbital geometry at Ω = 180°.

Figure 1.1: The geometry of a planetary orbit. (a) demonstrates the (X, Y, Z) polar coordinate system inclined to the orbital plane and the various angles (i, Ω, ω, f) used in Section 1.2.1. Adapted from Murray & Correia (2010). (b) presents the specific case where $\Omega = 180°$. The quantities $X_{\rm planet}$, $Y_{\rm planet}$, $Z_{\rm planet}$ (equivalent to X, Y, and Z in the text) and $r_{\rm sky}$ are shown on this plot and derived using Equations 1.2, 1.3, 1.4 and 1.5 respectively.

to the planet, as shown in Figure 1.1. If we resolve into Cartesian coordinates, as in Murray & Correia (2010) and Winn (2010), centred on the star, with the sky in the XY plane, the orbital plane intersecting along the X-axis ($\Omega = 180°$in Figure 1.1) and an observer looking along the Z-axis, we have the configuration shown in Figure 1.1b. We can describe the planet's position in this coordinate system by the following equations:

$$X = -r\cos(\omega + f), \tag{1.2}$$

$$Y = -r\sin(\omega + f)\cos i, \tag{1.3}$$

$$Z = -r\sin(\omega + f)\sin i, \tag{1.4}$$

where ω is the argument of periapse (the angle between the $-X$-axis and periapse on the orbital plane in this case) and i is the inclination (the angle between the sky plane and the orbital plane). We can project the planet-star distance, r in the orbital plane, onto the sky plane, given by:

$$\begin{aligned} r_{\mathrm{sky}} &\equiv \sqrt{X^2 + Y^2} \\ &= r\sqrt{1 - \sin^2 i \sin^2(\omega + f)} \\ &= \frac{a(1 - e^2)}{1 + e\cos f}\sqrt{1 - \sin^2 i \sin^2(\omega + f)}, \end{aligned} \tag{1.5}$$

using Equations 1.1, 1.2 and 1.3. From Figure 1.1b we see that for an eclipse to occur along the observer's line-of-sight, r_{sky} has to be at a local minimum, as this is the point where the visible on-sky distance between the planet and star is smallest. We can use the approximation that r_{sky} is minimal at $X = 0$, which gives the following solutions:

$$f_{\mathrm{tra}} = +\frac{\pi}{2} - \omega \quad \text{and} \quad f_{\mathrm{occ}} = -\frac{\pi}{2} - \omega, \tag{1.6}$$

for the situations where the planet passes in front of the star (a transit), f_{tra}, or the planet passes behind the star (an occultation), f_{occ}. If we take only the transiting scenario, we define the impact parameter b as r_{sky} at the point of mid-transit, which is expressed as a fraction by normalising by the radius of the star:

$$b = \frac{a\cos i}{R_*}\left(\frac{1 - e^2}{1 + e\sin\omega}\right). \tag{1.7}$$

The impact parameter of a transit is demonstrated in Figure 1.2. For a transit to occur, the planet has to obscure a fraction of the stellar face, or equivalently:

$$|b|R_* < R_* + R_{\mathrm{p}}. \tag{1.8}$$

If the planet only ever partially eclipses the star it is called a grazing transit, and must also satisfy $|b|R_* > R_* - R_{\mathrm{p}}$. Whereas for a full non-grazing transit the planet must satisfy the more

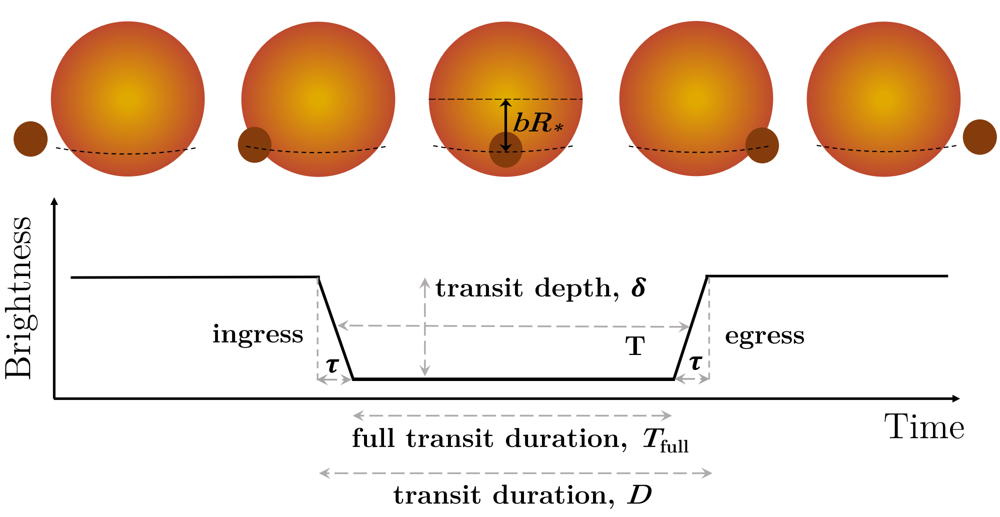

Figure 1.2: Illustration of a planetary transit and its resulting lightcurve. The beginning of a planet's eclipse of its host star is called the ingress, and the end of the eclipse is called the egress. The transit depth, δ, is the maximal dip in brightness once the planetary disc is entirely within the stellar disc. The transit duration, D, is the time between the planetary disc appearing to contact the stellar disc before and after a transit (from start of ingress to end of egress). The full transit duration, $T_{\rm full}$, is the time from end of ingress to start of egress (when the planetary disc is completely within the stellar disc). A useful timescale, T, is marked as the time between midpoints of ingress and egress. The impact parameter b is shown on the middle star and is a measure of the projected distance between the centre of the star and the centre of the planet at the point of mid-transit, normalised by stellar radius, R_*.

restrictive condition $|b|R_* \leq R_* - R_{\rm p}$. We can combine Equations 1.7 and 1.8 into:

$$|\cos i| < \left(\frac{R_* + R_{\rm p}}{a}\right)\left(\frac{1 + e\sin\omega}{1 - e^2}\right) = (\cos i)_{\rm tra}, \tag{1.9}$$

where $(\cos i)_{\rm tra}$ is the condition on $\cos i$ for a transit. Using the fact that $\cos i$ is a randomly distributed number between 0 and 1 (as an observer could be placed anywhere), this leads to:

$$p_{\rm tra} = \frac{\int_0^{(\cos i)_{\rm tra}} d(\cos i)}{\int_0^1 d(\cos i)} = \left(\frac{R_* + R_{\rm p}}{a}\right)\left(\frac{1 + e\sin\omega}{1 - e^2}\right), \tag{1.10}$$

as an expression for the geometric probability, $p_{\rm tra}$, of observing a transit, either grazing or full. When a planet orbits its star, it produces a cone of shadow that sweeps out a band on the celestial sphere. The closer the planet orbits to its parent star, the larger the area that the cone covers. When the planet moves closer to the star on an elliptical orbit, the transit can, therefore, be observed from a wider range of inclinations and vice versa.

If we assume the star is much larger than the planet ($R_* >> R_\mathrm{p}$), then the probability that a circular star-planet system ($e = 0$) is favourably oriented to observe a transit reduces to:

$$p = \frac{R_* + R_\mathrm{p}}{a} \approx \frac{R_*}{a} \approx 0.005 \left(\frac{R_*}{R_\odot}\right) \left(\frac{a}{1\,\mathrm{au}}\right)^{-1}, \tag{1.11}$$

where $R_\odot$ is the solar radius. This equation demonstrates that only a narrow range of planet-star inclinations have the required geometry for an eclipse. Essentially, to catch a transit we must be incredibly lucky, but we have an increased chance of detecting planets that orbit closer to their host star, or around larger stars.

Transit Durations

A transit starts when the planet initially begins to eclipse the star. The transition from when the planetary disc just "touches" the stellar disc to when the planet is first fully eclipsing (in a non-grazing transit) is called the ingress. The opposite scenario, when the planet goes from the last moment of full eclipse to completely lose "contact" with the stellar disc, is called the egress. For a circular orbit the ingress and the egress times, τ, are equal which generally does not hold in the eccentric case, although in reality the difference is very minor (of the order of seconds to minutes, Kipping 2008). See Figure 1.2 for a demonstration of a transit.

The transit duration is defined as the total time taken from the very beginning to the very end of transit (or the start of ingress to the end of egress). Using Kepler's second and third laws of planetary motion, Winn (2010) derive the following relation between the transit duration, D, and the orbital period, P (their Equation 16):

$$D = \frac{P}{\pi} \arcsin\left(\frac{R_*}{a} \frac{\sqrt{\left(1 + \frac{R_\mathrm{p}}{R_*}\right)^2 - b^2}}{\sin i}\right) \cdot \frac{\sqrt{1 - e^2}}{1 + e \sin\omega}. \tag{1.12}$$

This equation allows the transit duration to be extracted from knowledge of the orbital parameters (i, e, ω, b, P) and the sizes of the planet and star. Similarly, the full transit duration, T_full, (which includes only the portion of the transit when the planetary disc is entirely within the stellar disc) is defined as the time between the end of ingress and start of egress:

$$T_\mathrm{full} = \frac{P}{\pi} \arcsin\left(\frac{R_*}{a} \frac{\sqrt{\left(1 - \frac{R_\mathrm{p}}{R_*}\right)^2 - b^2}}{\sin i}\right) \cdot \frac{\sqrt{1 - e^2}}{1 + e \sin\omega}. \tag{1.13}$$

With expressions for these specific time intervals, Kepler's second law can then be applied to calculate any time interval on that orbit. For more details of these derivations see Winn (2010).

If we simplify by considering a transiting planet on a circular orbit (e=0), assuming that $R_\mathrm{p} \ll R_* \ll a$ and $\sin i \approx 1$, then we can compute T, the time between the halfway points of ingress and egress:

$$\begin{aligned} T &\equiv D - \tau \\ &= D - \frac{1}{2}(D - T_\mathrm{full}) \\ &\approx T_0\sqrt{1 - b^2}, \end{aligned} \tag{1.14}$$

where the characteristic timescale, T_0 is given by:

$$T_0 \equiv \frac{R_* P}{\pi a}. \tag{1.15}$$

Transit Depth

We can measure the subsequent dip in brightness during a transit and use this flux drop to probe various planetary parameters. The fraction of stellar flux blocked by the planet, or transit depth δ, can be represented as:

$$\delta \equiv \frac{F_\mathrm{out\,of\,transit} - F_\mathrm{in\,transit}}{F_\mathrm{out\,of\,transit}}, \tag{1.16}$$

where $F_\mathrm{out\,of\,transit}$ is the flux we observe out of transit and $F_\mathrm{in\,transit}$ is the resulting flux when the planet is fully transiting (maximum decrease in flux). Note that the resulting flux is not only stellar in origin; there is also a small fraction of reflected and emitted light coming from the planet. The stellar flux also varies as a function of time due to stellar variability (from surface inhomogeneity), which is discussed in more detail in Section 1.2.5. However, for simplicity we will assume it is constant, and that the planetary and stellar discs are both uniform in brightness (see the following section on limb-darkening for when this is not the case) with surface brightness $F_\mathrm{s,\,p}$ and $F_\mathrm{s,\,*}$ respectively. We see the planet's nightside (the side of the planet not illuminated by the star) during a transit, which means that there is very little reflected light. Therefore, we can make the assumption $F_\mathrm{s,\,*}\,A_* \gg F_\mathrm{s,\,p}\,A_\mathrm{p}$, where A_* and A_p are the projected areas of the stellar and planet discs respectively.

$$\begin{aligned} F_\mathrm{out\,of\,transit} &= F_\mathrm{s,\,*}\,A_* + F_\mathrm{s,\,p}\,A_\mathrm{p} = F_\mathrm{s,\,*}\,\pi R_*^2 + F_\mathrm{s,\,p}\,\pi R_\mathrm{p}^2 \\ F_\mathrm{in\,transit} &= F_\mathrm{s,\,*}\,(A_* - A_\mathrm{p}) + F_\mathrm{s,\,p}\,A_\mathrm{p} = F_\mathrm{s,\,*}\,\pi(R_*^2 - R_\mathrm{p}^2) + F_\mathrm{s,\,p}\,\pi R_\mathrm{p}^2 \\ \delta &= \frac{F_\mathrm{s,\,*}\,\pi R_\mathrm{p}^2}{F_\mathrm{s,\,*}\,\pi R_*^2 + F_\mathrm{s,\,p}\,\pi R_\mathrm{p}^2} \\ \delta &\approx \left(\frac{R_\mathrm{p}}{R_*}\right)^2. \end{aligned} \tag{1.17}$$

Therefore, the transit depth, δ, is directly related to the size ratio between the planet and its star. We can also see that the larger the planet, with respect to the star, the larger the transit depth, and the easier it will be to detect. Taking the solar system as an example, Jupiter would have a transit depth of ∼1%, compared to ∼0.01% for the Earth.

Limb Darkening

We previously assumed that the stellar disc is uniform in brightness. In reality, we see an optical phenomenon known as limb-darkening, where stellar discs appear brighter in the centre and dimmer towards the edges (or the limb). This gradient effect is due to the geometry of the stellar atmosphere and the fact that the effective temperature of the photosphere generally decreases as we move further away from the centre of the star. We make the Eddington-Barbier approximation that the majority of stellar light (with wavelength λ) reaching the Earth has originated from layers in the star's atmosphere with an optical depth τ_λ=1. In the centre of the stellar disc, radially-emitted light is directed towards the observer, and an optical depth of 1 probes deep into the atmospheric layers. At the edge of the stellar disc, light emitted (directed towards the observer) originating from an optical depth of unity is from atmospheric layers that are higher altitude, colder and less radiative, than the same depth in the centre of the disc. Therefore, the stellar surface will appear darker towards the edges and brighter in the middle.

The practical consequence of limb-darkening is that a planet transiting the stellar disc will have a smaller transit depth near the limb and a larger transit depth towards the middle. Limb-darkening, therefore, manifests as the transit "rounding-off" at the base, showing a U-shape rather than the classic flat-bottom.

Deriving Physical Parameters from Transits

The observation of a single transit yields the measurement of three main parameters; the transit duration D, the full transit duration $T_{\rm full}$, and the transit depth δ. If we catch at least two consecutive transits, this provides the orbital period P. Significant phase coverage of a planet's orbit is required to determine the orbital period precisely (to ensure we have not missed transits at fractions of this period). Combining these observables will allow us to derive a number of additional parameters, such as $R_{\rm p}/R_*$ (as in Equation 1.17), i, b, the scale parameter a/R_*, and the stellar density ρ_*.

Considering the simplest case of a transiting planet on a circular orbit (e = 0), where $R_* \ll a$, we already know:

$$\frac{R_{\rm p}}{R_*} = \sqrt{\delta}. \tag{1.18}$$

This allows us to simplify Equations 1.12 and 1.13:

$$D^2 = T_0^2((1+\sqrt{\delta})^2 - b^2) \quad \text{and} \quad T_{\rm full}^2 = T_0^2((1-\sqrt{\delta})^2 - b^2),$$

to obtain the following two equations:

$$b = \sqrt{\frac{(1-\sqrt{\delta})^2 - \frac{T_{\rm full}}{D}^2(1+\sqrt{\delta})^2}{1-\frac{T_{\rm full}}{D}^2}}, \tag{1.19}$$

and

$$\frac{a}{R_*} = \frac{2P}{\pi} \cdot \frac{\delta^{1/4}}{\sqrt{D^2 - T_{\rm full}^2}}. \tag{1.20}$$

These three values (Equations 1.18, 1.19 and 1.20) are then key to extracting the planetary system's physical parameters. Following from Kepler's third law of planetary motion, we have the following expression:

$$\frac{P^2}{a^3} = \frac{4\pi^2}{G(M_* + M_{\rm p})}, \tag{1.21}$$

where G is the universal gravitational constant, M_* is the stellar mass and $M_{\rm p}$ is the planetary mass. Once we know a/R_*, and we make a reasonable assumption that $M_* \gg M_{\rm p}$, then we can use the above expression to obtain the stellar mean density (Seager & Mallen-Ornelas 2003):

$$\rho_* = \frac{M_*}{\frac{4}{3}\pi R_*^3} \approx \frac{3\pi}{P^2 G}\left(\frac{a}{R_*}\right)^3. \tag{1.22}$$

While the stellar mean density does not provide any information about the planet alone, it is a useful constraint of the stellar parameters.

From our calculation of b we can also rearrange Equation 1.7 to obtain the inclination of the orbital plane:

$$i = \arccos \frac{bR_*}{a}. \tag{1.23}$$

Disadvantages

The main drawback of the transit detection method is that the alignment of the planetary system must be just right to be able to observe the transit from Earth. Additionally, as transits only occur once every orbital period, with long-period planets (periods of months or years), it is possible to observe a star for a long time and not observe any of its planets' transits. This is especially challenging for ground-based telescopes, which only can observe at night and whose operations are limited by sub-optimal weather conditions. Therefore, this method favours the detection of planets on short orbits, close to the star. This limitation is compounded by the probability of observing a planetary transit scaling with $\sim \frac{R_*}{a}$ (see Equation 1.11). The chance of an external observer detecting the Earth's transit is ~0.5%, but rises to 10% for a typical hot Jupiter at 0.05 AU around a Sun-sized star.

While an excellent method to determine the planetary radius, observing a transit alone does not tell us anything about a planet's mass. In the specific case of a multi-planet system, we can monitor transits over long baselines to detect slight changes in the times between transits, called transit timing variations (TTVs). Deriving TTVs can constrain planet masses (and are described in more detail in the following section). However, often to determine the mass (when we have limited observing time or are considering a single planet), we would use another detection method, such as radial velocity. This technique also only offers relative planetary parameters such as R_p/R_* and a/R_* which we use to extract physical variables (a, ρ_*, R_p etc.); therefore, their accuracy is dependent on how well the star is characterised.

Besides a planetary transit, several astrophysical sources can cause a light source to dim in brightness. Therefore, the transit method has a high rate of false positives. I will discuss these in more detail in Section 1.2.5. Additionally, if the star itself varies in brightness through stellar flares or rotational flux modulations, it significantly complicates the detection of planets.

1.2.2 Ancillary Products

Observations of planetary eclipses not only provide a method of measuring a planet's size, but also reveal a treasure trove of additional information.

Transit Timing Variations

We expect that a planet orbiting its host star, with period P, would produce regular transits every P days. However, this is only the case when the system is a single planet on a Keplerian orbit. If there are additional gravitational influences in the system (e.g., a multi-planet system or a planet orbiting a stellar binary), then the transits will deviate from being periodic (Agol et al. 2005; Holman & Murray 2005; Heyl & Gladman 2007; Agol & Fabrycky 2018). We can detect transit timing variations (TTVs) by monitoring transit ephemerides over long time periods. TTVs can provide evidence for undetected satellites or planets (perhaps non-transiting), and because they are a result of gravitational interactions, they can also be used to estimate planet masses (Carter et al. 2012; Hadden & Lithwick 2014; Jontof-Hutter et al. 2016).

Transit Spectroscopy

When a planet moves in front of its host star, a fraction of the light from the star filters through that planet's atmosphere. Different atoms, molecules and condensates in the planetary atmosphere will absorb light at various wavelengths and leave other wavelengths unaltered. Therefore, the magnitude of the resulting light at different wavelengths is a probe of the photochemistry within a planet's atmosphere. As a planet's atmosphere varies in opacity as a function of wavelength, the planetary radius will appear to vary accordingly. By finding this wavelength-dependence of the transit depth, we can build a spectrum, and by considering the difference between the resulting spectra in and out of transit, we can infer information about the planet's atmospheric composition (Deming & Seager 2017; Kreidberg 2018). By looking for key biosignatures, we can also make predictions about a planet's potential to host life (Seager et al. 2016; de Wit & Seager 2013). However, the presence of hazes and clouds in a planet's atmosphere severely impact precise atmospheric characterisation by dampening features in the transmission spectrum (Kreidberg et al. 2014; Deming & Seager 2017). The wavelength-dependent opacity of a planet's atmosphere determines the depth of spectral features, and when there are (optically thick) hazes and clouds, planetary atmospheres appear opaque, and the transit depth does not change, resulting in featureless spectra.

Eclipse Spectroscopy

Instead of observing a planet's transits, we can monitor its occultations (when it passes behind the star). When the star obscures the planet, the resulting spectrum is only stellar. When the planet re-emerges, we see an increase in brightness, stemming from its thermal emission and reflected light (Kreidberg 2018; Alonso 2018). A planet's atmosphere can emit light through reflecting and scattering stellar light (dominates at optical wavelengths), absorbing and re-radiating stellar light (dominates at infrared wavelengths), or through its intrinsic thermal processes (at infrared wavelengths, though this emission is generally negligible). In this method, we measure the planet's day-side (the face of the planet the star illuminates) brightness temperature or geometric albedo (defined as the proportion of incident light that is re-emitted) using the drop in flux before and during the secondary eclipse. Due to the wavelength-dependent opacity of a planet's atmosphere and vertical pressure-temperature profile, this technique can provide insights not only into the atmospheric composition but also its thermal structure (Stevenson et al. 2014; Kreidberg 2018).

Phase Curves

By using continuous time-series photometry to capture the entire orbital lightcurve (or phase curve) of a planet, we can detect time variations in the planet's emission (Parmentier & Crossfield 2018). A complete phase curve is shown in Figure 1.3. This technique allows for a wider longitudinal coverage than just transits or occultations so that we can map the different faces of a planet (Cowan & Fujii 2018). By generating a phase curve, we can explore additional aspects of the planet's atmospheric physics, such as day-night heat transport (Stevenson et al. 2014), seasonal variations, atmospheric evolution (Armstrong et al. 2017; Demory et al. 2016a) and rotational cloud mapping (Cowan & Fujii 2018; Demory et al. 2013).

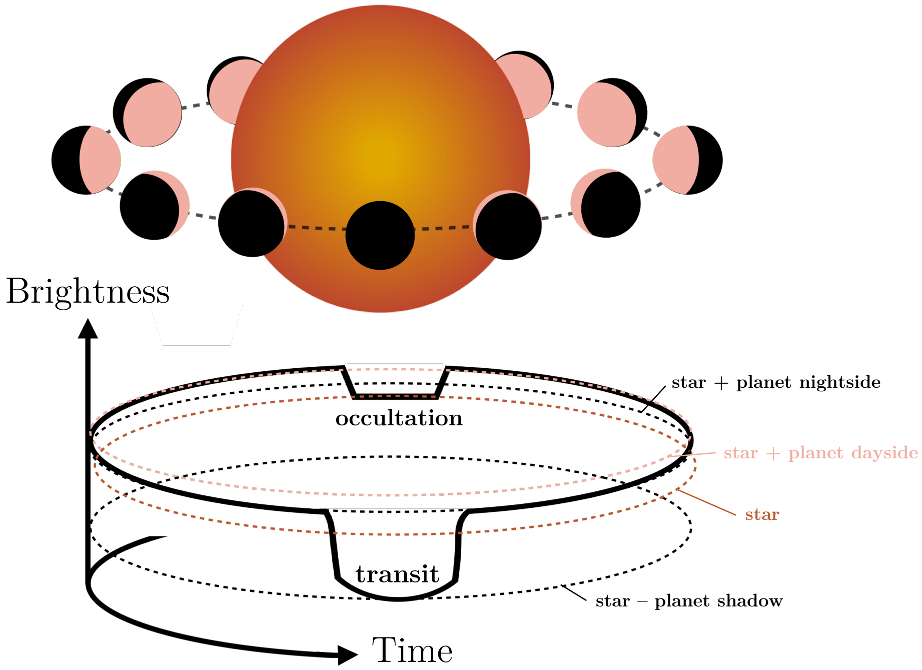

Figure 1.3: A demonstration of the main brightness variations we see during a typical orbit cycle. When the planet passes in front of the star, it blocks some of the star's light, producing a brightness dip known as the primary transit. When a planet is partially- or un-eclipsed by its host star, it will reflect some of the star's light. Whereas when the planet is fully obscured by the star then we see only light coming from the star. This produces a smaller dip in brightness called a secondary eclipse, or an occultation. This plot is inspired by Figure 1 from Winn (2010).

1.2.3 Searching for Transits from the Ground and the Sky

A transit was first observed for a planet orbiting the star HD 209458 by Charbonneau et al. (2000) shortly followed by the first exoplanet discovery of OGLE-TR-56b using the transit method (Udalski et al. 2002). Since then, 3438 of all 4551 confirmed exoplanets have been discovered with this method[b]. These transiting planets were predominantly found during the almost 10-year span of NASA's Kepler mission (Borucki et al. 2010), which is responsible for ~70% of these planets. The prosperity of Kepler, and its successor TESS, confirm the transit technique as the most prolific detection method to date.

In the 20 years since this first detection, transit surveys have become ubiquitous. There are several ways in which photometric surveys have been optimised to detect planets, by observing more stars, using faster cadences, monitoring stars for longer baselines, and increasing precision. Improved photometric precisions have meant that the type of planets we can detect has expanded

[b]From the NASA Exoplanet Archive as of 01 November 2021

dramatically. From being limited to only detecting the transits of planets larger than Jupiter, we are now beginning to scratch the surface of the detection of temperate, Earth-sized planets. There has been a significant focus on space-based telescopes to facilitate this. We are not limited by the day-night observation cycle in space, and we do not have to grapple with the Earth's rapidly changing atmosphere. Atmospheric contamination has meant that photometric precisions of the order of parts per million (ppm) are currently only achievable by space missions; however, ground-based observatories have demonstrated an ability to reach sub-mmag level measurements. Nevertheless, even today, Earth-sized planets in the habitable zone around Sun-like stars remain inaccessible with our current capabilities. One tactic employed by recent transit surveys is to look for small, terrestrial planets around different kinds of stellar hosts, such as cool red dwarfs, that bypass the instrumental limitations we face with Sun-like hosts. Alternatively, we can be patient for the upcoming decade, during which the launches of several new space telescopes will help bring the sought-after "Earth-twin" into reach.

Slightly less than 400 planets have been detected from various ground-based surveys, approximately a tenth of all planets and an eighth of the number detected by space missions. While the returns significantly differ between ground and space-based missions, there are compelling advantages to observing on the ground. Ground-based telescopes are much cheaper, more flexible, and can be easily repaired, updated or re-purposed. The 851 transiting planets for which we have both mass and radii measurements are plotted in Figure 1.4, demonstrating the different parameter spaces probed by ground and space-based missions. Therefore, utilising ground and space-based synergies will be essential to finding the best planetary candidates for future habitability studies.

A Timeline of Ground-Based Photometric Surveys

The first generation of ground-based transit surveys consisted of automated, small aperture (~0.1 m) wide-field surveys. This era's main surveys have been (in order of increasing planet detections): TrES (Transatlantic Exoplanet Survey, Dunham et al. 2004; Alonso et al. 2004), XO (McCullough et al. 2005), HATNet[c] (Hungarian-made Automated Telescope Network, Bakos et al. 2004), and WASP[d] (Wide-Angle Search for Planets, Pollacco et al. 2006). These surveys achieved photometric precisions of order 0.01 magnitude, sufficient to detect a Jupiter-sized planet around a bright ($V < 13$ mag) host star. See Section 3.1.4 for a description of magnitudes. The two most successful projects to come out of this first generation were HATNet and WASP. HATNet consists of seven 0.11 m telescopes located across the Northern hemisphere (five in Arizona and two in Hawaii). Since first light in 2003, HATNet has discovered 70 transiting

[c] https://hatnet.org/
[d] https://wasp-planets.net/

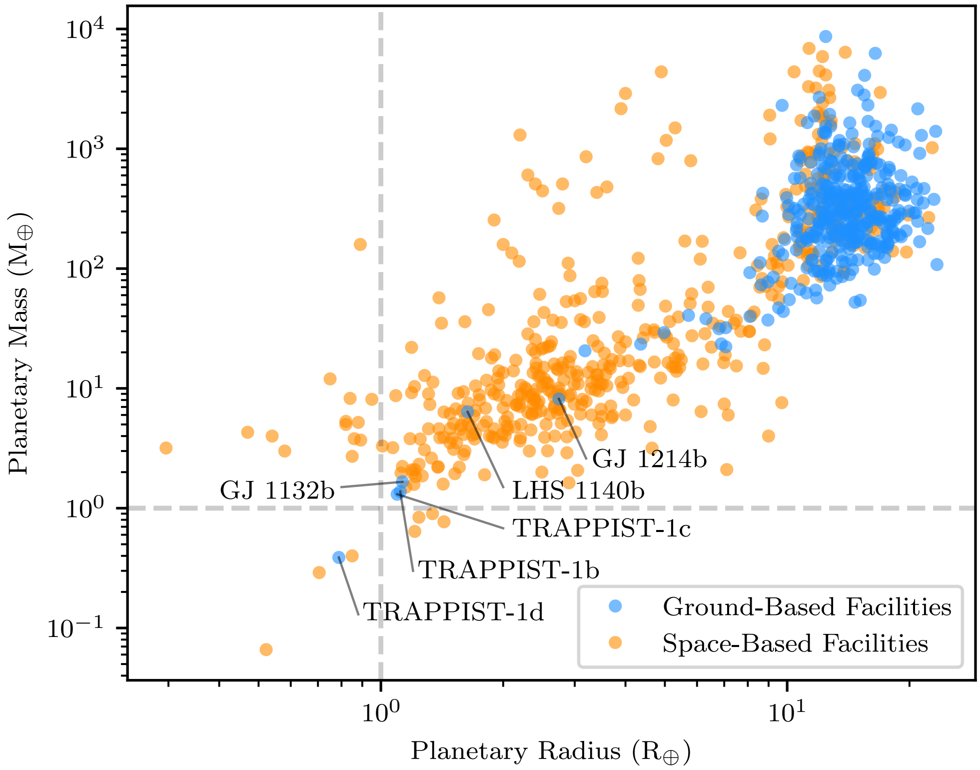

Figure 1.4: Mass-radius diagram for the 851 (out of a total 3438) confirmed transiting planets, which have both radius and mass measurements, as of 1st November 2021 on the NASA Exoplanet Archive. Of these 851 planets, 487 were detected by space missions, shown in orange, and ground-based telescopes detected the 364 planets in blue. Except for a handful of planets (such as the three highlighted TRAPPIST-1 planets, LHS 1140b, GJ 1132b, and GJ 1214b), the majority of ground-based detections are planets larger than $10R_{\oplus}$, and heavier than $50M_{\oplus}$. On the other hand, space-based facilities have the high precisions needed to detect much smaller and less heavy planets. Empirically, exoplanets follow an increasing mass-radius relation for planets lighter than $\sim$100 $M_{\oplus}$, and a poor dependence on radius for heavier planets (Weiss et al. 2013; Chen & Kipping 2016; Bashi et al. 2017). These works have suggested that as small planets gain mass, they increase in size by depositing gas in their outer layers. However, the relationship between planetary radius and mass is particularly dependent on the composition of the planet, which results in a broad spread of masses for a given radius. The heavier planets in this plot show a weak correlation with radius, likely due to their predominantly hydrogen-helium compositions and high gravitational self-compression, which results in an electron degeneracy pressure that roughly balances the Coulomb force (Zapolsky & Salpeter 1969; Guillot 2005; Seager et al. 2007; Bashi et al. 2017). Grey dotted lines show one Earth radius and mass.

exoplanets. There was also a later addition to HATNet of eight 0.2 m telescopes at three Southern sites (Chile, Namibia and Australia), called HATSouth[e] (Bakos et al. 2013). HATSouth has doubled the output of the HATNet project, discovering an additional 73 planets since its first light in 2009. There are plans underway for the HATPI[f] project, which will observe almost the entirety of the night sky from Chile, using 63 instruments. The extensive longitudinal coverage of this network reduces the impact of the day-night cycles and can allow HATNet to detect longer period planets than possible with only one site. However, the quality of observing sites and the total time coverage of stellar hosts are significant factors for successfully detecting these types of planets. WASP, which became fully operational in 2006, has two wide-field camera arrays of eight 0.11 m telescopes in the Northern (SuperWASP-North in the Canary Islands) and Southern (SuperWASP-South in South Africa) hemispheres. As the name suggests, WASP has an extensively wide-field coverage allowing it to survey the entire sky, which is why WASP is currently the most fruitful ground-based transit survey, claiming over 180 confirmed planets.

The past decade has seen massive advances in photometric precisions with ground-based observatories following the lessons learnt from these early transit surveys. These advances are due in part to improved instrumentation and a wide range of observing strategies. Some surveys decided to follow the wide-field technique of their predecessors, such as NGTS[g] (the Next Generation Transit Survey, Wheatley et al. 2018), QES[h] (Qatar Exoplanet Survey Alsubai et al. 2014), KELT[i] (Kilodegree Extremely Little Telescope, Pepper et al. 2007) and MASCARA[j] (Multi-site All-Sky CAmeRA, Talens et al. 2017). KELT, QES and MASCARA all target very bright stars searching for "hot Jupiters", Jupiter-sized planets orbiting very close to their host stars. However, the NGTS project chose to apply the experience of WASP to searching for much smaller planets. NGTS comprises an array of twelve 0.2 m telescopes situated at ESO (European Southern Observatory) Paranal Observatory, Chile. NGTS's lightcurves have demonstrated an unprecedented ten-fold increase in photometric precision, regularly reaching precisions of $\sim$0.1% (Wheatley et al. 2018). Moreover, as they focus on searching for transits around K and early M-type stars, they can detect Neptune-sized planets and smaller (West et al. 2019; Smith et al. 2021b).

Other surveys, such as MEarth[k] (Nutzman & Charbonneau 2008; Irwin et al. 2008), APACHE (A PAthway to the Characterization of Habitable Earths, Sozzetti et al. 2013), TRAPPIST (TRAnsiting Planets and PlanetesImals Small Telescope, Jehin et al. 2011; Gillon et al.

[e] `https://hatsouth.org/`
[f] `https://hatpi.org/`
[g] `https://ngtransits.org/`
[h] `https://www.qatarexoplanet.org/`
[i] `https://keltsurvey.org/`
[j] `https://mascara.strw.leidenuniv.nl/`
[k] `https://lweb.cfa.harvard.edu/MEarth/`

2011) and SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars, Gillon et al. 2018; Burdanov et al. 2018; Delrez et al. 2018b; Jehin et al. 2018), have opted for a much more targeted technique – a "quality not quantity" approach. These surveys instead optimise observations for a single target (or a couple of targets) each night. By focusing on cooler, fainter host stars (such as M-dwarfs and ultra-cool dwarfs), it is possible to detect much smaller planets with moderately-sized telescopes (≥0.4 m class). I will discuss this so-called 'M-dwarf opportunity' more in Section 1.3.1. The MEarth project comprises the MEarth-North observatory of eight 0.4 m telescopes in Arizona and the MEarth-South observatory of eight identical telescopes in Chile. MEarth aims to observe 2000 early-to-mid M-dwarfs with masses between 0.1–0.35$M_{\odot}$. While targeted surveys have smaller yields than wide-field surveys, the planets they find are extremely promising candidates for follow-up. MEarth's three planetary discoveries, GJ1214b (Charbonneau et al. 2009), GJ1132b (Berta-Thompson et al. 2015) and LHS1140b (Dittmann et al. 2017), are all terrestrial planets with sizes from 1–3$R_{\oplus}$ which have each been studied extensively (Kreidberg et al. 2014; Berta et al. 2012b; Louie et al. 2018; Bonfils et al. 2018; Edwards et al. 2020; Lillo-Box et al. 2020). The Ultra-Cool Dwarf mini-survey carried out on TRAPPIST also led to the discovery of a remarkable system of seven temperate, terrestrial planets (Gillon et al. 2016, 2017). I discuss the TRAPPIST project and this specific discovery in more detail in Section 2.1.1.

The launch of TESS in 2018, described in more detail in the following section, altered the planet-finding landscape significantly. TESS's continuous 28-day coverage (and longer for targets in overlapping sectors), high precision, ability to observe bright objects and the wide spectral range of its target stars has led to great advances in exoplanetary studies. As TESS is not optimised to observe the very cool, red dwarfs, the more targeted M-dwarf surveys are relatively unaffected, but wide-field surveys with large overlap have either aligned their strategies with TESS (Lendl et al. 2020; Gill et al. 2020a) or been retired (e.g., KELT). Collaborations with ground-based observatories can be indispensable to space missions through confirming detections, following up single transits (Trifonov et al. 2019; Gill et al. 2020b; Cooke et al. 2020), eliminating false positives (particularly an issue for TESS due to its relatively large pixel scale, Sullivan et al. 2015), and refining planetary parameters.

A Timeline of Space-Based Exoplanet Missions

CoRoT (COnvection, ROtation and planetary Transits, (Fridlund et al. 2006; Auvergne et al. 2009)), launched in 2006, was the first space mission with the goal of detecting rocky exoplanets using the transit method. The CoRoT satellite was a partnership between ESA, Austria, Belgium, Brazil, Germany and Spain, which aimed to detect planets several times larger than the Earth and perform asteroseismology of stars. CoRoT discovered 37 planets and brown dwarfs (see Section

1.3.3) in its six years of operation, and over 100 unresolved planetary candidates (Deleuil et al. 2018).

NASA's Kepler telescope (Borucki et al. 2010) launched shortly after in 2009. Kepler's primary mission was to survey 150,000 main sequence Sun-like stars for transits over approximately four years. Kepler offered, at the time, unprecedented photometric precision (Gilliland et al. 2011; Christiansen et al. 2012) that led to thousands of planetary detections, as well as making significant advances in exoplanetology (Borucki et al. 2011; Lissauer et al. 2011; Howard et al. 2012; Fressin et al. 2013; Batalha et al. 2013; Dressing & Charbonneau 2013, 2015) and stellar physics. There are still over 2000 planet candidates detected by Kepler that have not been confirmed. A modified 2-year K2 phase began in May 2014, adjusting for the failure of a control reaction wheel (Howell et al. 2014). The K2 mission focused on short-period planets around lower-mass stars, bridging the gap between Kepler and TESS. The Kepler telescope was retired in late 2018, bringing its almost 10-year service to a close.

Kepler's retirement paved the way for the newest generation of planet-hunting telescopes, TESS (Transiting Exoplanet Survey Satellite, Ricker et al. 2015) and PLATO (PLAnetary Transits and Oscillations of stars, Rauer et al. 2013). NASA and MIT's TESS satellite[l] launched in 2018. TESS's primary 2-year mission was to monitor 200,000 of the nearest and brightest stars for signs of transits. TESS evolves from the Kepler mission by having a sky coverage 400 times as large and observing targets 30-100 times as bright. Focusing on brighter stars increases the accuracy of derived planetary parameters and facilitates both ground-based follow-ups to confirm planet candidates and atmospheric characterisation. TESS is predicted to find over 14,000 planets in its lifetime (Barclay et al. 2018), approximately 250 of which will be Earth-sized, and around 70 within their host's habitable zone. As of 9th August 2021, the TESS mission has unearthed 4471 planet candidates[m], of which 154 have been confirmed. Notably, this includes the discovery of TOI-700d, the first potentially habitable Earth-sized planet discovered by TESS (Gilbert et al. 2020; Rodriguez et al. 2020). TESS is currently operating in a 27-month extended mission until September 2022. The planets that TESS has found (and will continue to find) will be excellent targets for follow-up studies in the coming decade.

CHEOPS (CHaracterising ExOPlanets Satellite)[n] is a European space telescope launched in 2019 (Broeg et al. 2013; Fortier et al. 2014). One of the main goals of CHEOPS throughout a 3.5-year mission is to observe known transiting planets and precisely measure their radii (with an accuracy of ~2–5%, Deline et al. 2020). With values for both radius and mass, it becomes

[l] `https://tess.mit.edu/`

[m] Taken from `https://tess.mit.edu/publications/`. Of these 4471 planetary candidates, 794 have radii $\leq 4R_{\oplus}$.

[n] `https://cheops.unibe.ch/`

possible to estimate a planet's bulk density and composition. The high precision of CHEOPS allows it to detect the shallow transits of Earth (for the very brightest G-type dwarf stars, Benz et al. 2020), taking us one step closer to detecting terrestrial planets around hosts like our Sun, to Neptune mass planets (1-4 $R_\oplus$). Though not explicitly focused on finding new planets, CHEOPS has already revealed a fifth planet in the TOI-1233 system (Bonfanti et al. 2021), as well as several unknown planets in the TOI-178 multi-planet system (Leleu et al. 2021). The TOI-178 planets have both radii and mass measurements, gained through collaborations with ground-based photometric and radial velocity teams. By accurately characterising known exoplanets, CHEOPS will help narrow down the most promising candidates for atmospheric spectroscopy, and constrain their planetary parameters more precisely.

PLATO[o] is a 4-year European Space Agency (ESA) mission planned for launch in 2026 (Rauer et al. 2013). PLATO aims to answer the key question of whether our Solar System is unique. It will do so by searching for transits around approximately one million stars, specifically detecting and characterising Earth-like planets orbiting in the habitable zones of Sun-like stars. Like TESS, PLATO will study relatively bright stars; however, it will be more sensitive to longer period planets due to an increased time spent on each field. PLATO will consist of 26 optical cameras: 24 standard operating cameras and a set of two "fast" cameras for particularly bright targets. This camera setup allows for sky coverage between 10 and 50%. In addition to detecting planets, PLATO will provide the tightest constraints on planetary radii to date (with an accuracy of ~3%), perform asteroseismology to fully characterise the stellar hosts (by extracting stellar masses, radii and ages), and work in synergy with ground-based radial velocity follow-up to obtain planetary masses. Therefore, PLATO will yield a well-characterised catalogue of temperate, terrestrial planets, with known radii, masses, bulk densities, compositions, and ages, as prime candidates for future habitability studies.

Atmospheric Characterisation

To fully understand a planet's potential to support life, we need to know the chemical composition of its atmosphere. Water vapour and the presence of biosignatures such as molecular oxygen, ozone, and methane, can provide clues to whether a planet is hospitable for life or whether biological processes are already happening. Biosignatures are substances with a biological origin that provide evidence for the presence of life (in the past or present). However, individually, many promising biosignatures can also have geological or photochemical origins, resulting in abiotic "false positives" and so must be interpreted with care. Transmission and eclipse spectroscopy (see Section 1.2.2) are powerful tools used to perform atmospheric char-

[o] https://platomission.com/

acterisation. Though neither missions' primary objective, the Hubble Space Telescope[p] and the retired Spitzer Space Telescope[q] have already made the first steps towards probing for various chemical "fingerprints" hidden in planetary atmospheres (Swain et al. 2008; Tinetti et al. 2007; Swain et al. 2009; Kreidberg et al. 2014; Dressing & Charbonneau 2015). However, as we move towards studying smaller, terrestrial planets, there is a need for greater spectroscopic precision and state-of-the-art technologies.

There are three major upcoming missions aiming to characterise the atmospheres of the best candidate planets; JWST (James Webb Space Telescope, Gardner et al. 2006)[r], LUVOIR (Large Ultraviolet Optical Infrared Surveyor, The LUVOIR Team 2019)[s], and ARIEL (Atmospheric Remote-sensing Infrared Exoplanet Large-survey, Tinetti et al. 2018)[t]. The ultimate objective of atmospheric characterisation is to find a planet where life could exist. These missions will obtain high quality transmission and emission spectra of the best planetary targets to scour for major atmospheric gases, metals and key biosignatures. JWST will be a particularly powerful tool. It will be able to determine (for clear solar composition planetary atmospheres) or constrain (for planets with cloudy atmospheres) the volume mixing ratios of dominant molecules (such as those of water, methane and carbon dioxide), atmospheric C/O ratio, metallicity and temperature-pressure profiles (Gardner et al. 2006; Greene et al. 2016). JWST's launch is planned for December 2021. JWST will have a 6.6 m primary mirror and several scientific instruments that give a wide spectral coverage from the visible to long-infrared (0.6-27 µm). While ARIEL will have a much smaller collecting area than JWST (1 m class primary mirror), it will have more observing time allocated to exoplanets (100% compared to 16% of JWST time). ARIEL's mission is to observe at least 1000 known exoplanets, with wavelength coverage extending from the visible to mid-infrared (~0.5-7.8 µm), to study and characterise their chemical composition. The large sample ARIEL will observe will help address questions on the formation and evolution of exoplanets. Therefore, the ARIEL and JWST missions will be highly complementary. LUVOIR is currently under development with a proposed launch date of 2039. LUVOIR's significantly larger mirror (15 m primary) will cause a marked increase in sensitivity from JWST and an overlapping wavelength coverage from near-UV to near-infrared (0.2-2 µm). Therefore, LUVOIR will build on the UV spectroscopy capabilities of Hubble. LUVOIR's ultra-high precision will make it the first space telescope to perform atmospheric characterisation of an Earth-sized planet in the habitable zone around a Sun-like star. LUVOIR will also assess the frequency of planets with environments conducive for life and explore the

[p]https://hubblesite.org/
[q]https://www.spitzer.caltech.edu/
[r]https://www.jwst.nasa.gov/
[s]https://www.luvoirtelescope.org/
[t]https://arielmission.space/

diversity of "habitable" planetary conditions.

The future of ground-based exoplanetary studies also looks highly promising as we move into the next generation of observatories: extremely large telescopes (ELTs). There are two ELTs currently under construction that will study exoplanets: the European Extremely Large Telescope[u] (E-ELT, 39.3 m, planned first light in 2027) at the European Southern Observatory (ESO) in Chile, and the Thirty Meter Telescope[v] (TMT, 30 m, planned first light in 2027), located on Maunakea in Hawaii. Once constructed, the E-ELT will become the world's largest optical and near-infrared telescope for at least the next decade. Both the E-ELT and TMT will be powerhouses in exoplanetary studies. They will host instruments with extremely high spectral resolution and spectral stability that will increase the sensitivity of radial velocity measurements to cm s^{-1} precisions to detect terrestrial-mass planets. They will perform physical characterisation of known planets through high angular-resolution direct imaging and spectroscopy. These new telescopes will also have the potential to achieve the high spectral resolutions necessary to characterise planetary atmospheres.

1.2.4 Contamination of Planetary Signals from Non-Astrophysical Origins

All exoplanet transit surveys must tackle a unique combination of errors affecting their photometry to optimise their planetary detection potential. This effect can limit the precision that can be reached or imprint time-varying photometric features into data, both of which complicate the retrieval of transits. The sources of photometric noise can be instrumental, atmospheric, or astrophysical in origin. It is also helpful to separate uncorrelated noise, also called "white noise", from systematics or correlated noise, called "red noise". Both white and red noise can arise from instrumental errors and atmospheric effects; however they must be treated differently. A key distinction is that while white noise can be reduced by averaging multiple data points, mitigating systematics (which can vary with time or wavelength) is much more complex. A detailed description of dealing with red noise is in Pont et al. (2006).

Photon noise provides a fundamental limit to the photometric precision that can be reached in space and on the ground. Due to the quantum nature of light, the detection of photons arriving at a detector is a random process. Each photon's arrival is independent of any other photon; therefore, the number of photons we detect follows a Poisson distribution. The number of photons goes as $N \pm \sqrt{N}$ with standard deviation:

$$\sigma_{\text{photon}} = \sqrt{N}. \tag{1.24}$$

[u] `https://elt.eso.org/`
[v] `https://www.tmt.org/`

σ_{photon} is also called photon noise, Poisson noise, or shot noise. The signal-to-noise ratio (SNR), $N/\sqrt{N} = \sqrt{N}$, can only be increased by collecting more photons, e.g., through larger telescopes, longer exposure times or observing brighter stars. As photon noise is white, its impact can also be reduced by binning multiple frames together. On the ground, however, the presence of an atmosphere makes it extremely difficult to reach the photon limit. There are many additional sources of error, from the atmosphere and instrument, that can dominate the photon noise, and must be taken into account to optimise our photometric precision.

From the Atmosphere

Observing with ground-based telescopes has the added complication that light travels through the Earth's atmosphere. The most obvious drawbacks are that two-thirds of observation time is lost due to the day-night cycle and that the ability to observe is subject to rapidly changing weather conditions. However, even when the observing conditions are favourable, the atmosphere can contaminate photometry in ways that are particularly difficult to disentangle. The primary sources of atmospheric error to consider are as follows:

- *Scintillation*: We are looking through a stratified atmosphere composed of many layers of air with differing temperatures and, therefore, densities. When wind causes these layers to mix, a phenomenon known as optical turbulence occurs. Optical turbulence is where the refractive index, determined by the air's density, varies temporally and spatially across the sky. The wavefront from an astronomical source (light from a distant star) can be considered as flat at the top of the Earth's atmosphere, but as it passes through the atmosphere, the waveform is diffracted and perturbed by the various layers of air (with spatially and temporally-varying refractive indices) it encounters. This turbulence has two effects. Firstly, the deformation of the wavefront affects the angular resolution that can be obtained by telescopes on the ground, as the wavefront that reaches the telescope is no longer 'flat'. Secondly, the light reaching the ground fluctuates in spatial intensity, due to local focusing and defocusing effects. This second effect is called scintillation and is why we observe stars "twinkling" here on Earth and not in space. Scintillation is known as a "second-order" effect as it results from the curvature of the wave and is dominated by high-altitude turbulence. Scintillation effects are enhanced towards the horizon, as the light must pass through many more layers of atmosphere to reach the ground, compared to the thinnest atmosphere vertically upwards (i.e., at zenith). It is possible to minimise scintillation and obtain regularly good seeing (seeing refers to how atmospheric turbulence degrades images e.g. through scintillation, through phase aberrations etc.) by observing from a location with low atmospheric turbulence. Observations minimising scintillation are thus best made from high-altitude observatories (above many layers of the atmosphere) where the predominant winds cross over large bodies of water such as the ocean (a very flat topography that radiates heat evenly), and where the weather patterns are relatively stable.

- *Sky Background*: There is always a background brightness level, even when there appear to be no nearby stars. This contamination is due to light diffusion in the atmosphere, caused by airglow or light pollution from neighbouring towns. Airglow is the faint emission of light in the upper atmosphere, due to several different processes. One such process is that atoms that have been photo-ionised by the Sun recombine, releasing photons. Another is that cosmic rays, very high energy protons and nuclei that travel through space, hit molecules in the atmosphere, which causes fluorescence. Additionally, oxygen and nitrogen in the atmosphere react with hydroxyl ions to release light. The only way to reduce airglow is to choose an observing site far from populated areas and at high altitudes, where the atmosphere is thin.
- *Differential Extinction*: Extinction results from absorption and scattering of light by dust and gas. The main components of atmospheric extinction are Rayleigh scattering by air molecules, particulate scattering, and molecular absorption. The tellurics that most strongly absorb incoming light are molecular oxygen, ozone and water. The longer a light ray must travel through the Earth's atmosphere, the more it will be attenuated. This path length is called airmass. Light from a target at zenith will travel through a different airmass to light from stars towards the horizon, so to first-order, they will experience relative brightness differences. There is also a second-order impact, where stars will experience different extinction levels depending on the spectra of their emitted light. Molecular oxygen and ozone absorb radiation strongly in the ultraviolet, whereas water absorbs strongly in the infrared. Therefore, a star emitting most of its light at redder wavelengths will appear fainter than a bluer star if there are significant amounts of water vapour in the Earth's atmosphere at the time. Therefore, by using targets close together on the sky of similar spectral types (see Section 1.3) we can minimise extinction effects. We can also reduce second-order extinction through complex detrending algorithms (as demonstrated in Section 2.2.7) and by using a narrow filter with little or no absorption lines.

  The interstellar medium will also cause extinction, affecting both ground and space-based telescopes. Interstellar dust absorbs and scatters blue light more than red light, causing an artificial "reddening" effect. Therefore, not only ground-based telescopes have to take extinction into account.

From the Telescope

All instruments come with their own sources of error. Instrumental errors can imprint photometric scatter or systematics into data; therefore, it is crucial to understand the detail of how these instruments work. For my research, the only detector I will consider here is a Charge-Coupled Device (or CCD). CCDs are present in most modern photometric telescopes as they are compact and highly efficient at detecting photons.

CCDs rely on semiconductors, most commonly silicon. In an insulator there is a large band gap between valence and conduction bands, so we would need to supply a huge amount of energy to promote an electron into the conduction band and create an electron-hole pair. The result is that insulators have very poor electrical conductivity. For metals, there are no separate valence and conduction bands which is why metals are such good electrical conductors. Semiconductors, with small band gaps, exist in between insulators and conductors. When a semiconductor is illuminated, its electrons can gain enough energy to conduct electricity simply from photons striking it, which why they make excellent photon detectors. The percentage of photons arriving at the CCD that produce photoelectrons determines a detector's quantum efficiency, a useful metric of its quality. The quantum efficiency is wavelength-dependent and is affected by various factors, such as the semiconductor material and anti-reflection coatings.

Even if two telescopes have the same type of CCD installed, they will produce subtly different outputs due to a combination of errors unique to each detector, telescope and method of operation:

- *Quantization Noise*: Each pixel can act as a capacitor storing a variable amount of charge and creating a small analogue voltage across it. The pixels across rows are linked so that when the CCD is "read-out" the charge is transferred step-by-step to the neighbouring pixels in each row. The charge in the end pixels is then fed to an amplifier, where an output analogue voltage is induced, amplified and measured. This is repeated for each column of pixels. This analogue output signal is converted into digital counts by an ADC (Analogue to Digital Convertor) in units of ADU (Analogue-Digital Units). The relationship between the number of electrons and counts is set by the gain, in units of e^{-}/ADU. The larger the gain, the more precision is lost when converting into discrete digital counts, called quantization noise.
- *Readout Noise*: Theoretically, when a photon strikes a pixel on the CCD, it should produce one electron. The detector then records this electron as one count. However, there are a combination of electronic noise sources to take into account when charge is measured, and when that signal is amplified and digitised (see above). If all pixels on a CCD contained the same number of electrons, during the read-out process a physical (and

so imperfect) amplifier would measure charge following a Gaussian distribution centred at the "correct" value, with a standard deviation (or average variation between pixels) given by the readout noise. Because the ADC cannot measure negative values (whereas noise fluctuations can be positive or negative), a constant offset voltage is applied to the capacitor, called the bias level.

- *Dark Current*: Thermal excitation promotes a small but non-negligible number of electrons into the conduction band. These electrons will contribute to the final output signal that is measured. This effect is known as dark current, and these electrons are indistinguishable from photoelectrons. The build-up of thermal electrons can be substantial, so we often cool CCDs to operate at low temperatures to limit the impact of dark current. Additionally, choosing a semiconductor with a larger band gap, like silicon, reduces the relative number of electrons excited by thermal energy compared to photon energy.
- *Pixel inhomogeneity*: Every pixel on a CCD will behave and respond to light marginally differently. Therefore, there are slight interpixel variations in response to the same light source. There are also three particular pixel defects: "dead" pixels that read very low values, "hot" pixels that read very high values, or traps that temporarily hold onto charge during readout. Hot pixels have an enhanced dark current compared to other pixels, usually due to a different temperature response.
- *Flat-fielding Errors*: We expect that if we observed a perfectly uniform light source, we would measure the same number of counts in every pixel across the entire image. Even after considering all the effects described above, each pixel will read a slightly different value. Several other effects combine to prevent us from observing a perfectly "flat field". Dust grains on the CCD window or filters block a fraction of the photons and appear as out-of-focus dark donut shapes. Dust grains can appear, move, or disappear during observation, spatially and temporally affecting the brightness of different parts of the field. The random movement of dust is, therefore, complicated to correct. Another effect is vignetting, where the edges of the image darken due to a loss of sensitivity when observing light off-axis. There will also be some marginal effects from optical imperfections, internal reflections etc., that can cause uneven illumination across the field.
- *Pointing Errors*: Telescopes usually utilise an autoguider or guiding software to ensure a telescope is pointing precisely at the correct source and that it follows that source during the night. If a telescope were static during an exposure, we would see light trails as stars move across the sky. Systematic errors caused by the drift of stars on the CCD (with a spatially inhomogeneous pixel response, as described above) can severely limit the precision of time-series photometry; therefore, fixing stellar positions at the

sub-pixel level is essential. Intentionally defocusing the telescope spreads the counts over more pixels, with the advantage of diluting the impact of pixel-to-pixel variations during slight drifts (to improve accuracy) and preventing bright targets from saturating. However, observing defocused can be a challenge for some guiding software (that rely on accurately determining the centroid of individual guide stars) or autoguiders to accurately point at the target and maintain stability throughout a night, if the defocus is very large, sky background is bright or the target is faint.

- *Saturation*: Saturation occurs when a pixel reaches some maximum value. This could be because each pixel has a maximum charge it can carry, called the full-well capacity. Once this well is full, the charge spills over into adjacent pixels in a process known as blooming. However, at low gain values, the saturation point is often set by the maximum value the ADC can measure. For example, an 8-bit ADC can only represent values between 0 and $2^8 - 1 = 255$.

The majority of these instrumental effects can be mitigated with proper calibration of science images during the data reduction process (as shown in Chapter 2).

1.2.5 Contamination of Planetary Signals from Astrophysical Origins

Searching for planets by detecting small dips in brightness is made much more challenging because stars themselves do not maintain constant brightness. On the timescales of minutes, hours and years, the stellar surface change, causing fluctuations in stellar flux. This stellar variability can complicate planet parameter retrieval or even create or obscure transit features.

Convection in outer stellar layers causes oscillations and granulations with lifetimes of the order of minutes to hours. This rapid variability provides a source of correlated noise into photometric lightcurves; however, the scale of this noise is likely to be far below current levels of photon noise (Sarkar et al. 2018).

The surface of a star does not have uniform brightness. It has magnetically active regions of cooler (and darker) stellar spots and hotter (brighter) stellar faculae. As the star rotates, these regions will come in and out of view, and the brightness will vary on the same timescale as the rotational period. Significant rotational modulations in photometric lightcurves make it difficult to extract a transit feature, especially if the rotation pattern is irregular or if the rotation period is similar to a typical transit duration. These magnetic surface features, both occulted and unocculted, can also affect a planet's transit depth and, hence, estimates of its radius. Spots and faculae unocculted by the planet will affect the baseline measurement of the star's brightness. We assume that the stellar surface is the same before and during the transit. Not only is this surface changing, but also, the planet will eclipse only part of the star's face.

This eclipsed region may look quite different to the rest of the star if it has a higher or lower coverage of stellar spots. This stellar inhomogeneity is also wavelength-dependent (due to the temperature contrast between spots, faculae and the star) which can imprint spectral features onto transmission spectra (Rackham et al. 2018, 2019).

Stellar flares can also complicate planetary detection. Flares are explosive events caused by magnetic recombination in the upper atmosphere of a star (Benz & Güdel 2010). This sudden eruption of magnetic energy usually occurs around active regions, such as stellar spots, and causes bursts of particles and electromagnetic radiation. The spectrum of this radiation resembles a black body with an effective temperature of 9000 K (Shibayama et al. 2013). Often, though not always, powerful flares will come accompanied by a coronal mass ejection (CME) event, where clouds of charged particles are directionally ejected from the star. Unlike magnetic surface features, flares are stochastic events occurring on minute-to-hour timescales that cannot be predicted.

Even when a transit is detected above the level of stellar variability, a few astrophysical false positives can mimic a planet. The following scenarios can create a transit in a photometric lightcurve (Fressin et al. 2013):

- Grazing stellar binaries, where two stars only slightly eclipse one another,
- A planet transiting a background star,
- A background eclipsing binary or an eclipsing binary in a triple system,
- Or a transiting red/brown dwarf.

Following up with different surveys and a range of detection techniques is one of the best ways to confirm a planetary candidate and rule out these false-positive scenarios.

1.3 Stellar Classification

A star's spectrum, or how the intensity of light it emits varies with wavelength, is encoded with information about the star itself. The strengths of different spectral lines, which represent different chemical and molecular abundances, can tell us about the temperature of a star's photosphere. We can then divide stars into groups called spectral types based on similar spectral characteristics.

The most commonly used stellar classification system is called the Morgan-Keenan system (Morgan & Keenan 1973). In decreasing temperature order, this system initially consisted of spectral types O (hottest), B, A, F, G, K, and M (coolest). Cooler stars predominantly emit light at longer wavelengths than hotter stars, so O-type stars mainly emit in the UV, unlike M-type stars that are brightest in the infrared. Each of these classes is further subdivided from 0–9, again in declining temperature order, so that M2 is hotter than M3. Each star is also assigned a roman numeral representing a luminosity class. However, I only consider main sequence stars in this work, given the class V. For reference, our Sun has spectral type G2V.

To define the main sequence, I will first briefly describe how stars form. Stars are made when clouds of gas and dust become too large to balance gas pressure and gravity. As the mass builds, gravity increases, so more and more material is pulled together to form a protostar. With increasing gravity, the material condenses (gravitational collapse), heating up as it does so. As the gas gets hotter, the atoms become excited and move more quickly, creating friction that generates more heat in a runaway process. Once temperatures are hot enough, hydrogen nuclei fuse to produce helium which generates enough energy to sustain this process and prevents further gravitational collapse. A star is in the main sequence when it has reached this equilibrium hydrogen-burning stage. When a star has depleted a significant amount of the hydrogen in its core, it will evolve off the main sequence. The timescale for this process (majority of a star's lifespan) and the amount of hydrogen available are dependant on the mass of the star. More massive stars will generally burn through their hydrogen supplies much faster than lower mass stars, giving them shorter main sequence lifespans. After the main sequence, stars can follow several evolutionary tracks depending on their mass. As the objects I consider in my research are ultra-cool dwarfs with main sequence lifespans of at least several hundred billion years (much longer than the age of the universe, Adams et al. 2005), I will not consider the evolution of these objects after leaving the main sequence in this thesis. See Salpeter et al. (1955) for more on main sequence evolution, and Adams et al. (2005) for the long term evolution of the coolest stars.

Since its creation, this spectral classification system has been extended to include several new spectral types, including the three spectral types L, T, and Y, specifically for the coolest sub-stellar objects that radiate primarily in the infrared.

1.3.1 M-dwarfs

M-dwarfs are defined as a spectral class spanning stellar types M0V–M9V, with temperatures between 2270–3850 K (Pecaut & Mamajek 2013). In the literature, the terms "low-mass stars", "red dwarfs" and "cool stars" are often used interchangeably to refer to this class of objects (Edgeworth 1946). This ambiguity around nomenclature can be confusing. The 2270–3850 K temperature range for M-dwarfs also includes sub-stellar objects, such as brown dwarfs (e.g., Rebolo et al. 1995) which are defined in Section 1.3.3, and overlaps with a grouping of cooler objects called ultra-cool dwarfs (UCDs, defined in Section 1.3.2). I will use these vague terms to more generally refer to any stellar or sub-stellar object with a spectral type later than M0. However, when necessary for clarity, I will divide M-dwarfs into three subgroups: early M-dwarfs, mid M-dwarfs and ultra-cool dwarfs. Early M-dwarfs will include stars of spectral types M0–M3, or with temperatures in the range 3210–3850 K. Mid M-dwarfs will include stars of spectral types M4–M6, or with temperatures in the range 2700–3210 K. Finally, the latest M-dwarfs fall into the category of ultra-cool dwarfs, as objects with spectral type M7 and later, or with temperatures cooler than 2700 K.

1.3.2 Ultra-Cool Dwarfs

Ultra-cool dwarfs are classically defined as objects with spectral types M7 and later, or with effective temperatures cooler than 2700 K (Kirkpatrick et al. 1995, 1999; Kirkpatrick 2005). These objects are extremely low-mass ($<0.1\,M_\odot$) and Jupiter-sized ($<0.12\,R_\odot$). The lifetimes of UCDs are several orders of magnitude longer than the Sun (Reid & Hawley 2006). Due to their very low temperatures, UCDs emit mostly in the infrared and, therefore, appear much redder than hotter stars.

1.3.3 Brown Dwarfs

UCDs consist of not only stellar but also sub-stellar objects not massive enough to sustain hydrogen fusion in their cores, such as brown dwarfs ($\sim$13–80 M_{Jupiter}). Brown dwarfs straddle the boundary between gas giant planets and stars. They begin life in a similar way to stars (see Section 1.3). However, the mass of a brown dwarf is too small to reach the temperature required to sustain hydrogen fusion, and it remains a ball of gas. Brown dwarfs, therefore, undergo a vastly different formation mechanism to gas giant planets (which begin as rocky cores that accrete gas and dust in a star's protoplanetary disc), and so while they may be similar in size, they have very different compositions, metal content and temperatures.

1.4 Ultra-cool Dwarfs as Planetary Hosts

Almost three decades after the first discovery of an exoplanet, the ultimate goal of planet detection is still to find an Earth-like planet orbiting in the habitable zone of a Sun-like star. However, until PLATO, this discovery likely will remain out of reach of our detection capabilities. Additionally, for the next decade, it will not be possible to study the atmospheres of these objects in detail. If we broaden our focus to searching for temperate Earth-sized planets around any star, this opens a compelling case for our coolest neighbours, ultra-cool dwarfs (UCDs).

The advantages of searching for planets around UCDs are numerous, as demonstrated in Figure 1.5. The transit depth scales with the ratio between the planet and host; therefore, for stars with sizes $\sim$0.1$R_\odot$ the eclipses of an Earth-sized planet would result in 1% dips in

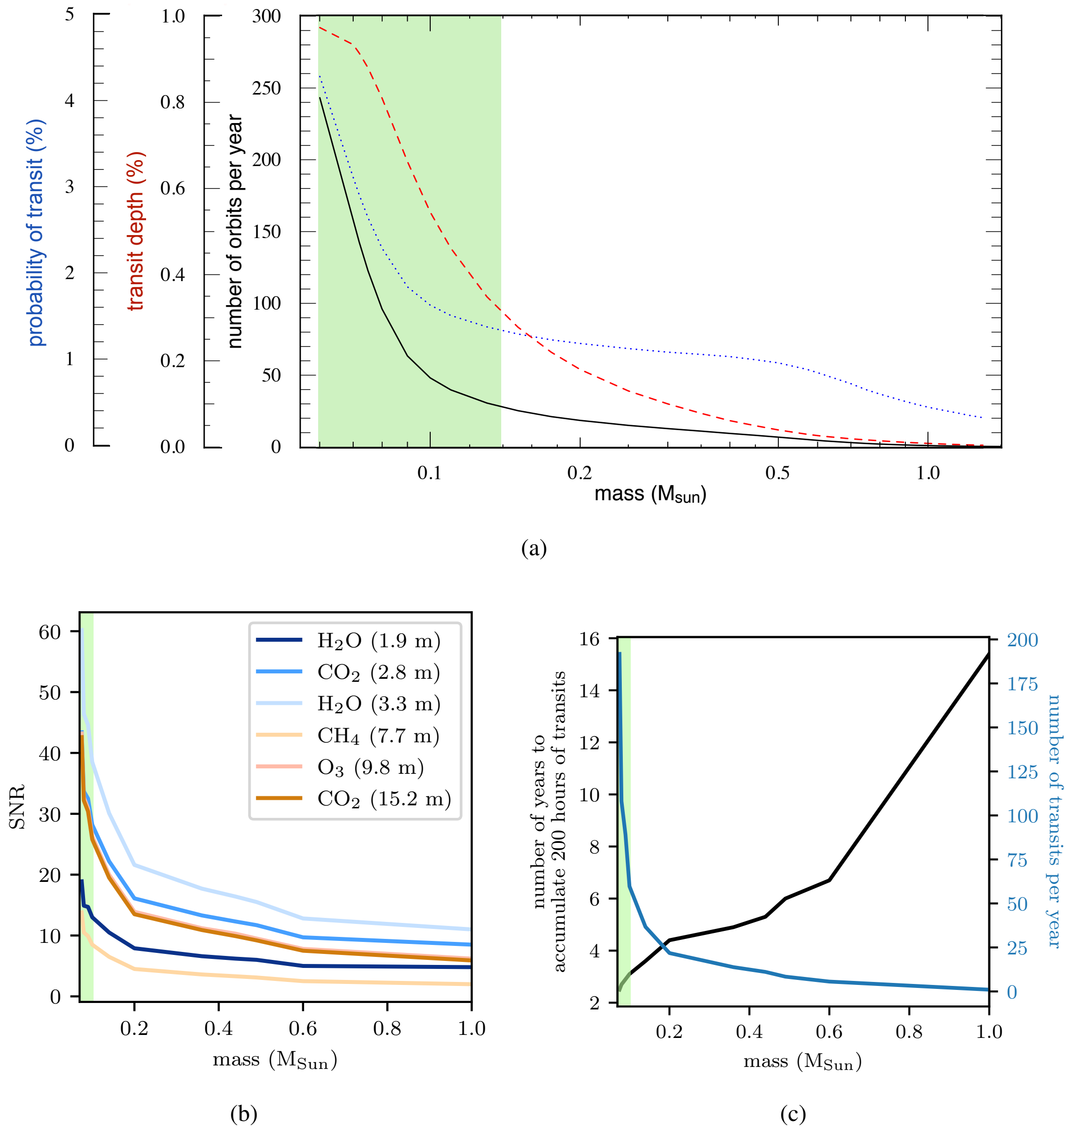

Figure 1.5: (a) Figure from He et al. (2017) showing the relationship between stellar mass and the geometric probability of observing an eclipse (dotted blue, see Equation 1.10), the transit depth (dashed red) and the number of orbits per year (black) for an Earth-sized planet with an equilibrium temperature of 255 K. The green highlighted region represents where spectral features can be explored with a signal-to-noise ratio ≥5. We see a sharp increase in transit depth, probability of observing a transit and number of transits per year for objects with masses <0.1 $M_\odot$, UCDs. (b) The signal-to-noise ratios (SNR) derived by Kaltenegger & Traub (2009) for several major spectroscopic features in the transmission spectrum of a transiting Earth, observed for a summed transit observation time of 200 hours with a 6.5 m space-based telescope (such as JWST). The shaded green region here marks the ultra-cool objects with masses ≤0.1 $M_\odot$. (c) The number of transits per year on average for an Earth-sized, temperate planet for different stellar hosts, and the number of years it would take to capture 200 hours of transits.

brightness (Equation 1.17). These transit depths are two orders of magnitude greater than those for planets around a Sun-like star and, therefore, much easier to detect. Additionally, the habitable zones around UCDs shrink inwards to periods of a few days, resulting in very frequent transits. There is also a much higher geometric probability of observing transits (Equation 1.11) because habitable zone planets are much closer to their host (~0.01 times the separation between the Sun and the Earth). The probability of observing a temperate, terrestrial planet's transit, therefore, rises from ~0.5% when we consider stars equivalent to the Sun, to ~3-5% for objects with masses less than 0.1 $M_\odot$. There is also a wealth of cool stars in our local stellar neighbourhood, as 61% of stars within 10 pc are M-dwarfs (Chabrier 2003; Reid & Cruz 2002; Henry et al. 2006; Kirkpatrick et al. 2012; Henry et al. 2018; Reylé et al. 2021). M-dwarfs are much more likely to host planets than Sun-like stars, and by extrapolation, UCDs may follow the same trend (Howard et al. 2012; Morton & Swift 2014; Dressing & Charbonneau 2015; Mulders et al. 2015; Gaidos et al. 2016; Hsu et al. 2020; Yang et al. 2020; Sabotta et al. 2021). Due to their low luminosities and small sizes, the detection of spectroscopic signatures in the atmosphere of a rocky, habitable zone planet becomes more favourable for low-mass stars than any other type of host star (Kaltenegger & Traub 2009; Seager et al. 2009; de Wit & Seager 2013). The large number of transits we could observe in a relatively short amount of time (see Figure 1.5) also facilitates atmospheric characterisation. In summary, not only are Earth-sized planets easier to detect around cool stars, but there are many more places to look, and the planets that we do find will be the best candidates for atmospheric characterisation by future telescopes, such as JWST.

Historically, searching for planets around UCDs has been instrumentally out of reach. UCDs are very red objects, meaning that they emit most of their light at near-infrared wavelengths. Therefore, UCDs are extremely faint objects in the visible, requiring either large telescopes or optimised near-infrared/infrared detectors to gather enough photons to obtain the necessary signal-to-noise ratio to detect even a ~1% transit. Surveys with relatively small (<0.5 m) telescopes, such as MEarth and TESS, have demonstrated a high detection potential for early and mid M-dwarfs that drops sharply for later-type objects (Sullivan et al. 2015; Barclay et al. 2018). The redness of these targets also poses a particular challenge for ground-based telescopes. The Earth's atmosphere contains several telluric absorption lines in the infrared, meaning that UCDs are more affected by differential extinction than other types of stars. The effect of second-order extinction placed a severe limitation on the type of objects that the MEarth survey could target (Berta et al. 2012a).

1.4.1 Seizing the "M-dwarf Opportunity"

Several planet detection surveys have focused on early and mid M-dwarfs in the past decade. Targeting these types of stars is a compromise between the precision constraints that surveys experience when observing extremely red objects, and the potential to detect temperate, terrestrial planets. Early-to-mid M-dwarfs can still provide a shortcut to detecting small (though larger than Earth-sized), close-in planets that would otherwise be out of reach. This "M-dwarf opportunity" (Deming 2008) was the driving factor behind the MEarth survey and had a significant impact on the development of the K2 and TESS missions.

The advantages of cool stars also extend to radial velocity surveys due to an increased ratio of planet to stellar mass. HARPS (High Accuracy Radial Velocity Planet Searcher, Pepe et al. 2000, 2002; Mayor et al. 2003) consists of two spectrographs, one based at ESO's La Silla Observatory in Chile[w] and the other in the Canary Islands (HARPS-North)[x], both performing high precision radial velocity measurements. From 2003 to 2009, HARPS performed a specialised survey (Bonfils et al. 2013) of 101 nearby early and mid M-dwarfs. This survey resulted in several detections of potentially habitable super-Earths around low-mass stars (Dittmann et al. 2017; Delfosse et al. 2013), including around our nearest neighbour Proxima Centauri (Anglada-Escudé et al. 2016). The more recent and more red-sensitive CARMENES[y] (Calar Alto high-Resolution search for M-dwarfs with Exoearths with Near-infrared and optical Echelle Spectrographs, Quirrenbach et al. 2010, 2014; Reiners et al. 2018) has been performing a radial velocity survey of over 350 M-dwarfs since 2016 from the Calar Alto observatory in Spain. While the CARMENES target list spans the entire M-dwarf spectral range, they predominantly focus on M3–M4 stars and include only ~10 UCDs (Quirrenbach et al. 2016; Jeffers et al. 2018; Quirrenbach & Consortium 2020). CARMENES has seen similar success to HARPS, already having detected 30 planets around cool stars (e.g., Luque et al. 2018; Ribas et al. 2018; Morales et al. 2019), and confirmed 17 more, in its five-year lifespan (Nowak et al. 2020; Trifonov et al. 2021). Since its first light in 2018, SPIRou[z] (SpectroPolarimètre InfraRouge, Artigau et al. 2014; Donati et al. 2018), a spectro-polarimeter installed on the Canada-France-Hawaii telescope, has been performing near-infrared radial velocity measurements to detect habitable super-Earths around mid M-dwarfs. In addition to these three projects, there are many other ongoing and planned RV surveys hunting for terrestrial exoplanets around early and mid M-dwarfs not mentioned here (Mahadevan et al. 2014; Kotani et al. 2014; Tozzi et al. 2016; Bouchy et al. 2017).

With statistically significant samples of early-to-mid M-dwarf planets from the Kepler,

[w] `http://www.eso.org/sci/facilities/lasilla/instruments/harps.html`
[x] `https://plone.unige.ch/HARPS-N/`
[y] `https://carmenes.caha.es/`
[z] `https://www.cfht.hawaii.edu/en/projects/SPIRou/`

Table 1.1: Table of the 12 *bona fide* planets around confirmed ultra-cool dwarfs.

Name	Host	Spectral Type	Minimum Mass ($M_\oplus$)	Radius ($R_\oplus$)	Discovery Paper
MOA-2007-BLG-192Lb	MOA-2007-BLG-192L	M7	3.2	–	Bennett et al. (2008)
OGLE-2016-BLG-1195Lb	OGLE-2016-BLG-1195L	M7	1.43	–	Shvartzvald et al. (2017); Bond et al. (2017)
TRAPPIST-1b	TRAPPIST-1	M8	1.017	1.121	Gillon et al. (2016, 2017)
TRAPPIST-1c	TRAPPIST-1	M8	1.156	1.095	Gillon et al. (2016, 2017)
TRAPPIST-1d	TRAPPIST-1	M8	0.297	0.784	Gillon et al. (2016, 2017)
TRAPPIST-1e	TRAPPIST-1	M8	0.772	0.910	Gillon et al. (2016, 2017)
TRAPPIST-1f	TRAPPIST-1	M8	0.68	1.045	Gillon et al. (2016, 2017)
TRAPPIST-1g	TRAPPIST-1	M8	1.148	1.148	Gillon et al. (2016, 2017)
TRAPPIST-1h	TRAPPIST-1	M8	0.331	0.773	Gillon et al. (2016, 2017)
Teegarden's Star b	Teegarden's Star	M7	1.05	–	Zechmeister et al. (2019)
Teegarden's Star c	Teegarden's Star	M7	1.11	–	Zechmeister et al. (2019)
TVLM 513b	TVLM 513-46546	M9	111–134	–	Curiel et al. (2020)

MEarth, TESS, HARPS and CARMENES surveys, it is possible to extrapolate their results to derive the first estimates of occurrence rates of planets orbiting these low mass stars. Mulders (2018) collated a number of these occurrence rate studies (Youdin 2011; Howard et al. 2012; Dong & Zhu 2013; Kopparapu et al. 2013; Petigura et al. 2013; Morton & Swift 2014; Mulders et al. 2015; Silburt et al. 2015; Ballard & Johnson 2016; Gaidos et al. 2016; Hsu et al. 2019, 2020) and found that they were in good agreement, predicting between 1 and 2 small planets (1–4 $R_\oplus$) per M-dwarf, with periods less than 50 d. Therefore, we find that terrestrial planets are abundant around M-dwarf stars, which explains the high yields of recent M-dwarf surveys. However, for now, the potential bounty of temperate Earth-twins around early and mid M-dwarf hosts exists just slightly out of reach.

1.4.2 The Planets Around our Coolest Neighbours

Despite the pronounced success of planet searches around early and mid M-dwarfs, very few of these surveys have been directed towards ultra-cool dwarfs. The TRAPPIST and SPECULOOS transit projects (discussed in more detail in the following chapter) are exceptions. These surveys have demonstrated the ability to reach photometric precisions necessary to detect terrestrial planets (Murray et al. 2020) and discovered the first planetary system around an ultra-cool star (Gillon et al. 2016, 2017). The success of the TRAPPIST and SPECULOOS surveys have also inspired a number of new UCD-focused transit surveys, such as the EDEN Project[aa] (Exoearth Discovery & Exploration Network, Gibbs et al. 2020) and others currently in development (García-Mejía et al. 2020; Tamburo & Muirhead 2019).

Due to the lack of planet searches aimed at UCDs, we only know of 12 planets orbiting UCDs to date, compiled in Table 1.1. The seven TRAPPIST-1 planets remain the only known

[aa] `http://project-eden.space/`

transiting planetary system around an UCD. However, there have also been successes with other detection methods. Recently, the CARMENES team made the first RV detection around an UCD (Teegarden's Star, a nearby M7-type star), announcing the discovery of two Earth-like planets (Zechmeister et al. 2019). Furthermore, only last year, the Saturn-like TVLM 513b, the first planet found orbiting an UCD through radio astrometry, was detected around the M9-type object TVLM 513-46546 (Curiel et al. 2020). Microlensing is also a powerful tool for probing ultra-cool dwarfs because it does not depend on the brightness of a host star (Gould & Loeb 1992). There have been two planets detected around UCD hosts during microlensing events; OGLE-2016-BLG-1195Lb (Bond et al. 2017; Shvartzvald et al. 2017) and MOA-2007-BLG-192Lb (Bennett et al. 2008; Gould et al. 2010; Kubas et al. 2012).

However, several additional planetary-mass objects have been detected orbiting UCDs, accompanied by serious concerns about their planet status. Detections from the microlensing technique are often limited by a lack of information about the host and lens objects. Some microlensing planet detections have been around hosts with ambiguous masses, where it is unclear if they are, in fact, UCDs (Sumi et al. 2016; Jung et al. 2018b; Miyazaki et al. 2018; Hwang et al. 2018; Ryu et al. 2019; Kondo et al. 2019; Zhang et al. 2020; Han et al. 2020). There have also been a couple of microlensing detections of giant planets around UCDs with small projected separations and mass ratios with their hosts (Han et al. 2013, 2016; Jung et al. 2018a). Close orbital distances hint at formation within the protoplanetary disc, but small mass ratios imply a binary-like formation mechanism through gravitational fragmentation (Lodato et al. 2005; Gaudi 2010; Shvartzvald et al. 2017); therefore, I do not include these potential "sub-brown dwarfs" ($<13\,M_{\mathrm{Jupiter}}$) in the count of UCD planets. Direct imaging has also resulted in several discoveries, such as 2M1207b (Chauvin et al. 2004), 2M044144b (Todorov et al. 2010), CFBDSIR J1458+1013b (Liu et al. 2011), VHS J1256-1257b (Gauza et al. 2015) and Oph 98b (Fontanive et al. 2020). However, all these directly imaged candidates have masses much greater than Jupiter and straddle the boundaries between giant planets, sub-brown dwarfs and brown dwarfs with no consensus of their planetary or stellar nature.

UCDs have been sparsely studied and, as a result, are poorly understood. This extends to their planetary populations: we know almost nothing about the frequency and diversity of the planets that orbit UCDs (Delrez et al. 2018b). While we predict that early and mid M-dwarfs are teeming with planets (see Section 1.4.1), we still know very little about whether or how these predictions might extend to ultra-cool objects. However, the discoveries of the planetary systems around TRAPPIST-1 and Teegarden's Star from very small UCD samples may suggest that UCDs host an abundance of multi-planet systems. A historical lack of planet-hunting surveys around UCDs has led to essentially unbounded planetary occurrence rates and, therefore, a broad range of predictions of their populations in the literature (Demory et al.

2016b; He et al. 2017; Sestovic & Demory 2020; Sagear et al. 2020; Lienhard et al. 2020). It follows that we can predict only limited accuracy detection yields for UCD surveys and conclude very little about the (potentially billions of) temperate, Earth-sized worlds hosted by our coolest neighbours.

1.4.3 Life around an Ultra-Cool Sun?

Despite their promise, serious questions remain about the habitability of planets around red dwarfs. For a start, planets in the habitable zone would orbit so close to their parent star that they would likely be tidally locked (Kasting et al. 1993; Barnes et al. 2008). The result of tidal locking is that the same side of the planet faces the star throughout the orbit. Hence one side of the planet will be in constant light and the other in perpetual darkness. This dichotomy could cause extreme temperature variations between the day and night sides of the planet and produce an unstable atmosphere. Though heat circulation, such as through ocean currents and winds, would play an influential role in these scenarios (Yang et al. 2014), the resulting environment could pose serious challenges for life trying to evolve.

UCDs are also especially active objects (West et al. 2015; Williams et al. 2015; Gizis et al. 2017; Paudel et al. 2018; Günther et al. 2020b), producing energetic stellar flares and large-scale photometric modulations, which results in treacherously variable conditions for their planets. Bombardment from flares could strip a planet's atmosphere (Lammer et al. 2007) and cause ozone depletion (Segura et al. 2010; Tilley et al. 2019). Several authors also find it unlikely that these extremely cool, red stars would provide their planets with enough UV photons to initiate specific prebiotic chemistry pathways (Rimmer et al. 2018) or enough visible light for photosynthesis to occur (Mullan & Bais 2018; Lehmer et al. 2018; Covone et al. 2021). On the other hand, it has been theorised that the frequent flares from active UCDs could deliver enough energy to their planets to drive these photochemical processes (Buccino et al. 2007; Ranjan et al. 2017; Rimmer et al. 2018; Mullan & Bais 2018; Lingam & Loeb 2019a), assuming the planetary atmospheres can withstand this high level of activity.

The dynamic relationship between a planet and its host is a prominent factor in evaluating how conducive a planet's environment is for life. However, for UCDs, this relationship is poorly constrained.

1.5 Conclusions

At the moment, we have only scratched the surface of the investigation of ultra-cool dwarfs and their planets. The ongoing SPECULOOS and TRAPPIST projects, and future UCD-focused transit surveys being developed (García-Mejía et al. 2020; Tamburo & Muirhead 2019;

Gibbs et al. 2020), will provide photometric measurements of these objects with unparalleled precision. Not only are these surveys predicted to detect many more temperate, terrestrial planets in the years to come (Sebastian et al. 2020), but monitoring these ultra-cool objects over long periods can help to shed some of the mystery surrounding their long-term behaviour. There remain many unanswered questions about UCDs, including how they generate magnetic fields, the architecture of their planetary populations, and whether life could originate or be sustained on their planets.

A central goal of exoplanet detection is to help us put the Solar System into context and address fundamental questions about life in the universe. With the rapid expansion of the field in the past 30 years and the giant leaps that both ground surveys and space missions will make within the next few decades, it seems that, at last, we may get some answers.

CHAPTER 2

THE SEARCH FOR HABITABLE PLANETS ECLIPSING ULTRA-COOL STARS (SPECULOOS)

The SPECULOOS (The Search for habitable Planets EClipsing ULtra-cOOl Stars) project aims to detect single transits from Earth-sized planets orbiting ultra-cool dwarfs. This ambitious task requires photometric precisions of at least ~0.1%. Therefore, to obtain the necessary high signal-to-noise ratio (SNR) lightcurves, and to deal with the specificity of our very red targets, I developed a specialised automatic pipeline to process and reduce the data from the SPECULOOS-South Observatory (SSO).

The content of this chapter is based heavily on the work presented in Murray et al. (2020). This chapter details the SPECULOOS project (Section 2.1), and the evolution of the SPECULOOS-South (SSO) Pipeline, from conception to completion. In Section 2.2, I outline each of the reduction, astrometric and photometric methods used by the SSO Pipeline. This includes an in-depth discussion of the development of a novel differential photometry technique, in Section 2.2.7. This pipeline also implements a correction for the effects of variable telluric water vapour absorption on photometry, described in Section 2.4. The research into the precipitable water vapour effect was carried out in collaboration with a fellow PhD student, Peter P. Pedersen. In particular, he was of great help in generating the water vapour grid in Section 2.4.1.

2.1 The SPECULOOS Project

The aim of the SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars) project is to search the nearest (within 40 pc) ultra-cool dwarfs for transiting, terrestrial planets. These planets will be prime candidates for future atmospheric characterisation with the James Webb Space Telescope (JWST) and extremely large telescopes. This project is led by the University of Liège (PI Michaël Gillon) and carried out in partnership with the University of Cambridge, University of Birmingham, Massachusetts Institute of Technology, University of Bern, Canary Islands Institute of Astrophysics and the European Southern Observatory.

The scientific case for SPECULOOS was demonstrated by a prototype survey carried out on the TRAPPIST-South telescope (described in Section 2.1.1). The SPECULOOS project itself is built on a global network of telescopes, outlined in Section 2.1.2. The target catalogue of ultra-cool objects SPECULOOS plans to survey are defined in Sebastian et al. (2021), and described in Section 2.1.3. I discuss SPECULOOS's observing strategy in Section 2.1.4 and how SPECULOOS has already contributed to exoplanetary studies in Section 2.1.5.

2.1.1 The TRAPPIST Prototype Survey

The motivation behind the TRAPPIST (TRAnsiting Planets and PlanetesImals Small Telescope, Jehin et al. 2011; Gillon et al. 2011) project was primarily to detect and characterise transiting exoplanets, and secondly to study comets and other small Solar System bodies. TRAPPIST began scientific operations in 2010, with one 0.6m robotic telescope situated at La Silla Observatory in Chile. In 2016 an identical 0.6m telescope, TRAPPIST-North, was built at the Oukaïmeden Observatory, in Morocco (Barkaoui et al. 2017, 2019). The project is led by the University of Liége (Belgium), and carried out in close collaboration with the Geneva Observatory (Switzerland).

As part of their exoplanet program, the TRAPPIST team performed photometric monitoring of a small number of the brightest ultra-cool dwarfs in the Southern hemisphere. This program, the Ultra-Cool Dwarf Transit Survey (UCDTS, Gillon et al. 2013), was aimed at detecting transiting planets and studying stellar variability. UCDTS also acted as a prototype for the planned ground-based transit survey, SPECULOOS, whose goal would be to perform a much more extensive search for transiting terrestrial planets around ~1000 UCDs. Due to the relative faintness of UCDs in the visible and presence of atmospheric absorption and emission lines in the infrared, the viability of such a survey was in question. However, in 2016 the TRAPPIST and Spitzer teams announced the discovery of seven Earth-sized, temperate planets transiting the M8-type star TRAPPIST-1 (Gillon et al. 2016, 2017). The success of UCDTS and promising nature of the TRAPPIST-1 system solidified the scientific case for a larger project, such as

SPECULOOS, focused on planetary detection around UCDs.

2.1.2 SPECULOOS Global Network

SPECULOOS is based on a network of identical, robotic telescopes spread across the Northern and Southern hemispheres. The largest facility, the SPECULOOS-Southern Observatory (SSO, Delrez et al. 2018b; Jehin et al. 2018; Murray et al. 2020) in Cerro Paranal, Chile, is composed of four telescopes (see Figure 2.1). These telescopes are named Io, Europa, Callisto and Ganymede[a]. Additionally, the SPECULOOS Northern Observatory (SNO, Delrez et al. 2018b; Niraula et al. 2020) has one telescope at the Teide Observatory in Tenerife (Canary Islands), and SAINT-EX (Search And characterIsatioN of Transiting EXoplanets, Demory et al. 2020) has one telescope at the San Pedro Mártir observatory in Mexico. As of 2019, all three SPECULOOS facilities were fully operational; the SSO began its official scientific schedule in January, followed by SAINT-EX in March and the SNO in June.

In addition, the two 60 cm TRAPPIST robotic telescopes of the University of Liège (one in Chile, the other in Morocco; Gillon et al. 2011; Jehin et al. 2011), while not being officially parts of the SPECULOOS network, devote a fraction of their time to supporting the project, focusing on the brightest targets.

The SPECULOOS-South Facility

The SSO consists of four robotic 1 m Ritchey–Chretien telescopes (see Figure 2.2), each equipped with a deeply depleted CCD detector which is optimised for the near-infrared. For the vast majority of our observations we use the $I + z'$ custom-designed filter (transmittance >90% from 750 nm to beyond 1000 nm) due to the faintness of our red targets in the optical wavelength domain. However, we are limited beyond 950 nm by the quantum efficiency of our CCD detector. We are also limited on the bluer end of the spectrum by the transmittance of the CCD window itself, which blocks all wavelengths below ~400 nm. Further technical information is shown in Table 2.1 and described in more detail in Delrez et al. (2018b). The overall system efficiency is determined by accounting for the CCD window transmission, the CCD's quantum efficiency, the atmospheric transmission, the reflectance of the mirrors and lenses, and the filter transmission curves. The spectral transmission of the SPECULOOS cameras, with all these various components, is shown in Figure 2.3, which corresponds to Figures 6 and 7 in Delrez et al. (2018b).

[a]The SSO telescopes are named after the four Galilean moons. Firstly, because this Jovian system mirrors the size ratio between Earth-sized planets and their UCD host. Secondly, as a tribute to the first objects discovered to orbit a body other than the Earth, which challenged the geocentric Ptolemaic model of the time.

(a)

(b)

Figure 2.1: Photographs of the SPECULOOS-South Observatory (SSO) in ESO Paranal, Chile. In (a) the four telescopes, Callisto, Ganymede, Europa and Io are shown from left to right. Photograph (b) shows three of the telescopes in operation at night, and the Very Large Telescope (VLT) lasers in the background.

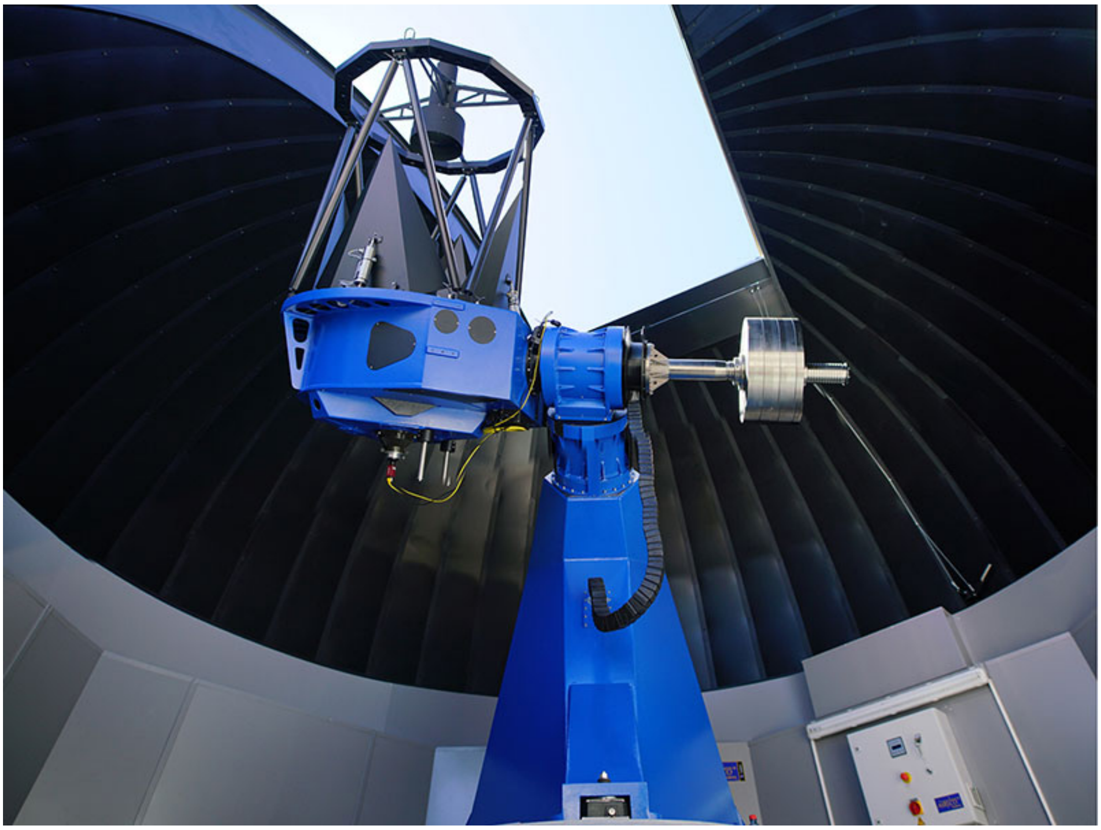

Figure 2.2: Photograph of one of the SSO telescopes (Europa) from inside the dome.

2.1.3 SPECULOOS Target List

The SPECULOOS target list is defined in Sebastian et al. (2021). To generate the target list, a catalogue was built initially from all M and L-dwarfs in the Gaia Data Release 2 catalogue (Gaia Collaboration et al. 2016a, 2018), within 40 pc and with a trigonometric parallax of $d \geq 25$ mas. This list is filtered to remove objects that are too bright, or with spectral types earlier than ~K9. The spectral types are calculated from stellar parameters, derived using distances and magnitudes in various filters from Gaia. The remaining objects are then cross-matched with the 2MASS catalogue (Skrutskie et al. 2006). Three separate calculations of effective temperature are made for each object based on empirical relationships between magnitudes, distances and temperatures[b]. Only the objects which have closely-matched positions and effective temperatures between the Gaia and 2MASS catalogues are kept in the

[b]The first is an estimate from Pecaut & Mamajek (2013) using the apparent G-band (Gaia-band) magnitude and the Gaia distance. The second uses another relationship in Pecaut & Mamajek (2013) based instead on the magnitude difference between the G and H wavelength bands. The third uses the relationship derived in Filippazzo et al. (2015), which is based on the absolute magnitude in the H-band. This magnitude can be computed from the apparent H-band magnitude measured by 2MASS and the distance from Gaia.

Table 2.1: Technical specifications of each telescope in the SSO

	Specification
Mirrors	1 m diameter primary with a $f/2.3$ focal ratio and 28 cm diameter secondary. Combined $f/8$ focal ratio. Both mirrors are coated with pure aluminium.
Camera	Andor iKon-L thermoelectrically-cooled camera
CCD Detector	Near-IR-optimized deeply-depleted $2k \times 2k$ e2v CCD detector
CCD Quantum Efficiency	$\sim$350 (near-UV) to $\sim$950 nm (near-IR). Peak quantum efficiency of 94% at 740 nm.
Field of View	12×12 arcmin^2
Pixel Scale	0.35 arcsec pixel^{-1}
Pixel Size	13.5 µm
Dark Current	$\sim$0.1 e$^-$ s^{-1} pixel^{-1} when the camera is operated at -60°C.
Readout Mode	Usually 1MHz readout mode with a pre-amplifier gain of 2 e$^-$ ADU^{-1} providing readout noise of 6.2 e$^-$
Gain	1.04 e$^-$ ADU^{-1}
Filter Wheel	Finger Lakes Instrumentation (model CFW3-10) allowing ten 5 x 5 cm filters.
Filters	All telescopes: Sloan-g', -r', -i', -z', $I + z'$,'blue-blocking' filters. Selected telescopes: broad-band Johnson–Cousins B, R_C and V filters, and the Sloan u' filter.

target list at this stage. Several missing M and L-dwarfs were added to this list using the catalogue in Bardalez Gagliuffi et al. (2019). The spectral type is estimated from Filippazzo et al. (2015) and is used to calculate the radius of each target using the Stefan-Boltzmann law and its mass using the relationship in Mann et al. (2019). Objects with radii that were too small ($R < 0.07R_\odot$, Dieterich et al. 2014), densities too small ($M < 0.2M_\odot$ and $R > 0.4R_\odot$), masses too large ($M > 0.125M_\odot$) or that were flagged as close binaries or suspected to be an incorrect cross-match, were discarded. The SPECULOOS target list is then taken as a subset of these $\sim$14,000 objects, that divide into three distinct scientific programs, outlined in the following paragraph. See Sebastian et al. (2021) for a more detailed description of the various stages involved in constructing the SPECULOOS target list.

The SPECULOOS target list is formed of 1658 nearby objects in three complementary programs. Program 1 focuses on the 366 dwarfs (including TRAPPIST-1) that are small and nearby enough to allow the detailed atmospheric characterisation of an Earth-sized temperate planet with JWST. Program 2 capitalises on the synergy between SPECULOOS and TESS by targeting the 171 mid M-dwarfs (with spectral types M5-M6.5) for which a significant (>5 sigma) detection of an Earth-sized temperate planet is within reach of TESS. Finally, Program 3 includes the remaining 1121 mid-to-late M-dwarfs that allow us to explore the planet occurrence

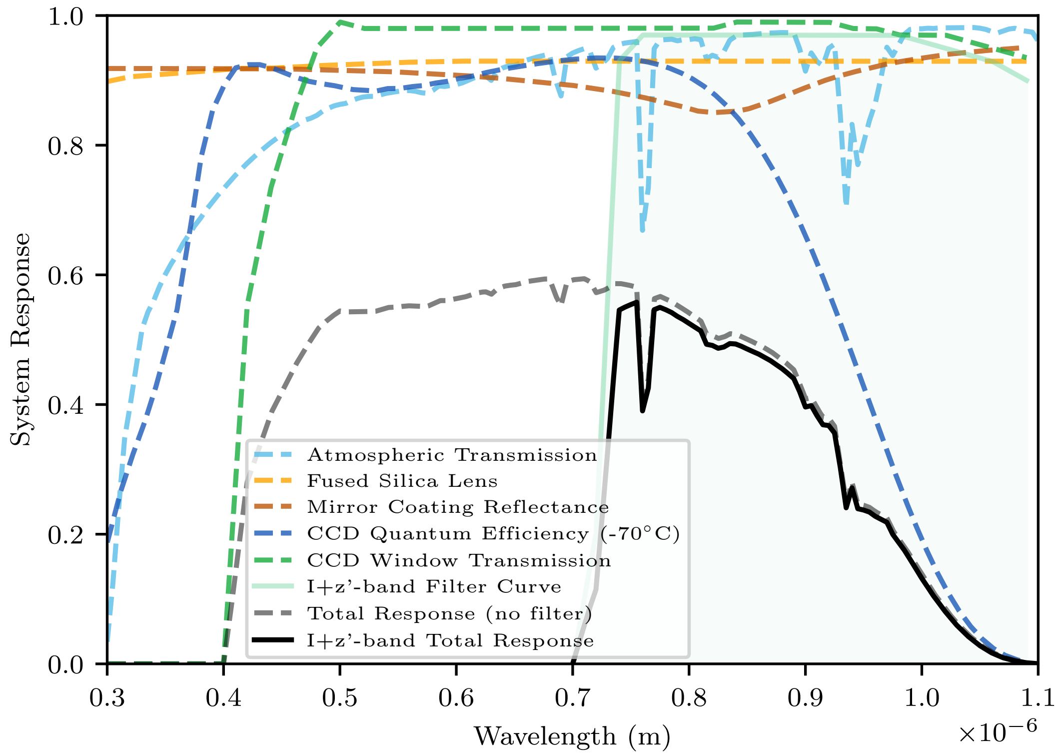

Figure 2.3: The various instrumental and atmospheric factors that combine to produce the overall system efficiency of SPECULOOS. This includes the atmospheric transmission in light blue (for an airmass of 1 and precipitable water vapour of 2.5 mm) from ESO's SkyCalc Sky Model Calculator (Jones et al. 2013; Noll et al. 2012), the CCD quantum efficiency in dark blue, the CCD window transmission in green, the reflectance of the primary and secondary mirrors in red and the transmission through the fused silica lenses in orange. The $I + z'$ filter transmission curve is shown by the shaded region in light green. The grey dotted line is the total response, with no filter in place, whereas the black solid line is the total response in the $I + z'$ filter.

rate for UCDs within our 40 pc sample.

2.1.4 Observation Strategy

Observations on the four telescopes are started remotely each night. Each telescope operates independently and in robotic mode following plans written by SPECULOOS's automatic scheduler, SPeculoos Observatory sChedule maKer (SPOCK, Sebastian et al. 2021). On average, 1–2 targets are observed by each telescope per night. Each target will be observed continuously for between several hours and an entire night (for however long the weather permits and the target is observable). Typically we monitor each target for 100–200 hours, depending on its SPECULOOS program, to efficiently probe different regions around that object (Sebastian et al. 2021).

The observing strategy of SPECULOOS diverges into the three programs described in the previous section. We aim to sample more of the temperate zone around Program 1 targets than around Programs 2 and 3, so the observing durations for Program 1 targets are typically longer. For most of Program 1's targets, habitable-zone planets would have periods of 8–10 days. To obtain an estimated ≥80% phase coverage of the habitable zone would require ~200 hours. For Programs 2 and 3, we are more focused on detecting short-period planets like TRAPPIST-1b (with the assumption that planets receiving 4 times the irradiation of Earth are temperate) and monitoring each target for ∼100 hours would provide an effective phase coverage of ≥80% for close-in planets (with periods ≤6 days). Therefore, our strategy is to observe each of the targets in Program 1 for 200 hours and those in Program 2 and 3 for 100 hours.

As this is a targeted survey and our targets are spread over the sky, there is typically only one target per field of view. During operation, each telescope uses the auto-guiding software, DONUTS (McCormac et al. 2013), instead of a traditional autoguider, to calculate guiding corrections directly from science images and to re-centre the telescope pointing between exposures. The first science image is used as a reference image, from which guide corrections are calculated. These corrections for each image are translated from X/Y shifts to telescope coordinates and sent to the mount to alter the pointing of the telescope in real time (this process occurs during observations and is repeated between each science exposure). Systematic errors caused by the drift of stars on the CCD (with inhomogeneous pixel response) can severely limit the precision of time-series photometry; therefore, fixing stellar positions at the sub-pixel level is essential (≤ 0.2 pixels with DONUTS, McCormac et al. 2013). DONUTS is also capable of auto-guiding on defocused stars, useful, for example, when we observe bright objects.

2.1.5 Discoveries with SPECULOOS

Since TRAPPIST-1, there have been no discoveries of a similar transiting planetary system with SPECULOOS. Sebastian et al. (2021) predict that SPECULOOS will yield 29 ± 4 planets (in 12 planetary systems) in its lifetime, with 8 ± 2 within the habitable zone of their host star. Several clarifications need to be noted alongside this predicted planet yield. This prediction is calculated from Monte-Carlo simulations of planets, assuming that every target in the target list hosts at least one planet, and the numbers of planets per system are drawn from the distributions in Miguel et al. (2019). The planetary parameters are (as in Delrez et al. 2018b) modelled on the TRAPPIST-1 system, and each planet has a 50% chance of having a period of 0.5–23 d or 23–1800 d. A planet is recovered if it is (a) transiting, (b) transits during simulated SPECULOOS observing time and (c) has a signal-to-noise ratio of at least 5, with an assumed noise floor of 500 parts per million. A simulated planet is defined as in the habitable zone if the received incident stellar flux reaching the top of the planet's atmosphere is between

0.2–0.8 times that of the Earth. These limits resemble the conservative habitable zone limits determined by Kopparapu et al. 2013 for an UCD - I refer the reader to their Figure 8 where they present the effective stellar flux against stellar effective temperature. These reduced habitable zone limits (compared to the Solar System) are due to the fact that the stellar spectra of cool stars are shifted towards longer wavelengths. Earth-like planets have lower planetary albedos (and so absorb more starlight) around cool stars than Sun-like stars due to reduced Rayleigh scattering and increased absorption of H_2O and CO_2 at near-infrared wavelengths. The runaway greenhouse effect determines the inner habitable zone (HZ) limit, and the outer is determined by the maximum greenhouse limit (where the atmosphere becomes opaque to outgoing infrared radiation). Planets at the inner limit have H_2O-dominated atmospheres; therefore, their increased absorption in the near-IR means that they can go into a runaway at lower stellar fluxes. At the outer HZ limit, planets around cooler stars absorb more CO_2 and have reduced Rayleigh scattering from CO_2 condensation (Rayleigh scattering increases the planetary albedo for Sun-like stars but is minimal for late-type stars), requiring lower stellar flux to sustain a habitable surface temperature (with the greenhouse effect). Therefore, SPECULOOS's planetary yields must be treated with caution, as there are assumptions built-in, and they include no information about the intrinsic activity of targets or unexpected noise sources. The most accurate predictions about SPECULOOS's ability to recover planets (and underlying UCD planet distributions) will come from the injection and recovery of simulated transits on real SPECULOOS lightcurves.

As of May 2020, we had observed around 10% of SPECULOOS targets for at least 50 hours. Within these predictions, Sebastian et al. (2021) found that we should have discovered 2 ± 2 planets. Therefore, this result is consistent with the current state of non-detections. If, however, in the coming years SPECULOOS does not find any additional planets, then this absence would have far-reaching consequences for planetary occurrence rates around ultra-cool dwarfs.

Since becoming operational SPECULOOS has also had success in a number of its complementary science goals, such as:

- Providing follow-up photometry to confirm planetary candidates (and interesting stellar systems) from NGTS (Günther et al. 2018; Smith et al. 2021a,b), WASP (Barkaoui et al. 2019; Nielsen et al. 2019), K2 (Niraula et al. 2020), TESS (Wells et al. 2021; Scanche et al. accepted; Timmermans et al. in prep; Pozuelos et al. in prep; Barkaoui et al. in prep) and CHEOPS (Leleu et al. 2021),
- Discovering novel ultra-cool systems, such as the first-ever detection of an eclipsing substellar binary in a young triple system (Triaud et al. 2020),
- Helping to shed light on the mysterious magnetic behaviour of ultra-cool dwarfs, by studying their flaring and rotating activity (Murray et al. 2022), examining the complex

modulation of M-dwarfs with multi-wavelength observations (Günther et al. 2020a), and using wavelength-dependant transit depths to constrain stellar spot coverage (Zhang et al. 2018),

- Refining the planetary parameters and transit timing variations of known planets (Ducrot et al. 2018; Bryant et al. 2021).

2.2 The SPECULOOS-South Pipeline

2.2.1 Pipeline Design

Simply put, a data pipeline is a series of data processing tools, where the output of one is fed as an input into the next, to take you from initial input A to final product B. In our case, for transit detection, the first input is a collection of images taken by a telescope of a particular stellar field. These images are subjected to a chain of complex procedures, each of which improves the output or extracts some useful information. In the end, we have a final time-series of the target object's brightness (a lightcurve), on which we can search for transits. This chapter will describe each of these steps in turn that take us from raw image to lightcurve.

As every survey presents unique calibration and photometric challenges, most will develop their own specific pipeline. Accordingly, I created a photometric pipeline for the SPECULOOS-South data, designing it to be fast, automatic, and modular. Depending on the targets and conditions of the night, we accumulate approximately between 250 and 1000 images per telescope per night with typical exposure times of 10–60 s, corresponding to between 4 and 16 GB of data. Having flexibility in the pipeline allows me to continuously perform various quality checks, extract feedback and use these outputs to optimise the performance of the survey. Modularity allows reprocessing certain stages of the pipeline with improved algorithms without requiring a complete rerun.

The structure and data format of the SSO pipeline is based on the architecture of the NGTS pipeline described in Wheatley et al. (2018). Similarly to NGTS, I built the SSO pipeline based on the CASUTOOLS[c] package of processing tools for image analysis, astrometric fitting and photometry (Irwin et al. 2004).

The various modules of the pipeline are illustrated in the flowchart in Figure 2.4. The science images are calibrated through bias and dark subtraction and flat-field division to remove various sources of instrumental error (Section 2.2.3). I then derive the astrometric solutions for each image to precisely know where the telescope is pointing (Section 2.2.4). If this is the first night of observation for a given field, these images are aligned and stacked to create a stacked image. The sources detected on this stacked image form a catalogue of bright objects in this field of

[c] http://casu.ast.cam.ac.uk/surveys-projects/software-release

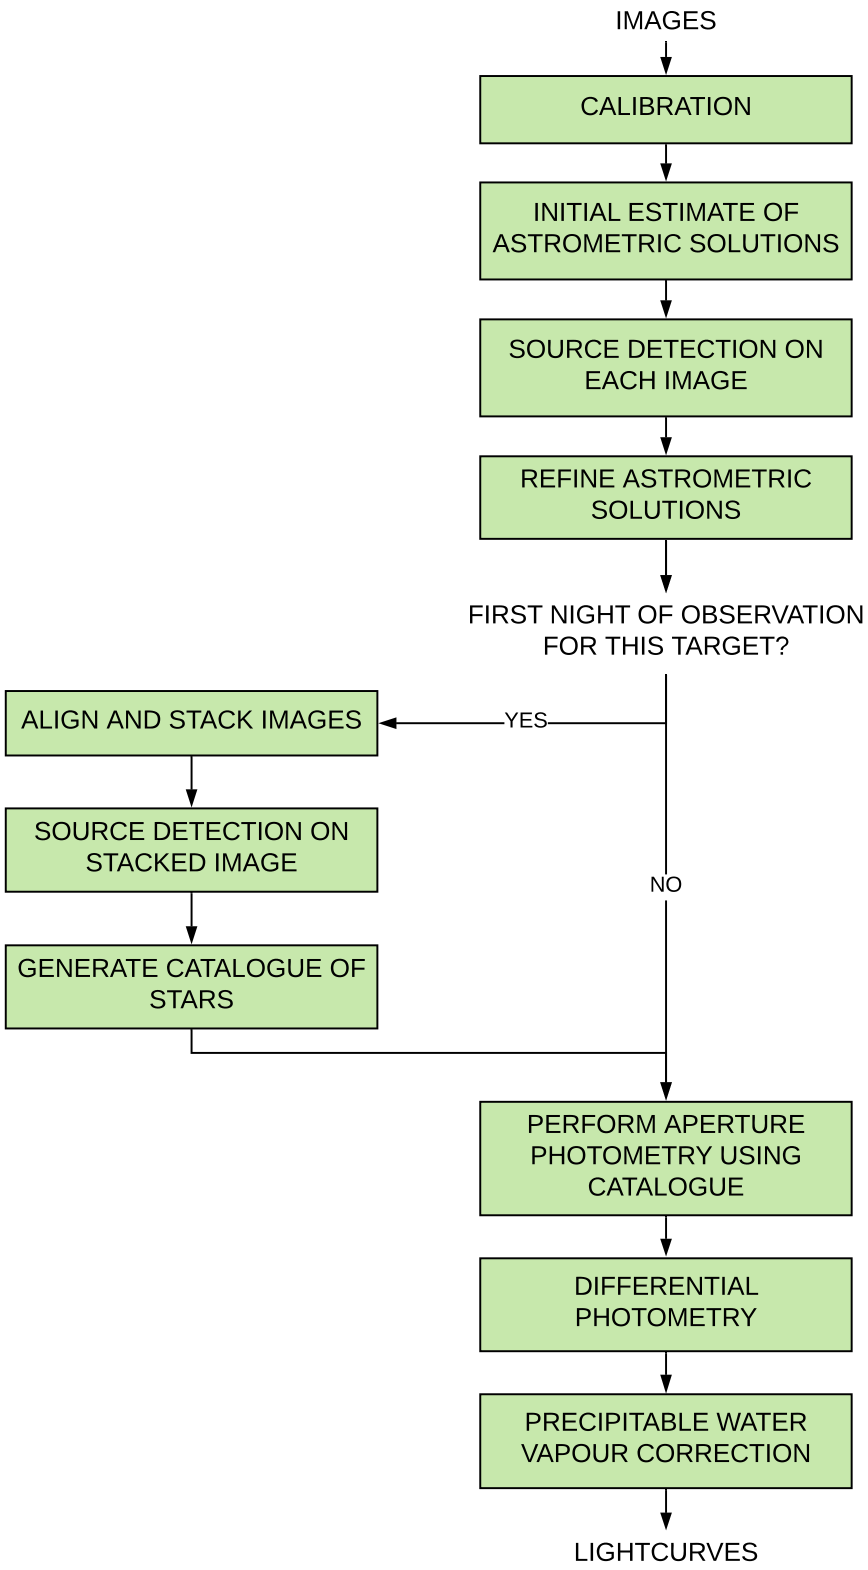

Figure 2.4: Simplified flowchart of the SPECULOOS-South Pipeline.

view, which can be cross-matched with known stellar catalogues (Section 2.2.5). Time-series of precise aperture photometry measurements are extracted for every source in the catalogue (Section 2.2.6). Then I use these measurements to generate differential lightcurves for all field objects, both for a single night and over multiple nights, to assemble a 'global' lightcurve (Section 2.2.7). Global lightcurves can be used to monitor the photometric variability of a target over longer timespans. Finally, I detrend for the systematic effect caused by variable precipitable water vapour (Section 2.4).

2.2.2 The ESO Archive

All raw images recorded by the facility are automatically uploaded at the end of the night to the online ESO archive[d]. These images are then (also automatically) downloaded to a server at the University of Cambridge and analysed by the SSO Pipeline. This data is not currently public. The SPECULOOS-South Consortium will make all SPECULOOS-South Facility reduced data products (images and extracted lightcurves) available to the ESO Science Archive Facility following the completion of the regular Phase 3 process[e].

2.2.3 Reduction

I use the first stage of the pipeline to mitigate instrumental systematics. As described in Section 1.2.4, there are several ways in which our detector can imprint non-astrophysical variations in brightness. These variations can result from pixel-to-pixel variations, readout noise, dark current, dust grains, vignetting, or non-homogeneous field illumination. To mitigate these undesirable effects, we capture calibration images each night in addition to the science images that are used to generate lightcurves. The three types of calibration images are described in detail below and demonstrated in Figure 2.5. I show the reduction of a raw science image in Figure 2.6.

Bias Subtraction

Bias images are images taken with zero exposure. In this sense, they are not "real" images. Instead, they provide a picture of the instantaneous underlying pixel fluctuations when no light reaches the CCD and no thermal electrons are excited. The bias image will show a small inter-pixel structure above some offset voltage (bias level), which I name the residual bias. The bias level is shown as the top and bottom regions of overscan in Figure 2.6a and is approximately 300 counts for the SSO telescopes.

[d] http://archive.eso.org/eso/eso_archive_main.html
[e] http://www.eso.org/sci/observing/phase3.html

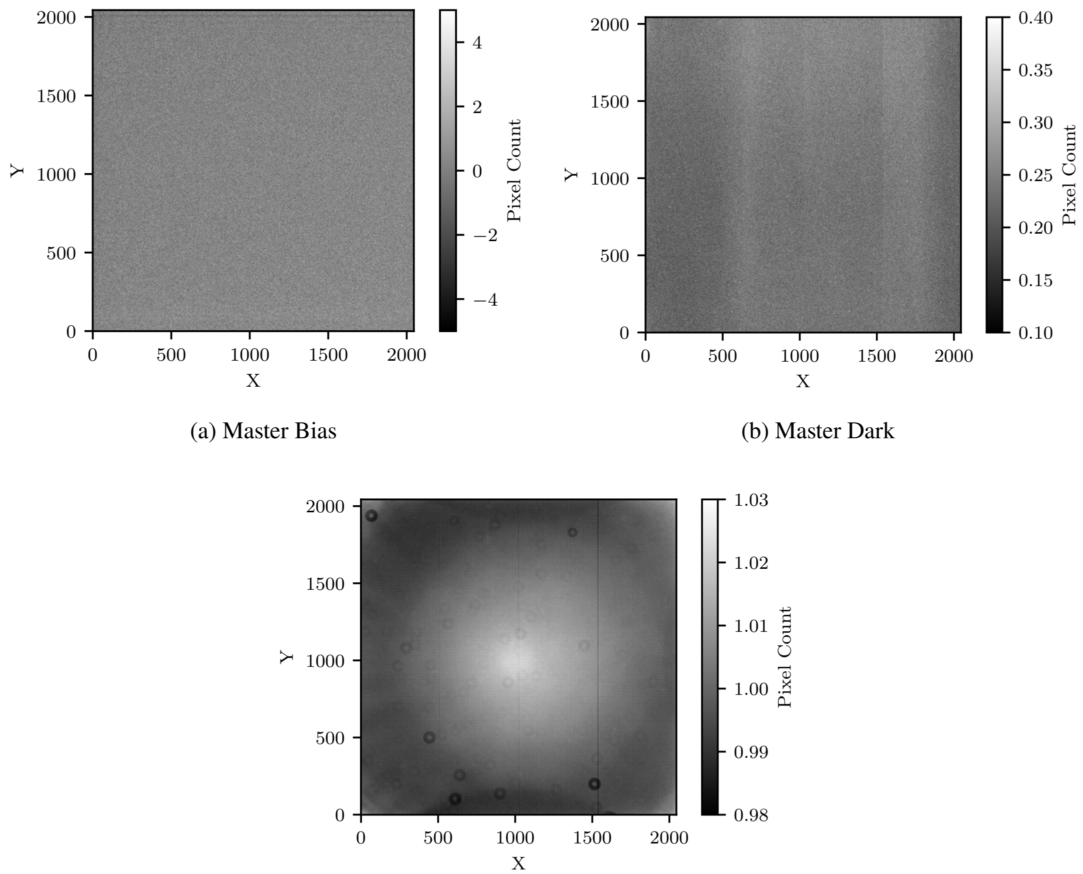

(a) Master Bias

(b) Master Dark

(c) Master Flat

Figure 2.5: Example master calibration images: (a) bias, (b) dark and (c) flat.

Ten bias images are taken each night at dawn, after the closure of the telescope dome. To correct for the pixel inhomogeneity, I combine these bias images into a "master bias". By averaging over multiple bias frames, we produce an image with less random noise than individual images. The master bias is taken to be the sigma-clipped median of all the bias images' residual bias. Sigma-clipping is the process in which we remove data points that are more than some multiple of the standard deviation above or below the median value. The sigma-clipped median is taken to be a good estimate of the average pixel-to-pixel variation. The master bias is shown in Figure 2.5a. The readout noise is taken as $\sigma/\sqrt{2}$, where σ is the standard deviation of the difference between two uncorrected bias images (which isolates only the noise introduced when the image was read from the CCD). On average, the SSO's readout noise is 6.3 e^{-}. Monitoring of the readout noise is demonstrated in Appendix A. Sudden variations or gradual changes in readout noise can provide information about the health of the electrical components.

Dark Subtraction

Like bias images, dark images are also taken at dawn, with the telescope dome and shutter closed. However, the difference is that dark images are taken with exposure times of 15, 30 and 60 s. Dark images can indicate the dark current (or the number of thermally excited electrons) generated during various lengths of time, typical for our observation exposures. By capturing dark images with exposure times different to science images, they will need to be scaled. In performing this scaling, I assume that the dark current increases approximately linearly with time. Dark images can also be used to identify spurious pixels, such as “hot pixels” that build up a higher dark current due to a faulty temperature response.

Similarly to the master bias, a master dark is built from the several dark frames captured each night. The dark images are first corrected for the bias level, then de-biased (by subtracting the master bias to remove time-independent pixel structure), and divided by the exposure time to give the rate of dark current build-up (in electrons per second). The master dark is then generated from the sigma-clipped median of these processed images. Taking the median, in this case, helps reduce the impact of cosmic rays. The master dark is shown in Figure 2.5b. The dark current is simply taken as the average value of the master dark (per pixel). Monitoring the dark current over time can provide feedback on the condition of the CCD cooling system. We cryogenically cool SPECULOOS’s CCDs to between -60°C and -70°C, which reduces the dark current to almost negligible. On average the SSO’s dark current is 0.3 e^{-} pixel^{-1} s^{-1} at -60°C. Monitoring of the dark current is demonstrated in Appendix A.

Flat-field Correction

By spatially mapping how the CCD responds to a uniform light source, we can correct the various effects that prevent us from achieving a perfect flat-field. The two methods of obtaining a uniform light source are by using a diffuse light source ("dome flats") or, as with SPECULOOS, using the twilight sky at dusk or dawn or both ("sky flats"). Using the twilight sky, we take flat-field images each night, which primarily capture two major sources of non-homogeneous light response: dust-grains and vignetting. Over time dust builds up and moves across the surface of the CCD window; therefore, it is important to monitor these dust particles by way of maintaining an up-to-date flat-field. Each filter used on a specific night must have corresponding flat images, as the filters themselves may have dust or heterogeneity.

I generate a master flat similarly to the master bias and dark images. Once again, I remove the bias level and de-bias each flat image, but here I also remove the dark current by subtracting

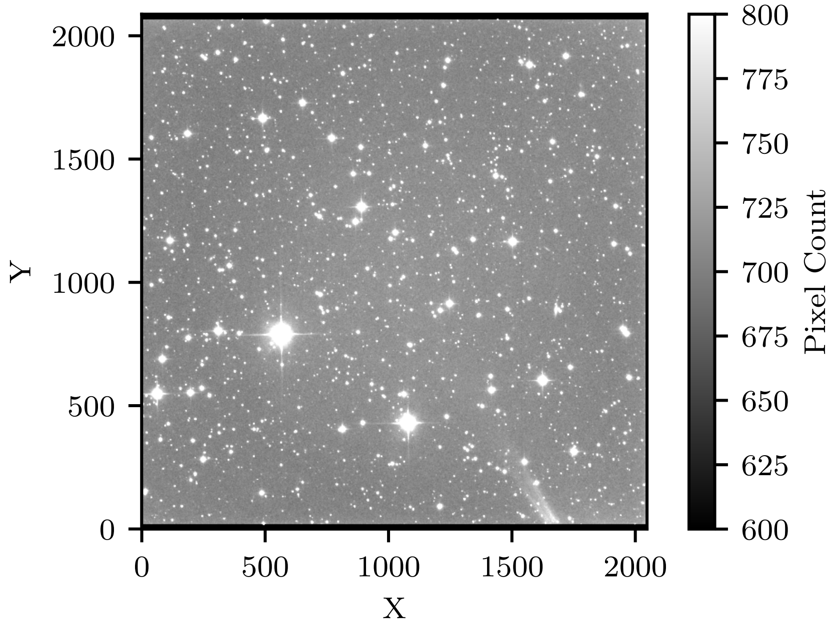

(a) Raw Image

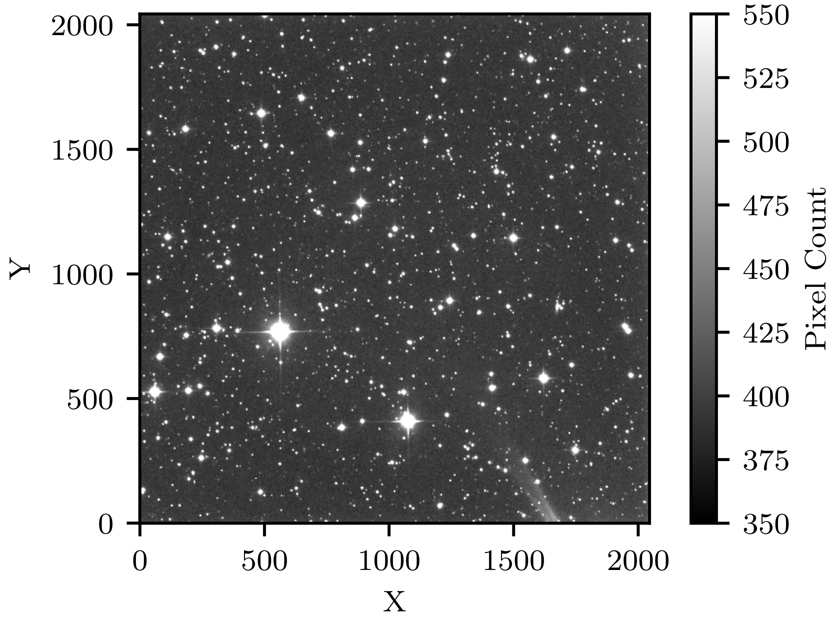

(b) Calibrated Image

Figure 2.6: (a) The raw image from the telescope. (b) The same image as (a) but with the overscan removed and calibrated using the master bias, dark and flat.

the master dark multiplied by the exposure time of the flat images. The master flat is then the sigma-clipped median of these flat images, shown in Figure 2.5c. Currently, only the master flats from each night are used for calibration. Since the flat images must be taken during the twilight sky, there is only time for, on average, 10 flat images to be taken before the sky becomes too light or stars begin to appear. The master flat images are monitored over time to assess their quality and flag significant variations (e.g., moving dust).

Mine is a rather basic flat-field correction, and several improvements could be made to this stage in the future. NGTS combine an entire season's worth of bias and flat-field master images to generate the best overall calibration master frames. These nightly flat-fields are monitored over time to detect any significant changes (Wheatley et al. 2018). The SuperWASP pipeline combines flat-fields from different nights using an algorithm with weights that decay on a timescale of 14 days (Collier Cameron et al. 2006). Including more flat-field images, such as from multiple nights, in creating the master flat is advantageous as it reduces the random noise. Additional elements that could improve the flat-field include a shutter map (to remove distortion effects from the camera shutter, which slows over time), a fringe map (to remove the stable spatial interference pattern seen strongly in the infrared) and a confidence map (to mask pixels with values outside of a specific tolerance range). Additionally, in theory, a method for automatically detecting moving features (such as dust grains), mapping their positions, and even correcting for the flux variations they cause would be very useful. However, this task proves extremely complex in practice as dust grains do not move uniformly or predictably, and since flat images are only taken at dusk and dawn, there is no way to know the various changes over the course of observations from the flats alone.

2.2.4 Astrometry

Despite the good performance of the telescope guiding with DONUTS, there remain very small drifts in object positions during the night, of the order of ~0.1 arcsec (0.3 pixels). Precise astrometric calibrations are needed for each image to place apertures for photometric measurements accurately. I use a local version of ASTROMETRY.NET code (Lang et al. 2010) to cross-match each science image with reference catalogues built from the 2MASS catalogue (Skrutskie et al. 2006) to find an initial approximate World Coordinate System (WCS) solution. This solution is then refined by using first IMCORE, to detect sources on the image, and then WCSFIT, to produce the final WCS solution, from the CASUTOOLS package.

IMCORE performs source detection by first computing a low-resolution background image. This process involves estimating background values for 64×64 pixel^2 sections by using an algorithm based on a robust (MAD) iterative k-sigma clipped median. These background values are then filtered to produce a low-resolution background image. Using bilinear interpolation, the local sky background of every pixel in the original image can then be derived. To identify a source, the algorithm searches for a connected series of 6 pixels with values higher than a user-specified threshold above the background. For astrometry, I want to use as many stars as possible; therefore, I use a low limit of 2-sigma above the background sky level to detect sources.

WCSFIT uses the initial WCS solutions to further correct each image's WCS solutions for translations, skews, scales and rotations by cross-matching the sources from IMCORE with the Gaia Data Release 1 catalogue (Gaia Collaboration et al. 2016a,b). At the time the SSO Pipeline was created, there was no facility to cross-match with the Gaia Data Release 2 catalogue (Gaia Collaboration et al. 2018); however, this will likely be a feature introduced into CASUTOOLS in the future.

2.2.5 Generating a Stellar Reference Catalogue

For each field of view that is observed (i.e. each target), the pipeline requires an input catalogue with the RA and DEC of the stars on which to extract aperture photometry data for each image. This catalogue is generated from a stacked image produced from 50 images in the middle of the night (to reduce the airmass and sky background), taken on a target's first night of observation. I then have a unique catalogue for each field of view which is then referenced across all the subsequent nights that target is observed, to track these stars over long periods of time.

The IMSTACK and IMCORE programs from the CASUTOOLS package (Irwin et al. 2004) are used in generating this catalogue. For each of the 50 science images IMSTACK aligns (using the WCS solutions from WCSFIT) and stacks these images to produce the final stacked image.

IMSTACK defines a WCS reference grid using the first image and subsequent images are then aligned and resampled onto this grid. The sigma-clipped mean of the pixel values from all images, scaled by their exposure times, is computed and recorded in the output stacked image. Outliers (defined by threshold values of 5σ) are removed from the averaging. IMSTACK uses a bi-linear interpolation approach where an input pixel is divided into the four pixels on the output grid that surround the input equatorial position, as this can reduce systematic errors (Mighell 1999). The fraction in each output pixel corresponds to the amount of overlap of the input pixel. The final stacked images are crucial in the creation of the catalogues that define each field of view. Therefore, quality checks implemented by the automatic pipeline help to ensure the stacked image is created on a night with good seeing and atmospheric conditions, and ideally no defocusing, to increase the accuracy of the source positions on the field.

IMCORE then performs source detection on the stacked image to create a catalogue of the stars in the field of view. This time, however, IMCORE searches for sources with more than 6 contiguous pixels containing counts 8-sigma above the background sky level. This higher threshold limits the detected objects to $I + z'$-magnitudes brighter than ~21. The background sky level present in the stacked image will vary depending on the angular proximity and phase of the moon. However, I did not see any noticeable variation in the number of stars in the catalogue corresponding to the moon cycle, potentially due to the small pixel size of our CCDs.

This catalogue is cross-matched with Gaia Data Release 2 (Gaia Collaboration et al. 2016a, 2018) to apply proper motion corrections on a night-by-night basis. To cross-match, I perform a "cone search" (returns all objects in a catalogue encompassed within a circular area of a given radius centred on a given RA and Dec) with a radius of 30 arcseconds around every detected source. I then correct each of the matched objects in the Gaia DR2 catalogue for its proper motion to find its position at the time of observation. Finally, I take the object predicted to be closest to the catalogued source, as long as it has an on-sky separation smaller than 0.00001 radians. If the separation is larger, I discard the match. It is possible to have sources for which there is no match in the Gaia DR2 catalogue, or have multiple sources match with the same Gaia object. This can happen if a star is extremely faint (and, therefore, is not catalogued in Gaia DR2), or if it has a very high proper motion so that its Gaia match is not found in the original 30 arcsecond cone search. The latter issue will become increasingly common as more time passes from the Gaia DR2 release. In the future there is the facility to cross-match with additional catalogues, such as 2MASS, and to update the cross-match to use positions from the newest Gaia Data Release 3 (Gaia Collaboration et al. 2021).

2.2.6 Aperture Photometry

Aperture photometry involves centring apertures on each star (equivalent to blocking all light except that which passes through a small circular hole) and summing the pixel counts above the average background sky level. IMCORELIST, a fourth CASUTOOLS program, is used to perform aperture photometry on each science image. It carries out essentially the same process as IMCORE but requires an input list of equatorial positions, provided by the catalogue, to define the positions of the apertures. IMCORELIST takes photometric measurements of each source on every image for 13 "soft-edged" apertures sizes which are multitudes[f] of the user-defined radius r_{core} (I define the radius as the average derived seeing over a year of nightly observations with the SSO, which is 4 pixels or 1.4 arcsec). Theoretically, the optimum aperture radius to use is ≈ seeing, as this provides a good balance of systematic centering errors (larger for smaller apertures) with the fact that signal-to-noise ratios decrease for larger apertures because they contain more background flux (Mighell 1999). However, the pipeline chooses the optimal aperture for a given night by binning every 5 minutes and selecting the aperture which minimises

[f]The 13 apertures used are multiples (1/2, $1/\sqrt{2}$, 1, $\sqrt{2}$, 2, $2\sqrt{2}$, 4, 5, 6, 7, 8, 10 and 12) of r_{core}.

the average RMS of data points within each bin (to represent the photometric scatter of the target's differential lightcurve) multiplied by the overall RMS of the binned lightcurve (which represents the correlated noise in the target's differential lightcurve). This method was chosen to avoid both minimising genuine photometric structure (e.g., stellar variability), and adding correlated noise into the lightcurve, for example from the changing seeing and airmass during the night if an aperture is chosen that is too small.

Measuring the Seeing

Seeing is an extremely useful indication of how well we can resolve images, based on the temporally-changing turbulence of the Earth's atmosphere. It is possible to derive the seeing directly from the science images, by using the average stellar profile. King (1971) demonstrated that the central core of ground-based stellar profiles are approximately Gaussian; therefore, to a good approximation a star can be fit with a standard 2D elliptical Gaussian function. To measure the seeing I average multiple stars in the field to get the "shape" of a typical star (or point spread function, PSF). I then fit this shape with an elliptical Gaussian to extract the ellipticity and the full width of the Gaussian at half maximum (FWHM), which becomes the seeing.

To assess the PSF (and measure the seeing) for a particular frame, each science image is split into a grid of 3x3 squares. For each grid section, I isolated and stacked the stars from the Stack Catalogue contained within that area. The FWHM values are then calculated following the method described below. The image is split into 9 squares because it shows how the shape of the PSF varies across the field of view in case any spatially-dependant instrumental systematics are present. The shape of the PSF can also provide the limitations of the model to describe the data with increasing defocus of the image. As mentioned earlier defocusing our images can reduce saturation and flat-fielding errors, however a 2D Gaussian is unlikely to fit well for very defocused images (where point sources appear donut-shaped). Monitoring the average ellipticity of the stellar profiles also provides a useful quality check.

The equation of a 2D elliptical Gaussian is as follows:

$$f(x, y) = A \exp\left(-(a(x - x_0)^2 + 2b(x - x_0)(y - y_0) + c(y - y_0)^2)\right) \tag{2.1}$$

with coefficients:

$$\begin{aligned} a &= 2\sigma_{\mathrm{x}}^2 cos^2\theta + 2\sigma_{\mathrm{y}}^2 sin^2\theta \\ b &= -4\sigma_{\mathrm{x}}^2 sin^2\theta + 4\sigma_{\mathrm{y}}^2 sin^2\theta \\ c &= 2\sigma_{\mathrm{x}}^2 sin^2\theta + 2\sigma_{\mathrm{y}}^2 cos^2\theta \end{aligned} \tag{2.2}$$

In these equations A is the Gaussian's amplitude, (x,y) is the coordinate system in the plane of the image, (x_0,y_0) are the coordinates of the centre of the Gaussian, σ_{x} and σ_{y} are the corresponding spreads of the Gaussian in the x and y directions, and θ is the clockwise angle the Gaussian is rotated.

Therefore, there is the following relation between the FWHM, σ_{x} and σ_{y}:

$$\begin{aligned} FWHM_{\mathrm{a}} &= 2\sigma_{\mathrm{y}}\sqrt{2ln2} \\ FWHM_{\mathrm{b}} &= 2\sigma_{\mathrm{x}}\sqrt{2ln2} \end{aligned} \tag{2.3}$$

where b is along the short axis and a is along the long axis of our ellipse. I used a least squares optimisation to fit the average stellar profile to this function and by optimising for the parameters (A, θ, σ_x, σ_y, x_0, and y_0) I calculated the FWHM.

2.2.7 Differential Photometry

Differential photometry is a technique based on the assumption that stars of similar brightness, on-sky location and colour in a field of view will experience a common photometric pattern, due to shared atmospheric and instrumental effects.

It is worth noting here the difference between *absolute* and *differential* lightcurve. A star's absolute lightcurve is defined as a time-series of the raw flux values extracted from aperture photometry. In other words, it is the sum of all electron counts within an aperture placed on that star for every frame. Even if the absolute lightcurve is normalised, it still suffers from the systematics attributed to atmospheric and instrumental brightness fluctuations. In this context 'normalised' means the lightcurve is divided by its median value. Alternatively, that star's differential lightcurve is the normalised absolute lightcurve divided by a reference lightcurve in an attempt to remove these systematics.

For the SSO, I developed an algorithm to automatically choose and combine multiple comparison stars. This ensures that the final differential lightcurves are reproducible and avoids the time-intensive, manual selection of stars and potential observer bias. Statistically, it is optimal to average over as many comparison stars as possible to reduce the photometric scatter in the final differential lightcurves. However, when using large numbers of comparison stars you have to contest with the fact that the lightcurves from faint stars have lower signal-to-noise ratios and including them can reduce the quality of your differential lightcurves. Using a weighted average of multiple stars is a way to mitigate this effect if weightings are carefully assigned (see following section).

The algorithm I implemented in the SSO pipeline is based on a concept described in Broeg et al. (2005). This iterative algorithm automatically calculates an 'artificial' comparison lightcurve (ALC) by weighting all the comparison stars based on their variability and removing those which are highly variable. To optimise the pipeline for SSO data, I implemented several major changes from the algorithm developed by Broeg et al. (2005). The basic algorithm is

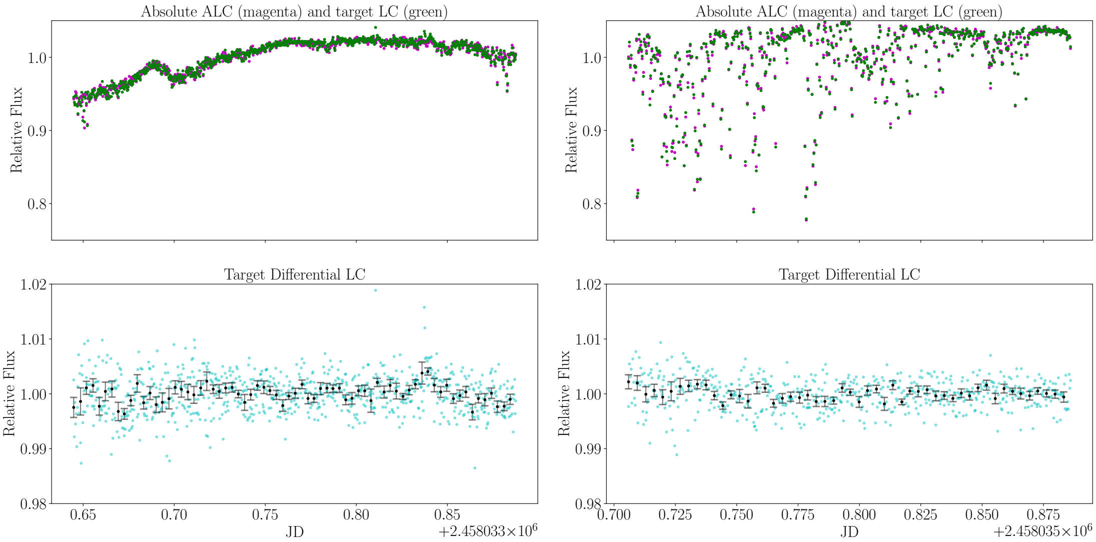

(a) (b)

Figure 2.7: Demonstration of the differential photometry algorithm on a bright M8V (J = 10.4 mag) target star, observed by Europa during its commissioning phase, comparing the results on a relatively clear night (a), 6th October 2017, and a cloudy night (b), 8th October 2017. The top plots show the artificial lightcurve (magenta) compared to the target's absolute lightcurve (green), for both nights the optimal aperture is 11.3 pixels. The bottom plots show the target's final differential lightcurve (unbinned points in cyan and 5-min binned points in black), produced by dividing the target's absolute lightcurve by the artificial lightcurve. The differential lightcurve for (a) shows a small flare-like structure (JD 2458033.84), which would be difficult to extract from the absolute lightcurve.

described, along with my changes, in the following subsections. I demonstrate the need for differential photometry and its implementation on observation nights of different photometric quality in Figure 2.7.

The Creation of an 'Artificial' Comparison Star

The following method is similar to that described in Broeg et al. (2005) where each object (excluding the target and any saturated stars), i, is assigned a weight, $W_{\mathrm{var},i}$, determined by its variability.

1) The initial weights are defined as:

$$W_{\mathrm{var},i} = 1/\sigma^2_{\mathrm{photon},i}, \tag{2.4}$$

where $\sigma_{\mathrm{photon},i}$ is the photon noise of star i; therefore, in this step $W_{\mathrm{var},i}$ is set to be equal to the average flux for each object. These weights are then divided by their sum such that they sum to 1. At this initial stage, therefore, the stars are weighted simply on their brightness.

2) An artificial lightcurve (ALC) is constructed from the weighted mean of the normalised absolute flux (F) of each of the n objects in the field, at each frame j:

$$ALC_j = \frac{\sum_{i=1}^{n} W_{\mathrm{var},i} F_{ij}}{\sum_{i=1}^{n} W_{\mathrm{var},i}}. \tag{2.5}$$

3) Every star's normalised absolute flux, F, is divided by this ALC to produce its differential lightcurve.

4) The weight for star i is replaced by:

$$W_{\mathrm{var},i} = 1/\sigma_i^2 \tag{2.6}$$

where σ_i is the standard deviation of the entire night's differential flux (equivalent to the differential lightcurve) for star i. This step acts to weight down intrinsic stellar variability.

Stages (2), (3) and (4) are repeated with these new weights until the weights are constant to within a threshold of 10^{-5}.

Initial variability cut

From testing, it became clear that if there was variability in the brightest stars, which are highly weighted during stage (1) of this algorithm, then the initial ALC estimate would be significantly affected. If these objects are not removed, in the next iteration, more stable stars would be down-weighted and any with similar time variability structure weighted up. This results in a runaway effect, down-weighting the less variable comparison stars. Therefore, I included a basic variability check prior to generation of the initial ALC by sigma-clipping across all stars' normalised lightcurves for each frame. For every point in time, this process masks any star's flux that differs significantly from the average of all other stars in the field. If any object has >20% of its values masked it is determined that this object is variable, and it is removed.

Distance

Due to spatially varying atmospheric and optical effects, I added an additional weight based on projected distance from the target star, using the formula:

$$W_{\mathrm{dist},i} = \frac{1}{1 + \left(\frac{a s_i}{s_{\max}}\right)^2} \tag{2.7}$$

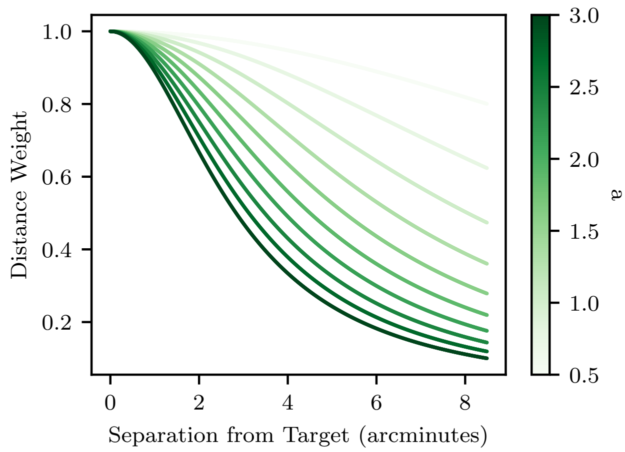

Figure 2.8: The distance weighting, Equation 2.7, for a star against its on-sky separation from the target, for different values of a.

where $W_{\mathrm{dist},i}$ is the distance weight of star i, s_i is its on-sky separation from the target star, $s_{\max}$ is the maximum distance of any star from the target and a is a parameter optimised for each night. I chose this form to be finite and relatively flat near the target object and decay slowly as the distance on sky increases. This functional form is shown in Figure 2.8 for a range of values of a from 0.5–3. The value of a is chosen to minimize the 'average spread', or photometric scatter, of the target's differential lightcurve (as defined in Section 2.2.6). a does not change significantly between nights, SPECULOOS's small field-of-view likely minimises spatial variations, typically favouring lower values around 0.5. We normalise these weights to sum to 1, and combine the distance weights and the variability weights, $W_{\mathrm{var,\,i}}$, to produce the final weights used in the ALC:

$$W_i = W_{\mathrm{var},i} W_{\mathrm{dist},i} \tag{2.8}$$

Once again, these weights are normalised and replace the weights in step (4) of the iteration process.

Removal of the faintest stars

Ideally, I would use as many comparison stars as possible in the creation of the ALC (weighted appropriately). However, including a large number of faint comparison stars increases the red noise in the ALC. This effect is particularly clear on nights where the atmospheric transmission varies by more than 30%, suggesting passing clouds or poor weather conditions which limit our ability to conduct precise photometric measurements. Red noise is likely imprinted into lightcurves when faint stars are included in the ALC when there are significant brightness variations that produce the same decrease in absolute flux for a large number of stars in the field. Possible sources of these 'absolute brightness variations' are errors in our background subtraction algorithm, or some instrumental error that has not been accounted for, that would cause variations in flux across the course of the night affecting all stars similarly (apparently irrespective of brightness, colour or spatial position). These variations then scale relatively when each star is normalised. This would not be an issue if only a small number of the stars of similar brightness (to each other and to the target) were used. I, therefore, also trialled a threshold that could be adjusted each night, to remove a certain number of faint stars. This threshold value has a minimum value of $m_{\text{target}} + 3$ (removing all stars that are 3 magnitudes fainter than the target, m_{target}) and a maximum value that leaves the 20 brightest (non-saturated) comparison stars remaining. The threshold is chosen automatically from this range so as to minimize the 'average spread' of the target's final differential lightcurve (as defined in Section 2.2.6).

However, even when the very faintest objects had been removed, the ALC still consisted of a large variety of stars with different magnitudes. The observing strategy is also optimised for the target. The exposure time is chosen to maximise the number of photons we collect from the target, without it saturating. Therefore, the "average" star in the field is likely to be fainter than the target. This difference in magnitude means that any absolute brightness variations would be relatively larger for this "average" star (which is normalised by a smaller median flux) than for the target. When the target's absolute lightcurve is then divided by the ALC, this imprints a small amplitude inverse structure into the target's differential lightcurve. Every star in the field experiences the same effect, and so will display the same structure in its lightcurve, with an amplitude dependant on its magnitude. There are two potential, simple solutions to this issue: I could choose not to normalise the lightcurves to avoid scaling absolute features, or I could only use stars with very similar magnitudes to the target. However, in both cases there are drawbacks. For the former, combining multiple stars without normalising is equivalent to adding an additional weight dependent solely on each star's brightness. While this would reduce the red noise, it introduces the risk of bright, variable stars dominating the ALC. I plan to test the latter solution in the future, by implementing more restrictive bounds on the flux threshold described in the previous paragraph. However, due to SPECULOOS's relatively small field of view, often we observe fields where the SPECULOOS target is the brightest star in the field and there are only a small number of usable comparison stars.

Instead, I explored adding an extra step after the differential photometry: a "magnitude-dependent correction", which I demonstrate in Figure 2.9. If there are significant brightness variations throughout the night then there will be a "common field structure" in each of the lightcurves that scales with the magnitude. To find the common field structure I took the

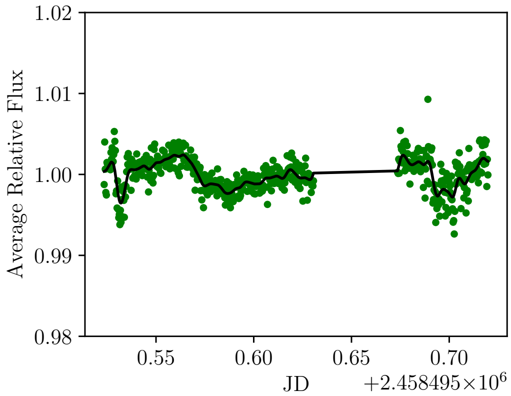

(a) Average night residual structure

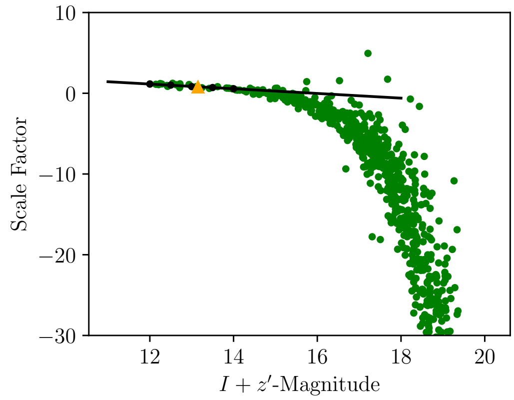

(b) $I + z'$-Magnitude against scale factor

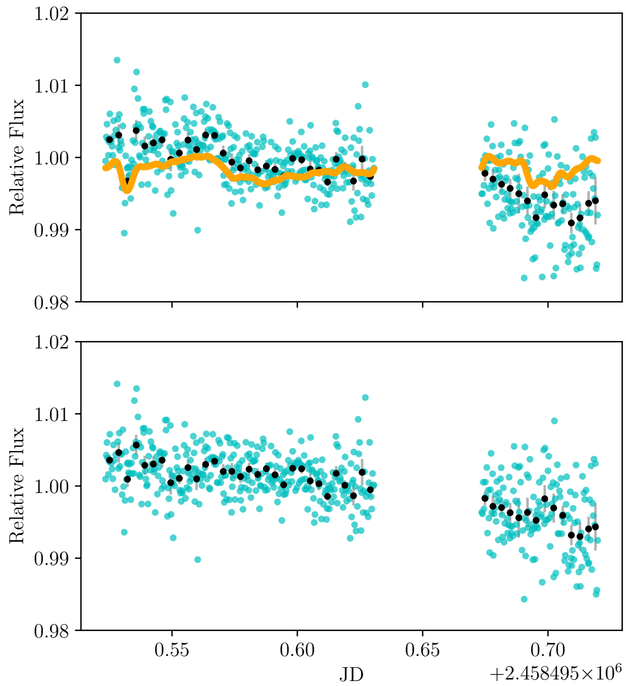

(c) Before and after magnitude-dependent correction

Figure 2.9: (a) The average residual structure in a nightly lightcurve from 11th January 2019, calculated from the weighted sum of the five highest-weighted comparison stars, Gaussian-smoothed in black. (b) The $I + z'$-magnitude of every star in the field against its scale factor, calculated from the least-squares fit to the residual structure shown in (a). The target is marked with an orange triangle (using its fitted scale factor). I bin the magnitudes of comparison stars with magnitudes similar to the target every 0.5 mag (black points). A straight line relation (black) is then fit to the binned magnitudes to calculate the target's predicted scale factor. As $|a| \propto 1/F_a$ (where F_a is absolute flux) and magnitude is $\propto -\log F_a$, $|a| \propto 10^{I+z'\mathrm{mag}}$. The residual structure is inverted for stars fainter than the "average" magnitude of the comparison stars, a is zero at this magnitude and becomes increasingly negative for fainter stars. This is reflected in (b), where a is approximately constant for the brightest objects ($I + z'$-mag < 14), and decreases rapidly for the faintest objects ($I + z'$-mag > 16). (c) Example of the target's correction on a single night. The top plot shows the uncorrected target lightcurve, with the scaled, smoothed residual structure in orange. The bottom plot shows the corrected lightcurve once the residual structure has been divided out.

weighted average of the differential lightcurves of the highest weighted (non-saturated) 5 stars in the field for every frame (see Figure 2.9a). The highest weighted stars are usually the brightest, and, therefore, have the lowest photometric scatter and will exhibit similar magnitude-dependent residuals. I smoothed this structure with a Gaussian filter and found the best fit (using a least squares optimisation) to every differential lightcurve in the field that had a non-zero weighting. This resulted in a series of scale factors for all the stars in the field. By comparing each star's scale factor to its magnitude there is clearly a magnitude dependence that 'saturates' for the brightest stars and almost exponentially increases for the faintest stars, as shown in Figure 2.9b. To ensure this correction was reliable I did not use the optimised scale factor, as the best fit might remove genuine stellar structure (such as rotation patterns and transits). Instead, I fit the relationship between the scale factor and magnitude, so that for a given magnitude we would know the scale of the common field structure. This curve is not straight-forward to fit, however, binning in 0.5 magnitudes and applying a linear fit for a small magnitude range around the target ($m_{\mathrm{target}} + 1 : m_{\mathrm{target}} - 2$, where m_{target} is the magnitude of the target) successfully reproduced the local shape. The predicted scale factor, rather than the fitted scale factor, for the target is then used to divide out the common field structure, as in Figure 2.9c.

Colour

By design, the SSO's targets are usually among the reddest stars in the field of view, and so there is always a colour mismatch between the target star and the comparison stars (see Figure 2.10), resulting in second-order differential extinction effects. The redder comparison stars in the field are often significantly dimmer than the target. It was, therefore, necessary to resist the temptation to implement a strict cut of the bluest (and brightest) stars, which would significantly increase the photometric scatter in the ALC, and subsequently the target's differential lightcurve. Instead I decided to correct the differential extinction from first principles, in a way that avoids removing our best comparison stars, in a stage of the pipeline post-differential photometry (in Section 2.4).

2.3 Global vs. Nightly Lightcurves

Rather than treating every night of data independently, it is possible to perform the differential photometry process in place in the pipeline (see Section 2.2.7) on longer duration photometric time-series. This is particularly useful for studying the photometric variability and rotation of targets over time periods longer than a night.

To create the global lightcurves, I apply my differential photometry algorithm to the entire time series at once, which can span several nights, weeks or months. To ensure any observed

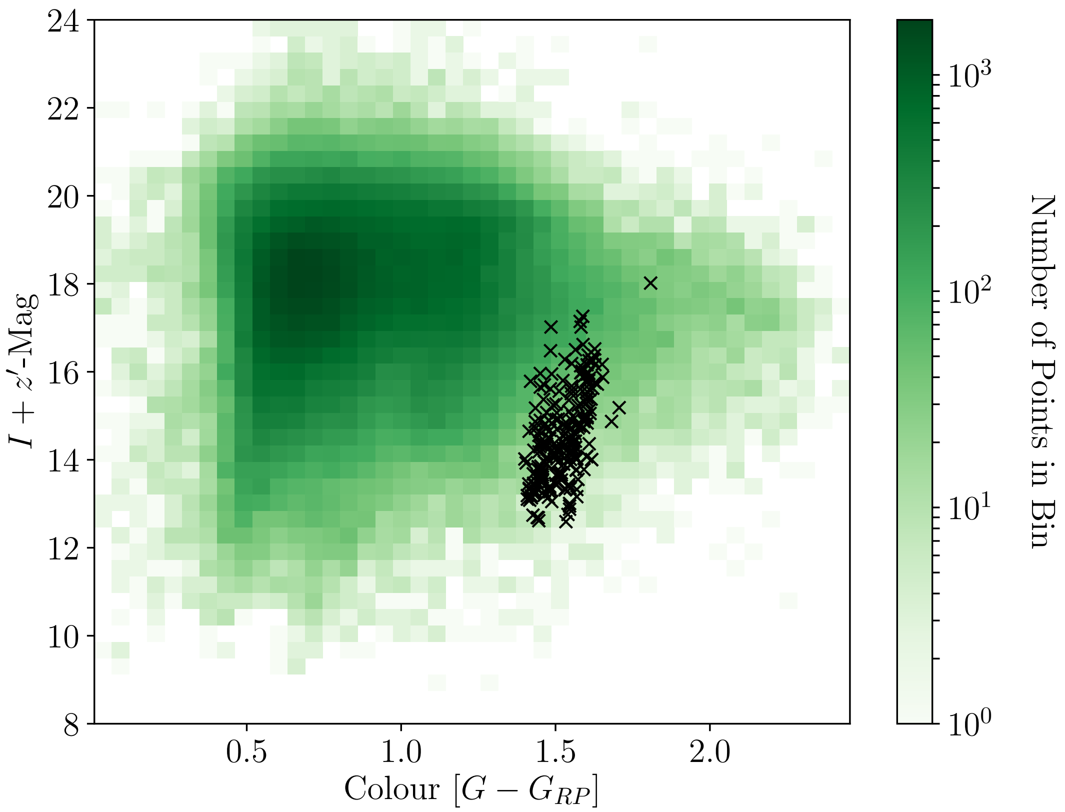

Figure 2.10: $I + z'$-Magnitude against Gaia colour, $G - G_{\mathrm{RP}}$, for all catalogued stars in every observed field of view (on all SSO telescopes) from April 2017 – September 2019. The SSO targets are marked by black crosses

changes in flux between nights are caused by real astrophysical variability (and not as a consequence of the differential photometry process) I use the same comparison stars, weightings and aperture across all nights. This decision, however, reduces the ability to optimise per night, which may result in residuals in the target's final differential lightcurve, which are particularly obvious on nights with sub-optimal observing conditions.

Choosing the optimal aperture for the global lightcurves is a non-trivial process. The optimal aperture changes from night to night, mostly due to seeing variations affecting the size of sources on the field of view. In practice, the optimal aperture of the series has to be large enough to avoid losing stellar flux on the nights with larger seeing. Increasing the aperture size, however, increases the background noise within that aperture, which disproportionately affects the faintest stars. This effect is partially mitigated by the cut and correction I implemented for the faintest stars (see Section 2.2.7). A larger aperture size also increases the risk of "blending" the target's flux with that from any nearby field stars. Therefore, for global lightcurves, I choose the aperture manually, usually by selecting the aperture which has the lowest correlation with

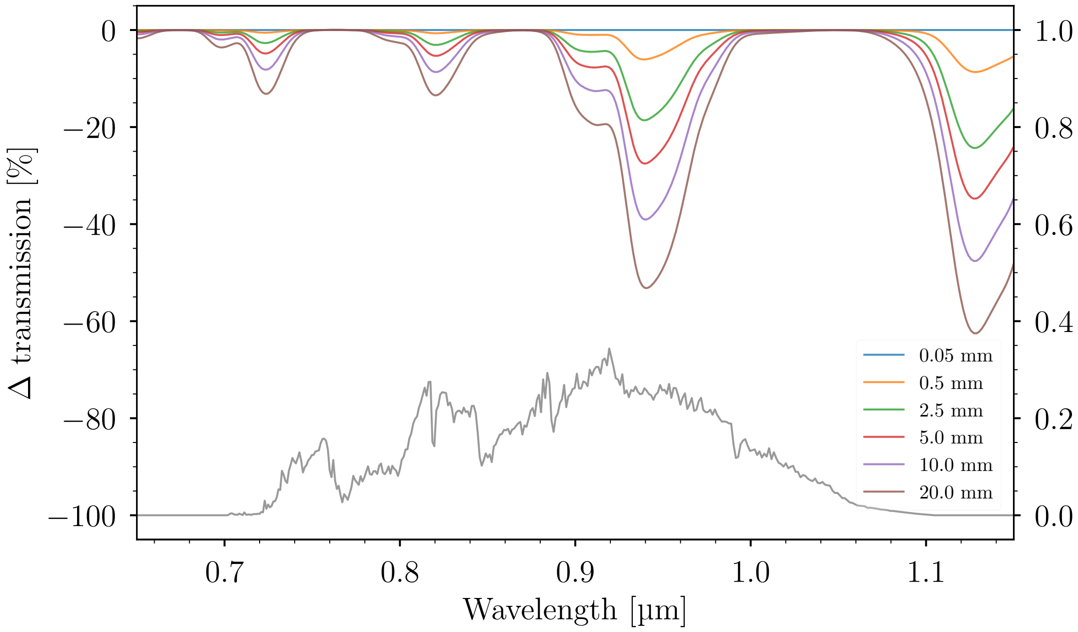

Figure 2.11: Percentage change in atmospheric transmission (left-hand axis) with different increments of PWV (from 0.05 mm) and, in grey, the spectrum of TRAPPIST-1 (from PHOENIX, Husser et al. 2013) as observed through the SSO's $I + z'$ filter, taking into account the overall system efficiency (right-hand axis), as described in Delrez et al. (2018b) and presented in Figure 2.3. This plot was generated with assistance from P. P. Pedersen.

the FWHM.

2.4 The Precipitable Water Vapour Effect

SPECULOOS faces additional photometric challenges to most other ground-based transit surveys, as we are observing very red objects in the near-IR. For the vast majority of our observations, we use the $I + z'$ photometric filter. This wavelength range is strongly affected by atmospheric water absorption lines, and to a much lesser extent by OH radical absorption and emission (airglow) lines. The atmospheric transmission varies strongly with the amount of precipitable water vapour in the Earth's atmosphere (see Figure 2.11), which can be measured from the ground. Despite the fact that Paranal is an exceptionally dry site (situated in the Chilean Atacama Desert), with a nightly median PWV of ~2.4 mm and 45 nights a year less than 1 mm of PWV (Kerber et al. 2014), it can experience large variations in PWV. This includes pronounced seasonal variations (Kerber et al. 2010), variations of up to 20 mm over long time-scales, and even as much as 13 mm during a single night of observation (see Figure 2.12).

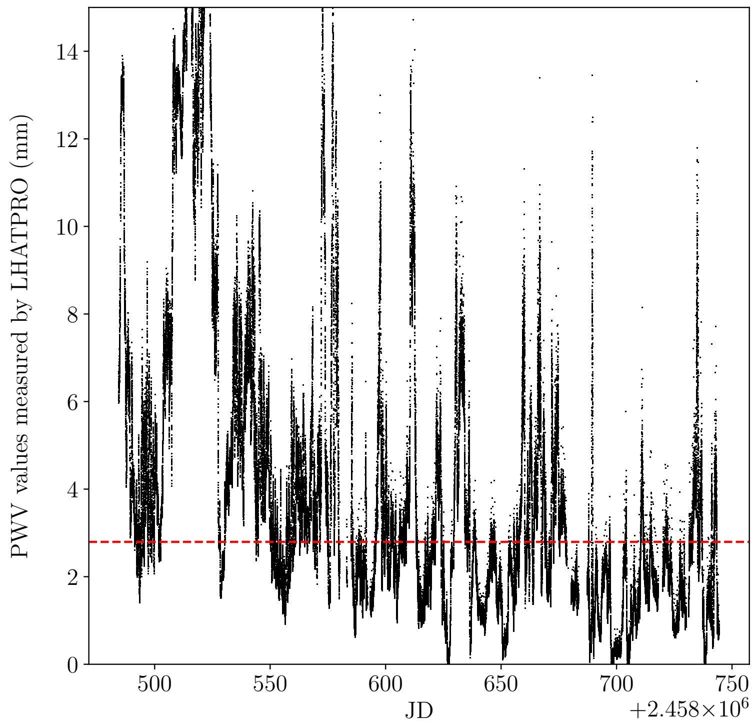

Figure 2.12: PWV variations in Paranal, measured by LHATPRO from 2019 January 1 to 2019 September 18. The median value over this timespan of 2.795 mm is shown by the dashed red line.

By construction of the SPECULOOS's UCD survey, there is always some mismatch in spectral type (and thus colour) between the target and comparison stars used to perform differential photometry. Since redder wavelengths are more readily absorbed by water than bluer wavelengths, when the amount of PWV in the atmosphere changes then objects of different spectral types (whose spectral energy distributions peak at different wavelengths) will experience differing amounts of atmospheric absorption (see Figure 2.13). Temporal variations in PWV can, therefore, imprint second-order extinction residuals on the target differential lightcurves during differential photometry of order $\sim$1% (Baker et al. 2017) or more, when the change in PWV is substantial. These residuals can be a serious limitation for sub-millimag precision surveys, especially as they can be of the same order of amplitude as the transit signals the SPECULOOS team is looking for.

To differentiate the photometric variations in the differential lightcurves related to changes in PWV from those of astrophysical origin, I implemented a correction as part of the automatic pipeline. First, I needed access to accurate, high cadence PWV measurements, which are provided by LHATPRO. LHATPRO (Low Humidity and Temperature PROfiling radiometer) is a microwave radiometer optimised for measuring PWV (from 0 mm to a saturation value

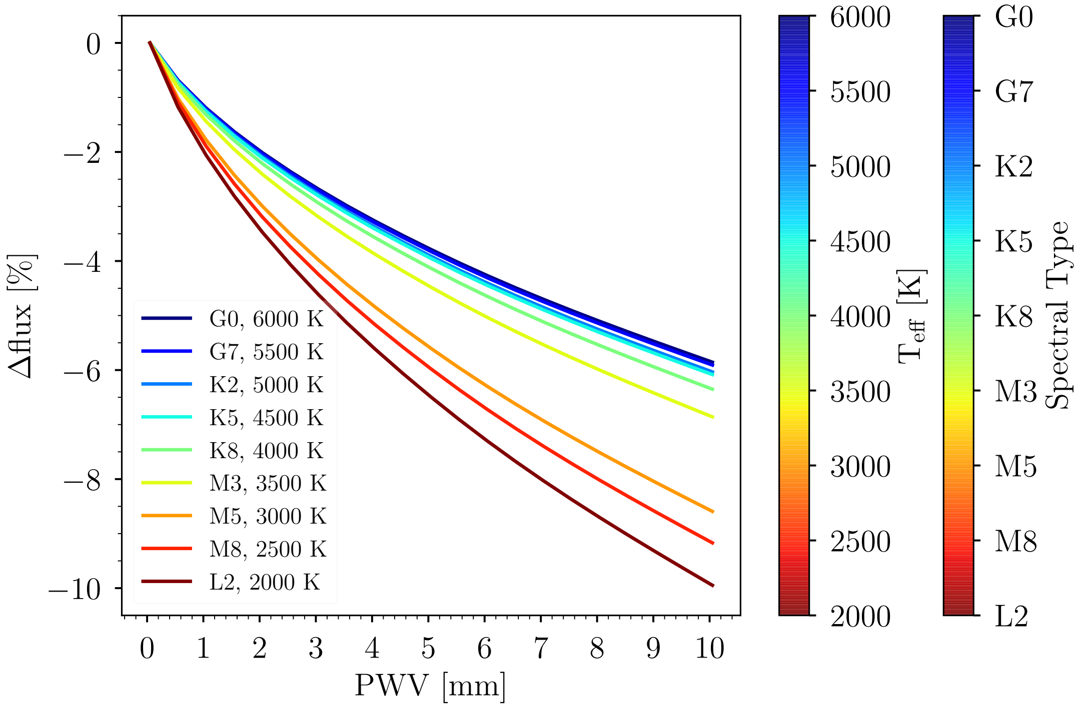

Figure 2.13: Demonstration of the differential flux effect in the $I+z'$ band with changing PWV. For example, an M8 target star will experience a 9% flux drop for a PWV change from 0.05 to 10 mm, whereas G- and K-type comparison stars (the difference is minimal between hotter stars) will only experience a 5–6% flux decrease. This plot was generated with assistance from P. P. Pedersen.

of 20 mm, within an accuracy of ~0.1 mm and with internal precision of 30 μm) situated on a platform at the Very Large Telescope on Cerro Paranal (Kerber et al. 2012). The LHATPRO instrument measures the column of water vapour at zenith approximately every 2 minutes, performs a cone scan at 30° for 2.5 minutes every 15 minutes and a 2D all sky scan for 6 minutes every 6 hours. Due to this cone scan there are peaks in the PWV measurements, which are removed, creating small gaps and discontinuities. I use a cubic spline to interpolate between the remaining PWV values to get a smooth lightcurve correction. As the gaps are on such a small timescale (of the order of ~5 min) I did not view it as a concern to the correction.

By using these PWV values, I can then model the effect of the atmospheric absorption with high time resolution on objects of different spectral types (Section 2.4.1). This allows me to correct the differential PWV effect between the target and comparison stars (Section 2.4.2).

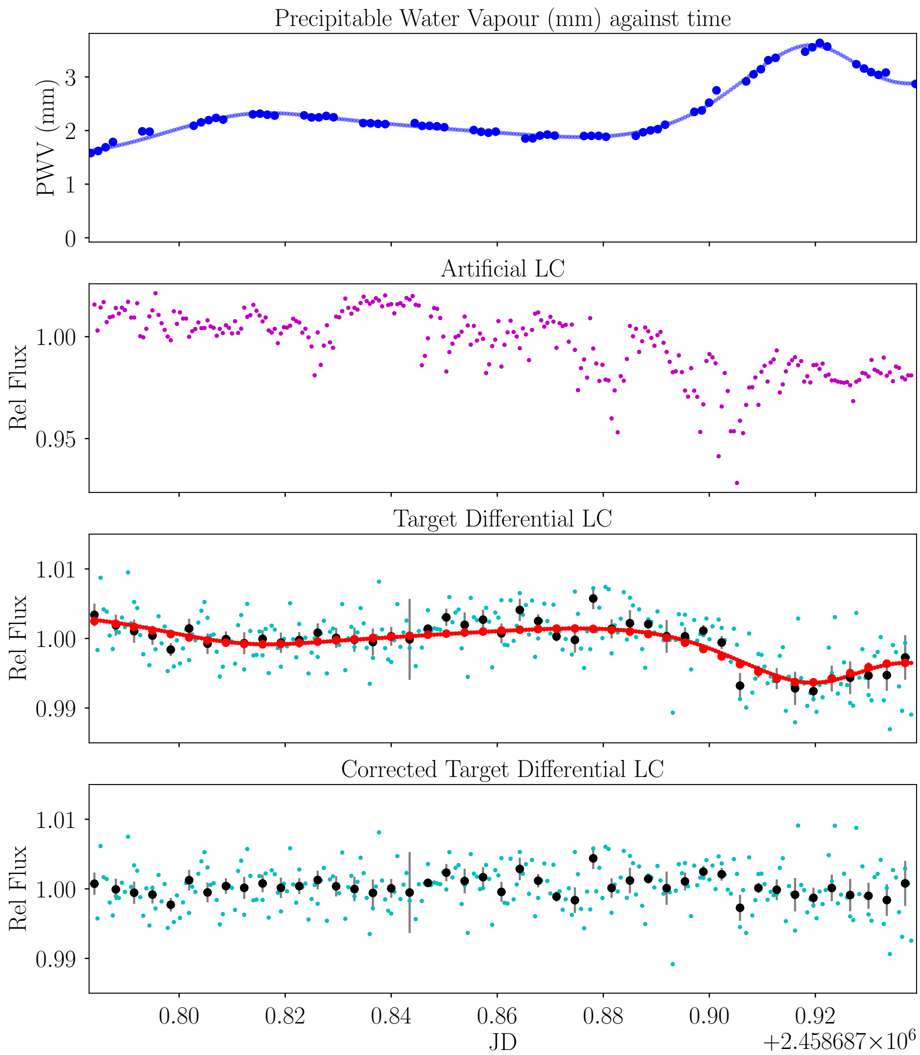

Figure 2.14: **Top**: PWV (mm) measurements from LHATPRO for the night of 22nd July 2019, with peaks removed. The cubic spline interpolation is shown by the blue line. **Upper Middle**: The artificial lightcurve generated by the pipeline for this night. **Lower middle**: unbinned differential lightcurve (cyan), with 5-min binned points (black), for an M7-type target. A shallow transit-like feature is visible at the end of the lightcurve. The expected differential flux effect of PWV is shown in red. **Bottom**: the corrected differential measured lightcurve in cyan and 5-min binned points in black. I obtain this lightcurve by dividing the original differential lightcurve by the calculated differential flux effect from PWV. The transit-like feature was due to PWV changes and is no longer visible in the corrected lightcurve.

2.4.1 Quantifying the PWV Effect

The work in this subsection was carried out in collaboration with P. P. Pedersen. To model the effect of the PWV on differential lightcurves, we calculate its ‘expected’ effect on our measurements for objects of different spectral types, observed with the $I + z'$ filter, at different values of PWV and airmass:

$$f_{I+z'} = \int W(\lambda, X, V)\, R_{I+z'}(\lambda)\, S(\lambda, T_{\text{eff}})\, \mathrm{d}\lambda \tag{2.9}$$

where $W(\lambda, X, V)$ is the water absorption spectrum at airmass X and precipitable water vapour V, $R_{I+z'}$ is the instrument response (including the bandpass for filter $I + z'$, CCD quantum efficiency, CCD window, transmission of the lenses, and reflectivity of the mirror coatings), and $S(\lambda, T_{\text{eff}})$ is the synthetic stellar spectrum generated from PHOENIX (Husser et al. 2013). This stellar spectrum is dependent on the surface gravity, metallicity, and effective temperature T_{eff} of the star. For simplicity we assumed the stars have solar metallicity ($[Fe/H] = 0$).

The water absorption spectrum is provided by the SkyCalc Sky Model Calculator, a tool developed by ESO and based on The Cerro Paranal Advanced Sky Model (Jones et al. 2013; Noll et al. 2012). This tool provides a library of atmospheric transmission curves for a continuous range of airmass values and discrete PWV values of 0.05, 0.1, 0.25, 0.5, 1.0, 1.5, 2.5, 3.5, 5.0, 10.0, 20.0 and 30.0 mm. We interpolate between these value to create a smooth 4-D grid of all possible values of PWV, airmass, T_{eff}, and $f_{I+z'}$, which can be used to correct any object’s differential lightcurve in any frame.

2.4.2 Implementing the PWV Correction into the SSO Pipeline

To correct a target’s differential lightcurve from the impact of PWV, the quantitative effect on both the target and the artificial reference star need to be known. For this purpose, I estimate an effective stellar temperature for the artificial reference star from a weighted mean of the temperatures (extracted from Gaia DR2) of all the comparison stars in the field, using the weights computed by the pipeline in Section 2.2.7. The fact that not all of the comparison stars will have a corresponding Gaia DR2 temperature will have little effect on the correction, as most of the calibration stars are G- and K-type. The differential effect between these spectral types is marginal, even for large changes of PWV (see Figure 2.13).

Having a correct estimate of the target’s effective temperature, however, is more critical. Inaccuracies in this temperature can lead to over, or under, corrections of the flux. Gaia does not provide reliable values for stellar effective temperatures below 3000 K (Andrae et al. 2018; Dressing et al. 2019); therefore, for every target in our target list its temperature has been carefully estimated, as described in Section 2.1.3 (Sebastian et al. 2021). These temperature

Figure 2.15: Cumulative histogram of the amplitude change in a target's differential lightcurve induced by PWV variation on typical transit time-scales (1h). I recorded PWV variations between consecutive 1 hr bins from 1st January – 18th September 2019, and used these variations to generate the corresponding differential flux variations for a 2650 K target object and 4600 K comparison star. I calculated the amplitude variations that we would have a 95% chance of seeing at least one of, on a daily (±0.7 mmag), weekly (±2 mmag), monthly (±4 mmag) and annual (±8.1 mmag) time-scale, marked by the dashed black lines.

estimates are used as input parameters for the pipeline to compute the photometric effect of PWV changes on each target and its corresponding ALC. Finally I divide the PWV effect on the target by the PWV effect on the ALC to generate a differential PWV effect. Then I correct the target's differential lightcurve by dividing by this differential PWV effect. The application of this method is shown in Figure 2.14.

2.4.3 Impact of the PWV Correction

Correction of the PWV effect is a prerequisite to obtain precise differential photometry and to detect shallow transits. This effect impacts the lightcurves over both short (single-night) and long (multi-night) time-scales. During observation of a single night, residuals in the target differential lightcurves may mimic a transit-like signal, even with modest PWV variations of ~1 mm (see Figure 2.14).

By analysing all the PWV measurements from our first ten months of operation, I estimated the likelihood of observing a corresponding differential flux effect large enough to be mistaken for a transit. By averaging the PWV values from 1st January – 18th September 2019 in hour bins (typical time-scale of a transiting planet), the PWV variations between consecutive bins will result in a calculable differential flux effect. In this example I considered a 2650 K target (M7V) and 4600K (K4V) artificial lightcurve. From the cumulative histogram of these differential flux effects (see Figure 2.15) I approximated that there was a 95% chance of observing at least one flux variation (δF) larger than x, using a Binomial distribution:

$$P(\delta F \leq x)^n = 0.05 \tag{2.10}$$

where n is the number of flux variations ($n + 1$ hour bins) and $P(\delta F \leq x)$ is the probability of observing a flux variation less than x. I estimated that there was a 95% chance of seeing at least one amplitude variation of ~1 mmag every night, ~4 mmag every month, and ~8 mmag every year. While these larger variations in the lightcurves may not always resemble transits, they are significant enough to affect our detection of a transit, demonstrating the clear need for an implemented correction. Over multiple nights of observation, correcting for this effect is an absolute necessity to isolate intrinsic variability of our targets from atmospheric transmission changes due to variation of PWV from one night to another, such as in Figure 2.16.

2.4.4 Limitations

I have identified a couple of limitations in the PWV correction, that could potentially leave some residual structures in the final differential lightcurve. The LHATPRO instrument saturates at 20 mm at zenith which will limit the accuracy that can be achieved for very high PWV, especially for high airmass. There is also a ~200 m vertical distance between the VLT platform (2635m) and the SSO facility. Additionally, the LHATPRO instrument measures the water vapour at zenith (where the atmosphere is thinnest) instead of along our line-of-sight. All these factors may result in underestimating the amount of PWV affecting our observations. The effect on our photometry is, however, likely to be small; Querel & Kerber (2014) found that PWV over Paranal was spatially homogeneous down to elevations of 27.5°, such that measuring PWV along zenith is sufficient for most astronomical applications. Concerningly, this homogeneity was found to decrease with rising levels of water in the atmosphere, as they found the PWV variations were reliably 10–15% of the absolute PWV. Therefore, our correction is likely to be most effective at zenith where we don't have to consider spatial variations, and more effective at low values of PWV (<2 mm), where the variations across the sky are of the order ~0.1–0.3 mm. An investigation of the impact from these various effects on our precise photometry is currently underway (Pedersen et al. in prep).

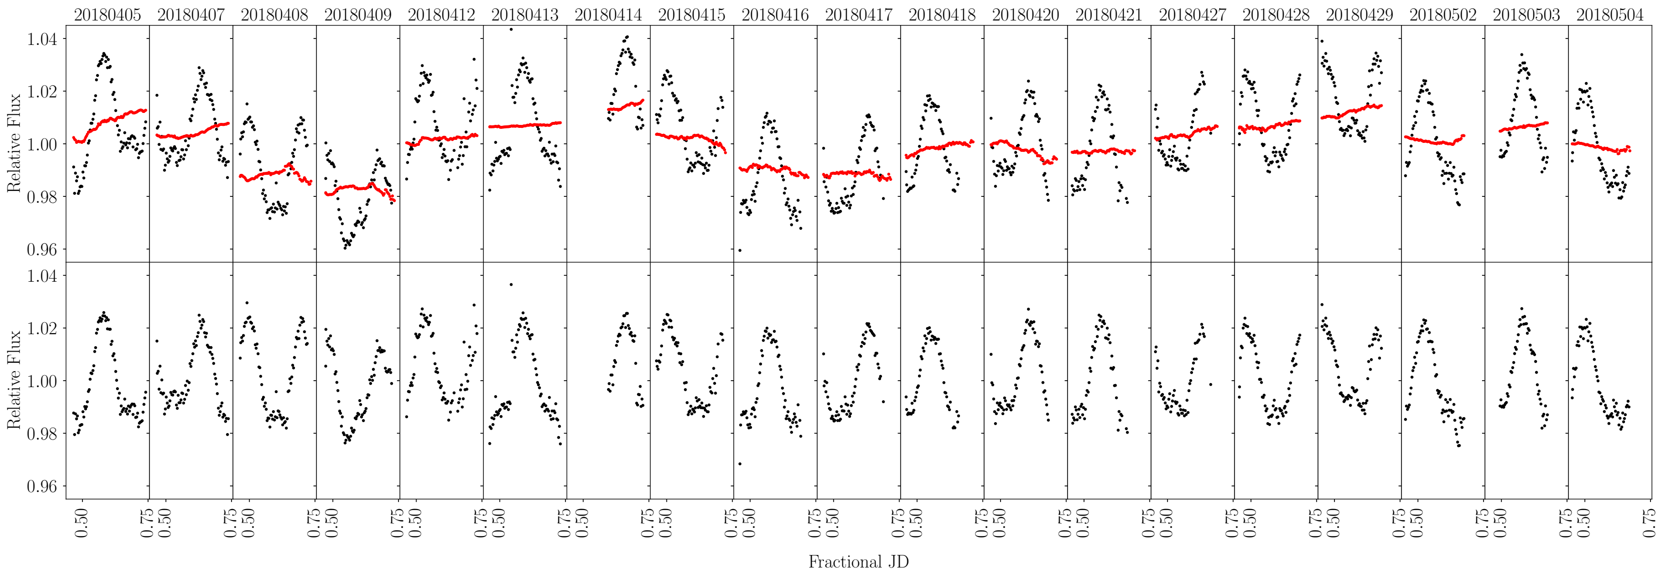

Figure 2.16: **Top:** Global $I + z'$ 5-min binned differential lightcurve for an M8-type variable target (LP 609-24, $J = 12.33$ mag) is shown in black, observed from 5th April – 6th May 2018, during the commissioning phase of Callisto. The calculated differential flux effect from PWV is shown in red (5-min binned). This target exhibits both nightly and multi-night variability. **Bottom:** Water vapour corrected differential lightcurve (5-min binned). While the night-to-night variability remains, the longer time-scale variations were a result of the PWV changes between nights and were removed during the PWV correction.

2.5 Conclusions

In this chapter, I introduced the SPECULOOS survey, a global network of 1m-class robotic telescopes searching for transiting terrestrial planets around the nearest and brightest UCDs. This chapter also illustrated the practical and successful implementation of the automated SSO Pipeline. I developed this pipeline to efficiently process the observations of SPECULOOS-South and to deal with the specialised photometric requirements of ultra-cool dwarf targets. The SSO Pipeline includes a novel differential photometry algorithm with carefully calibrated weighting schemes for comparison stars, and a correction of the effect of varying telluric water vapour. I will assess the photometric quality of our data, and the impact of my work towards optimising the SSO's photometric precision in the next chapter.

While the photometric precisions reached by ground-based transit surveys have improved dramatically over the past 20 years, these facilities are not yet able to detect the shallow 0.01% transit depths produced by an Earth-radius planet orbiting a Sun-like host. Limited by the Earth's rapidly changing weather and atmospheric conditions, current state-of-the-art facilities, such as NGTS and SPECULOOS, have demonstrated the ability to reach photometric precisions of 0.1%. Ground-based transit surveys have previously shown a trade-off between two factors; the size of detectable planet and the photometric quality. While observing Sun-like objects in the visible reduces the systematics caused by the Earth's atmosphere, it limits the smallest detectable planets to Neptune-sized. On the other hand, observing redder objects, such as mid-to-late M-dwarfs, which are faint in the visible and, therefore, must preferentially be observed in the IR or near-IR, allows for the detection of super-Earth and Earth-sized planets. Observing these objects, however, comes with substantial challenges. In the near-IR the varying wavelength-dependent opacity of the Earth's atmosphere has a serious impact on the incoming light. Specifically, second-order extinction effects due to highly variable absorption by atmospheric water vapour has previously limited the quality of the photometry for red dwarfs, as experienced by MEarth (Berta et al. 2012a). In this chapter, I presented a method of modelling and correcting the effect of precipitable water vapour (PWV) during differential photometry. Not only does this correction eliminate the chance of spurious transit-like signals caused by short time-scale changes in PWV, but it significantly reduces the red noise in the photometry.

Several publications have already addressed this telluric water vapour problem when observing cool stars in the near-IR (Bailer-Jones & Lamm 2003; Blake et al. 2008; Blake & Shaw 2011). The MEarth survey has a similar 715–1000 nm bandpass and also witnessed induced photometric systematics that could mimic an exoplanet transit, due to variations in atmospheric water vapour (Berta et al. 2012a). These systematics were also a limiting factor in the type of

M-dwarfs ($< 0.15\,R_\odot$) that could have rotation periods extracted from MEarth (Newton et al. 2018).

Despite identifying the issue, I have not found any implemented correction of telluric water vapour, directly from first principles, in a large-scale survey in the literature. However, I do note that MEarth developed an alternative method of correcting the water vapour effect (Irwin et al. 2011), by medianing all M-dwarf lightcurves gathered by their 8 telescopes (at each site) in half-hour time bins to create a “common mode” lightcurve. A scaling factor is then calculated for each star, determined by a least-squares optimisation. While this method has proved successful to a survey like MEarth (Berta et al. 2012a), which observes dozens of stars every 30 minutes, for the SSO, which only ever observes a maximum of 4 M-dwarf targets at once, this technique is more limited. Correcting the water vapour from the transmission spectra directly offers the advantage that it is determined from an independent dataset (LHATPRO), and removes the chance of over-fitting real structure. Therefore, I developed a first-principles model to correct for this differential effect which I implemented in the SSO Pipeline. This work highlights also how beneficial it is to have access to high time resolution, high precision PWV measurements. However, this correction has a wider impact than just the SSO lightcurves. For example, it could be applied to any future transit survey observing redder stars in the near-IR, including earlier M-dwarfs, or more generally, in long-term photometric variability studies of red objects.

Not every facility has access to expensive water vapour radiometers and so there has been substantial development of alternate methods of measuring the PWV. Instruments like aTmCam (Li et al. 2012, 2014) and CAMAL (Baker et al. 2017) use a set of imagers to take simultaneous observations of bright calibration stars with different narrow-band filters chosen to be in-band and out-of-band for water. Along with measurements of local surface pressure and temperature, GPS receivers have also been used to estimate the atmospheric PWV to accuracies of 0.11–1 mm (Bevis et al. 1992; Duan et al. 1996; Blake & Shaw 2011; Castro-Almazán et al. 2016; Li et al. 2017). I have shown in this chapter that even modest changes in PWV of 1 mm are sufficient to limit our detection efficiency and can mimic a transit from an Earth-sized planet, therefore, accurate PWV measurements are essential. As an alternative to correcting the effect, it is possible to minimize the impact of water bands in the near-IR (and the photometric consequences from changing PWV) by reducing the filter band-pass, but at the cost of losing stellar photons and thus requiring larger telescopes.

The SSO Pipeline was designed to run every day automatically, with little to no human intervention. This pipeline has now also been applied to process data from our Northern facility each night. As the developer and maintainer, I have led one peer-reviewed publication on the SSO Pipeline, which has generated 12 citations. In seven of these papers lightcurves generated by this pipeline contributed directly to the results. The lightcurve products from each night are automatically synced with the online SPECULOOS Portal (designed by P. P. Pedersen) which displays interactive lightcurves from all SPECULOOS facilities, including SAINT-EX. This tool is used by the entire SPECULOOS team to inspect the data by eye and manually flag interesting features and potential transits. Feedback and quality checks from the pipeline show the results of pipeline or instrumental changes in real-time. They can also help assess whether SPECULOOS is meeting the photometric precision goals necessary for the detection of small planets, which I will delve into in the following chapter.

CHAPTER 3

Assessing the Photometry and Performance of SPECULOOS-South

Since the start of scientific operations, I have tracked the quality of the SSO's photometry. The ability of the automatic SSO pipeline to provide consistent and reproducible results has allowed me to carry out daily monitoring of the photometric performance and health of the overall system. Using pipeline feedback, I optimise the photometric pipeline and assess whether the facility is reaching the expected performances set out by the survey goals. The content of this chapter includes some of the work presented in Murray et al. (2020) and the definition of the data package described in Murray et al. (2022). This chapter details the various stages involved in assessing the performance of the SSO Pipeline and the photometric quality of the SSO survey during its first few years of operation. I also put the photometric precision of SPECULOOS into context with other ground-based surveys and space missions to show how, for observing red dwarfs, SPECULOOS is at the forefront of the field.

3.1 Global Observing Statistics from the SSO

For this study, I defined a dataset spanning from 1st June 2018 to 23rd March 2020. While this start date is before official scientific operations began, it marks a point of stability in the commissioning phase. After this date the SPECULOOS team performed no major maintenance to the observatory; DONUTS, our auto-guiding software for precise pointing, was in use; and the operating strategy had been finalised (see Sebastian et al. 2021 for details). The end of the data package was defined as the date that ESO Paranal Observatory shutdown due to the coronavirus pandemic.

In the period between 1st June 2018 and 23rd March 2020, there were 661 potential nights of observation. Ganymede, the final telescope installed in the SSO, started commissioning on 30$^{\text{th}}$ September 2018, therefore, between the four telescopes there is a cumulative total of 2523 nights. However, the SSO lost 23–24% of observing time from (predominantly) bad weather and human error.

3.1.1 Targets Observed in the SSO Data Sample

Over the 1931 (combined) nights observed in this sample, the SSO has observed 176 unique photometric targets for at least one night. These observations have typical exposure times of 20–60 seconds. From Sebastian et al. (2021) I confirm that none of the target list objects are known binaries. I remove any objects that are not part of SPECULOOS's usual "survey mode", such as targets observed for follow-up and monitoring of TRAPPIST-1's transits. I note that in SPECULOOS's "survey mode", there are no simultaneous observations with multiple telescopes during the course of a night. During this time period the operational temperature of the CCD was increased from −70 to −60°C in October 2018. The choice to raise the temperature of the CCD was taken due to the effect on the quantum efficiency of the detector, to improve the SSO's sensitivity at the red limit, while the small increase in dark current was not found to noticeably affect the differential lightcurves. This temperature change introduces an offset in the differential flux between the nights before and after the change was made. Whilst this offset has no consequence on night-by-night analysis, it affects long-term photometric trends. Therefore, when recovering rotation periods (in Section 4.3), I split the lightcurves that straddle this temperature change and analyse the before and after sections independently. As this temperature change happens during the first few months of the dataset, the majority of targets were either observed entirely before or after it. Therefore, only a handful of target lightcurves need to be split. If a target has been observed by multiple telescopes I do not combine their lightcurves as each will experience slightly different instrumental systematics. However, I do check that the rotation periods I estimate in Section 4.3 are present in observations from all telescopes.

Choosing only the objects which have been observed more than 20 hours with one telescope, provides us instead with 154 targets. The observation time is defined as the sum of the span of all observations (start of night to end of night), where any gap longer than 15 minutes is excluded. These objects cover a range of M and L-dwarfs in spectral type, extending from M4 into the early L-dwarf regime (with masses of 0.07–0.2 $M_{\odot}$), as shown in Figure 3.1. All objects in this sample will, therefore, be fully convective, as the convection limit occurs around 0.35 $M_{\odot}$ (Chabrier & Baraffe 1997). While M4, M5 and M6 objects are not classically

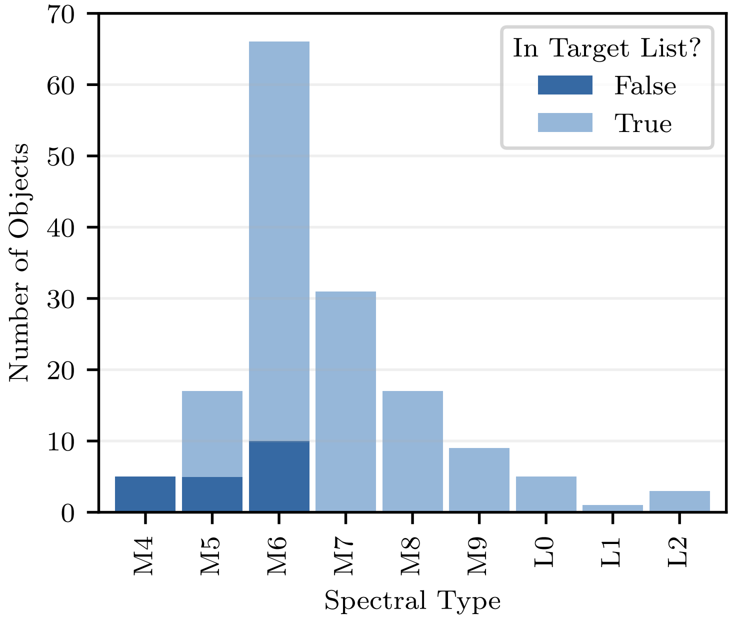

Figure 3.1: Spectral type distribution of the 154 objects in the SPECULOOS-South data sample, observed for more than 20 hrs from 1st June 2018 to 23rd March 2020. I separate objects based on whether they are in the SPECULOOS target list defined in Sebastian et al. (2021).

considered ultra-cool dwarfs, I include them in this sample to explore any differences between mid-M, late-M, and L dwarfs. The majority of the data sample (139 objects, ~85%) are in the SPECULOOS target list, as defined in Sebastian et al. (2021).

This target list is divided into three main scientific programs, which are described in more detail in Sebastian et al. (2021) and Section 2.1.3. The objects in the SSO sample which are not in the target list are exclusively objects with spectral types M4–M6. It is likely these objects were part of the commissioning phase of the telescopes or were removed from the target list (due to reclassifying their spectral type) before official scientific operations began. For the objects in the target list, I extract radii, masses, effective temperatures, and spectral type classifications from Sebastian et al. (2021). The stellar parameters for non-target-list objects are calculated in the same way.

In this sample 80% of the objects have been observed for less than 77 hours (see Figure 3.2), and 50% have been observed for less than 51 hours.

3.1.2 Cleaning the Lightcurves

Once high precision lightcurves have been produced for each of the targets in the SSO data sample, they have to be 'cleaned'. Cleaning involves removing 'bad' observations, or frames in

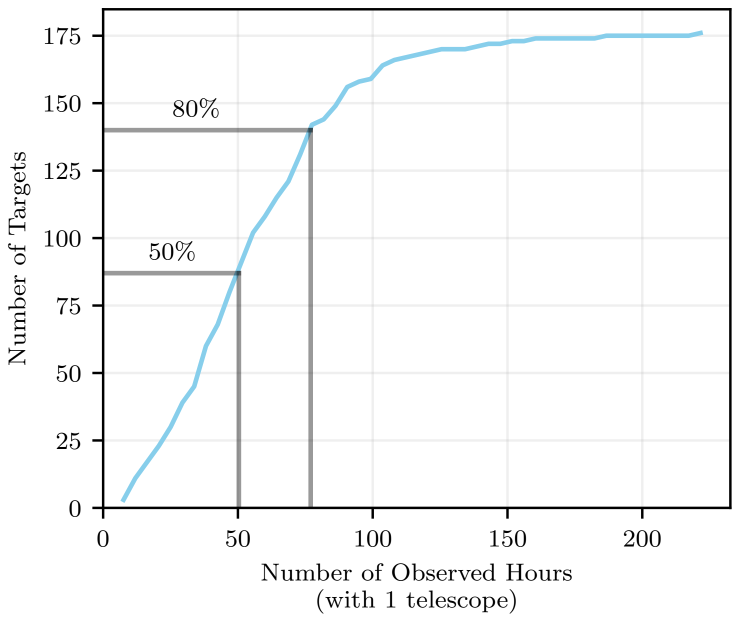

Figure 3.2: A cumulative histogram of the number of hours observed for the objects in the SSO data sample. Of the 176 objects, 13% (22 objects) have been observed for less than 20 hours, 50% have been observed for less than 51 hours, and 80% have been observed for less than 77 hours.

which all stars in the field would behave spuriously, leaving only good quality lightcurves that can be searched for transits. In the context of this thesis, 'bad' observations are defined as those where the atmospheric conditions of the night did not prevent us observing, but significantly affected the quality of the lightcurves. In this case distinguishing real stellar variability or planetary transits from ground-based systematics would be extremely difficult.

To reduce the impact of these sub-optimal observations on future lightcurve analysis I implement a bad weather flag, defined as follows. This flag is not related to any specific external monitoring of the weather. While in theory differential photometry should allow us to correct for any change of atmospheric transmission, empirically there is a practical limit to this assumption. I found there was a threshold for the local RMS of a data-point in the ALC, above which the local RMS of the corresponding data-point in the target's differential lightcurve increased dramatically. The local RMS of a given data-point in the lightcurve is defined as the RMS measured when considering a time range (or box) of ±0.005 d (~7.2 min) around that point in time. I use a relatively small time range to assess the photometric scatter to reduce the impact of real stellar variability or structure. Combining many nights of data allowed me to determine a threshold of 8% to flag (not remove) bad weather in the lightcurves (see Figure 3.3), above which I cannot consistently obtain high photometric precision. The bad weather

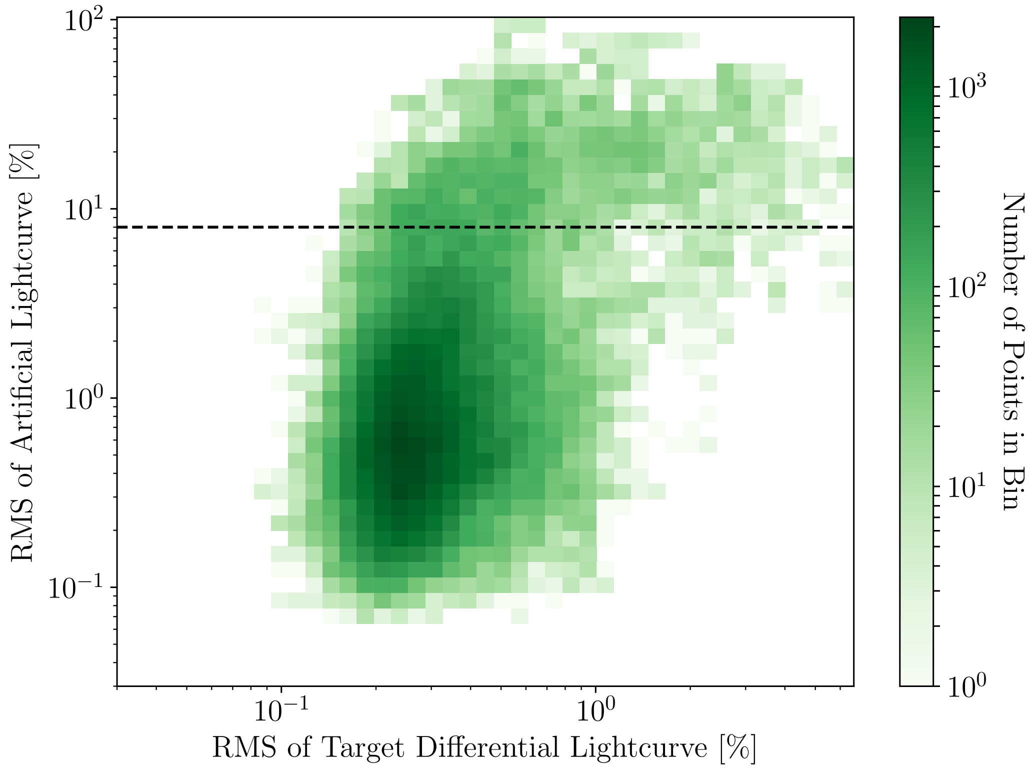

Figure 3.3: Local RMS of unbinned artificial lightcurves (ALC) against local RMS of unbinned target differential lightcurves for all Io's observations from 1st January 2019 to 18th September 2019. For this analysis, there is no water vapour correction or removal of cosmic rays hits, flares or variability, which may cause times when the RMS of the ALC is low, but the target lightcurve has an RMS of a few per cent. There is a much larger variation in quality of a target differential lightcurve when the local RMS of its corresponding ALC exceeds the threshold of ~8%, shown by the black dashed line.

flag is used to remove several data points of the global lightcurve in Figure 3.4.

I also include strict cuts on specific observation parameters. I exclude data points in the lightcurves where the sky background level is greater than 4000 counts per pixel, the airmass is greater than 2.5 or the full width at half maximum (FWHM) of the point spread function (PSF) is greater than 7.5 pixels (i.e., a seeing greater than 2.5 arcseconds). I demonstrate a couple of these cuts in Figure 3.4.

From 1st May to 19th June 2019 the SSO experienced an issue with moving dust on the CCD of Ganymede. While stationary dust can be easily corrected by flat images taken at the start and end of each night, if dust moves across the CCD window during the night, this leads to residuals in the flat-field correction, and structures in the final differential lightcurves. These structures can mimic planetary transits, though the times affected by dust are easy to identify from the raw images. However, as I do not currently have a robust correction technique for moving dust, I removed all observations taken by Ganymede during this period. This amounts to less than 2% of total observations, and so the impact is marginal.

3.1.3 Blended Targets and Crowded Fields

The "blending" of a target's flux with that of a nearby star (see Appendix B) is a major obstacle to obtaining high precision photometry. In this scenario, at small apertures a fraction of the target's flux is lost, and at larger apertures additional flux is included from the nearby star. Therefore, when a target is blended all apertures show a correlation between their flux and the FWHM of the night, and so there is no optimal aperture because every lightcurve has some contamination. Not only is the blending of the target object an issue, but also if the field of view is extremely crowded then a high proportion of the comparison stars are likely to be blended as well. While, in a normal field, these stars would be weighted down for large apertures (due to their high photometric variation), if most stars in the field are blended then there are only a small number of usable comparison stars. This limitation can degrade the quality of the ALC, and also lead to large variations between apertures in the number of comparison stars and the photometric quality, even when the target itself is not blended.

PSF-fitting photometry would be necessary to extract high precision lightcurves under these conditions. However, as this has not been developed for SPECULOOS, I am currently unable to produce good quality photometry for blended targets. In the SSO sample all objects that I have manually flagged as blended have already been removed. A future improvement to the pipeline would be to include an automatic blended flag for every star in the field.

3.1.4 Quality Checks

There are a number of outputs from the SSO Pipeline that can be used to monitor the degradation of the health of electrical components and photometric quality over time. The parameters that can be used as quality checks include the readout noise, dark current, and transformations between different photometric systems. The readout noise and dark current are as defined in Section 2.2.3 and demonstrated in Appendix A. The transform is described below. Flagging sporadic values can alert us to instrumental issues. For example, monitoring of the readout noise and photometric transform was used to identify when the gain had mistakenly changed on Europa, between July and September 2019. Additionally, when the operational temperature of the CCD was increased from −70 to −60°C in October 2018, the subsequent increase in dark current could be tracked.

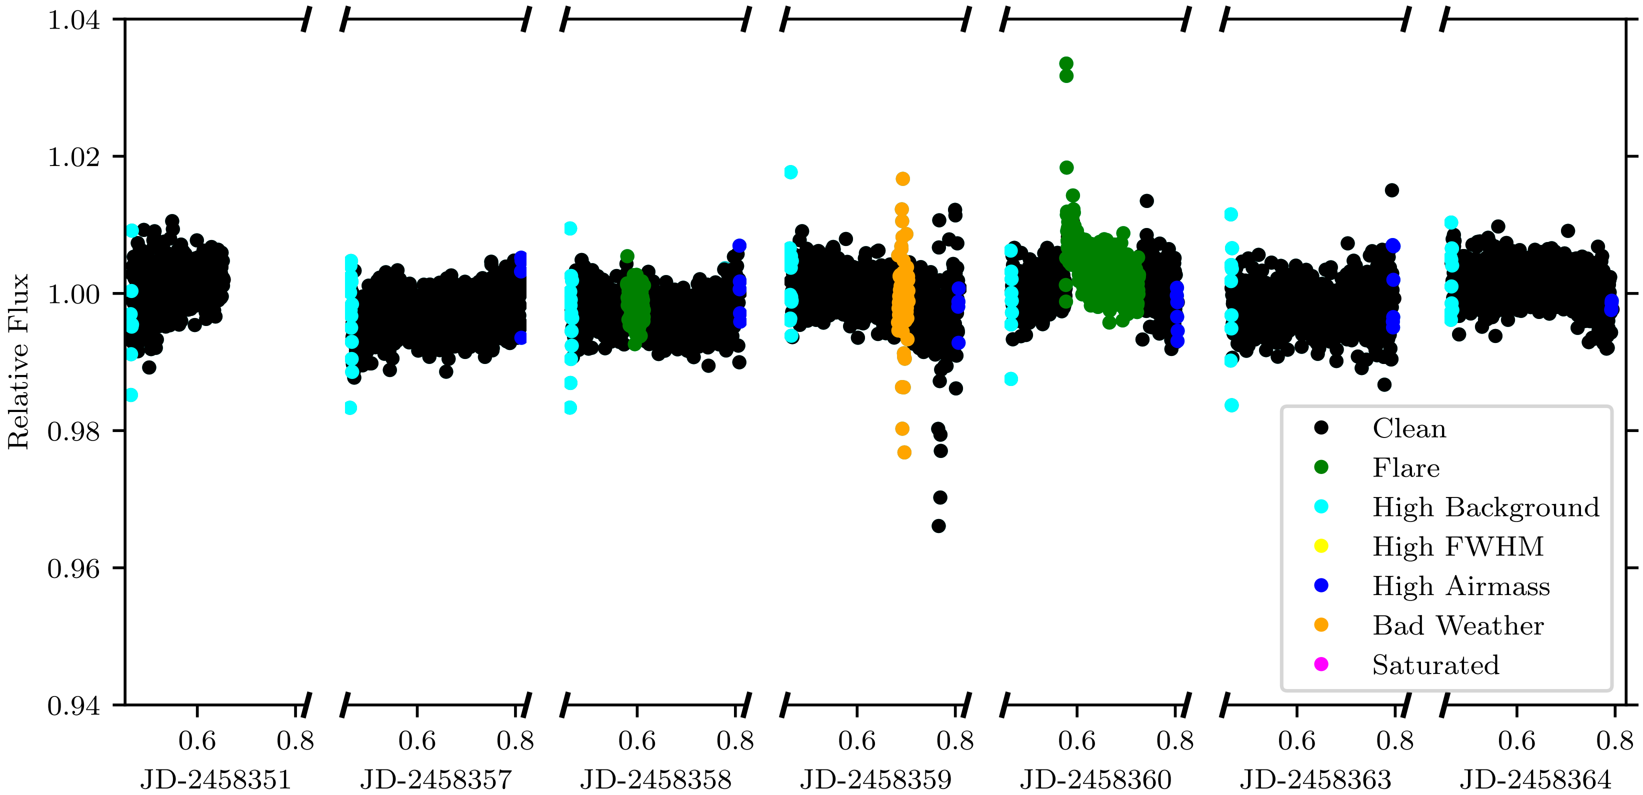

Figure 3.4: Demonstration of the need for cleaning the lightcurve post-differential photometry. The black points are the unbinned data points that would remain after the lightcurve has been cleaned. A couple of the observation parameter flags are triggered in this global lightcurve: flares detected by the automatic detection pipeline described in Chapter 4 (in green), high sky background (in cyan), high airmass (in dark blue) and bad weather (in orange). There are several other flags, such as saturated points, or high FWHM, which are not triggered by this lightcurve. For improved visibility, I have removed the time between nights.

At −70°C the dark current was on average less than 0.04 $e^{-}s^{-1}pixel^{-1}$, whereas, at −60°C the dark current ranges from 0.2–0.4 $e^{-}s^{-1}pixel^{-1}$ (depending on the telescope). Quality checks can also be used to detect slight differences between the four telescopes.

It is also essential to track feedback about the photometric quality of the SSO. Over time the CCD sensitivity will decline due to dust building up on the window and deterioration of the mirror coatings. One way to monitor the average gradual flux decrease is to consistently compare to magnitudes measured by other surveys.

Converting Flux to Magnitude

The apparent magnitude, or simply magnitude, of a star is a measurement of its brightness from the Earth. Magnitude is often used in astronomy as an alternative to flux. How bright a star appears is dependent on its intrinsic brightness (or luminosity), how far it is away from the Earth, and any transmission effects the light experiences on its path to the Earth's surface.

The relationship between stellar flux and magnitude is as follows (Carroll & Ostlie 1996):

$$m_1 - m_2 = -2.5 \log \frac{F_1}{F_2}, \tag{3.1}$$

where m_1 and m_2 are the apparent magnitudes of two stars with corresponding flux densities F_1 and F_2. Two stars with a difference of 5 magnitudes would have a flux ratio of 100 (one star would appear 100 times brighter than the other). The apparent magnitude of an object that produces a count rate of one count per second defines the "zero-point" magnitude, used to calibrate an instrument's sensitivity.

Each survey uses its own narrow passband (or filter) to collect photometric information. As stars emit light with wavelength-dependant intensity, each filter will measure a different brightness. The brightness we measure will also depend on the instrumental set-up. Therefore, we need a calibration to convert between these different instrumental set-ups. The magnitude that a detector would measure in a specific wavelength band (for example, $I + z'$) above the Earth's atmosphere would follow:

$$m_{I+z'} = -2.5 \log \int_0^\infty F(\lambda) R_{I+z'}(\lambda) d\lambda + C_{I+z'}, \tag{3.2}$$

where $m_{I+z'}$ is the magnitude, $C_{I+z'}$ is a constant which sets the zero-point of the magnitude scale, F_λ is the stellar flux and $R_{I+z'}$ is a sensitivity function describing the instrumental response as a function of wavelength (fraction of stellar flux that is detected), all for a given filter passband and instrumental set-up ($I + z'$ in this case). For a ground-based observatory, we would also have to take the wavelength-dependent atmospheric transmission into account.

Transforming between Photometric Systems

We can convert magnitudes between similar photometric systems using empirically-derived corrections. The corrections can also be a valuable tool to track the degradation of SPECULOOS detectors. Dust builds up in the telescope, and the mirror coating will degrade over time, affecting the camera's sensitivity to light. This can be improved by regularly cleaning the mirror, or re-coating it.

The transformation between $I+z'$-mag and G_{R}-mag (the magnitude in Gaia's red passband) can be easily derived by using an equation for a straight line where the correction, m_{T}, is the y-intercept:

$$m_{G_{\mathrm{R}}} = m_{I+z'} + m_{\mathrm{T}}. \tag{3.3}$$

To measure the correction accurately I use as many stars as possible in each field, to lower the relative contribution of each star. However, I needed to introduce a number of cuts to remove stars which might have a variable magnitude or are significantly different from other stars in the field. Therefore, I remove any stars which exhibit stellar variability, are within 20 pixels of the CCD edge, are too blue or too red ($G - G_R < 0.8$ mag and $G_B - G_R > 0.25$ mag), or are faint ($G > 19$ mag). As all stars in the field have different colours we will see slight variations in m_{T} for each star. I mitigate this effect by restricting the colours of the sample of stars I am fitting to a small colour range (as described above), and assuming that the straight-line fit will represent an "average" star in the field. The frames used also have to be of good quality, so I only include observations which have a focus value close to the mean (to remove highly defocused fields), which are not flagged as bad weather, where the airmass is less than 1.75, and the background sky level is less than 1000 ADU pixel^{-1}.

Correcting for Extinction

I also need to consider that the magnitude measured on the ground is not the magnitude that one would measure in space. I have already discussed the impact of the Earth's atmosphere on our lightcurves in detail (see Section 2.4), however, that is in the context of second order extinction effects in differential photometry, where we compare stars of different stellar types with light travelling through the same approximate region of the atmosphere. Here I am comparing a star to itself (so I do not have to worry about different stellar spectra) but at a different time, when the atmosphere will not be the same. Objects will appear to move across the sky during the night and the more atmosphere that light has to travel through, the more it will be attenuated. Therefore, first order extinction will have a significant effect on magnitude measurements. This can be seen by a variation of the transformation over the course of the night, correlating with the airmass.

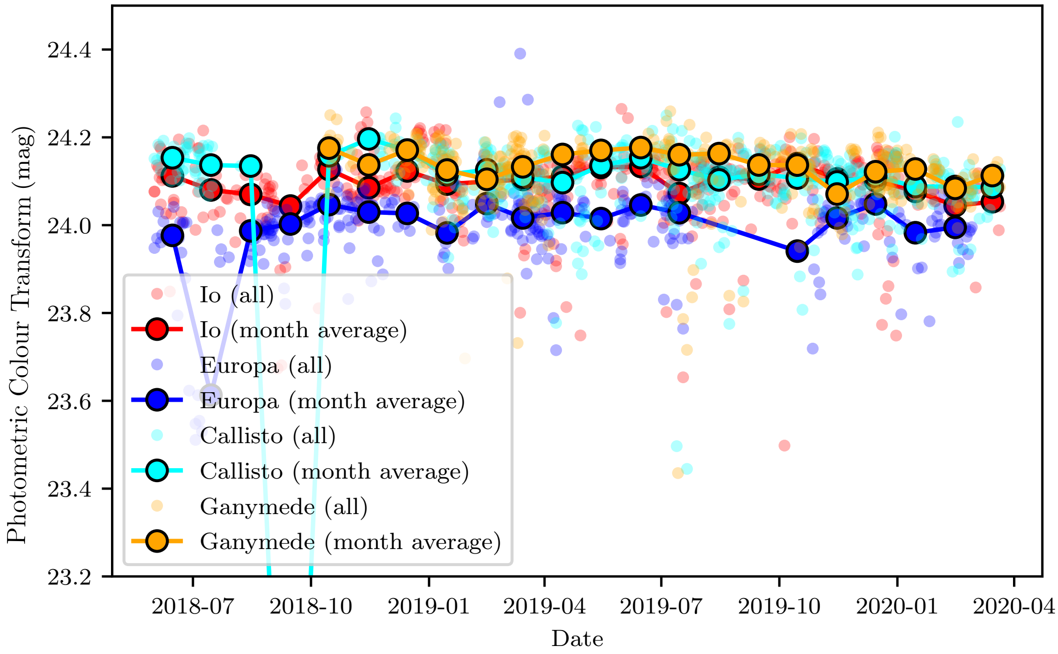

Figure 3.5: The photometric transforms between the $I + z'$-band and Gaia's G_{R}-band for each telescope against date for the time span of the SSO data sample. Individual nights' corrections are plotted as small, semi-transparent circles, whereas the monthly averages are plotted with larger, opaque circles.

I can represent this by adding an airmass-dependent term[a] to the calculation of the transform: $k(X - 1)$, where X is the airmass, and k is the extinction coefficient that I can calculate empirically. As we cannot measure $m_{I+z'}$ above the Earth's atmosphere, I only consider the magnitude difference from a minimum airmass of 1 (at zenith). By measuring the gradient of a plot of the correction against the airmass I can derive the extinction coefficient and adjust the correction for different airmasses. In collaboration with D. Sebastian, I find that $k = 0.05\, m_{I+z'}$ per airmass.

I measure the photometric transform by taking the average of almost all the stars in the field. This will produce a second order extinction effect I have not accounted for. If we consider a field with on average 'bluer' stars than another field, the atmospheric absorption will be different from a field that has on average 'redder' stars. However, by removing the reddest and bluest stars, and averaging over all stars in a field and all nights in a month (during which we observe many different fields), this effect should be very small. The times when Ganymede was affected by dust have been removed as the dust blocks flux, artificially reducing stellar magnitudes.

The calculated corrections for each of the four telescopes are shown in Figure 3.5, over the

[a] `http://slittlefair.staff.shef.ac.uk/teaching/phy217/lectures/principles/L04/index.html`

defined duration of the SSO dataset. I calculated the median correction as 24.16, however, it is clear from Figure 3.5 that there are slight offsets between each telescope. There can be large variations in the correction if we observe a field that is not ‘photometric’, such as one that is extremely crowded, or defocused.

3.2 Photometric Precision of the SPECULOOS-South Observatory

3.2.1 Calculating the Theoretical Noise

We can model the theoretical noise limit with a model taking into account the major contributing noise sources. This noise model (Merline & Howell 1995) accounts for several different contributions: Poisson noise from the star, read noise from the detector, noise from background light, noise from dark current, and atmospheric scintillation. If I assume the sources of noise are independent and Gaussian then the variances can be summed to give the total expected standard deviation of each data point:

$$\sigma_{\mathrm{total}} = \sqrt{\sigma^2_{\mathrm{p},*} + \sigma^2_{\mathrm{p,sky}} + \sigma^2_{\mathrm{sci}} + \sigma^2_{\mathrm{dark}} + \sigma^2_{\mathrm{R}}}, \tag{3.4}$$

where $\sigma_{\mathrm{p},*}$ is the photon noise of the source, $\sigma_{\mathrm{p,sky}}$ is the photon noise of the sky background, σ_{sci} is the scintillation noise, σ_{dark} is the dark current noise (equivalent to the photon noise in thermally excited electrons), and σ_{R} is the readout noise.

The expected fractional scintillation variance, $\sigma^2_{\mathrm{sci,frac}}$, can be found from a modified version of Young’s approximation (Young 1967; Osborn et al. 2015):

$$\sigma^2_{\mathrm{sci,frac}} = 10 * 10^{-6} \cdot C_{\mathrm{Y}}^2 D^{-4/3} X^3 t_{\mathrm{exp}}^{-1} e^{-2h/h_0}, \tag{3.5}$$

where C_{Y} is a constant dependant on the observing site, D is the diameter of the telescope in cm, X is the airmass, t_{exp} is the exposure time, h is the altitude of the observatory and h_0 is the scale height of atmospheric turbulence, for which I used 8000 m for Paranal as in Osborn et al. (2015) and O’brien et al. (2021).

Each of the contributing error sources in Equation 3.4 has a different dependence on observables; therefore, I need to convert the measurable parameters into the same format:

$$\sigma_{\mathrm{total}} = \sqrt{S_* G t_{\mathrm{exp}} + S_{\mathrm{sky}} G t_{\mathrm{exp}} n_{\mathrm{pix}} + (\sigma_{\mathrm{sci,frac}} S_* G t_{\mathrm{exp}})^2 + d t_{\mathrm{exp}} n_{\mathrm{pix}} + R^2 n_{\mathrm{pix}}}. \tag{3.6}$$

The gain, G, is a conversion factor between photon counts and electrons. The exposure time of an image, t, provides the timescale over which photons are collected and n_{pix} is the number of pixels in an aperture. S_* is the total number of photons s^{-1} collected in an aperture around the

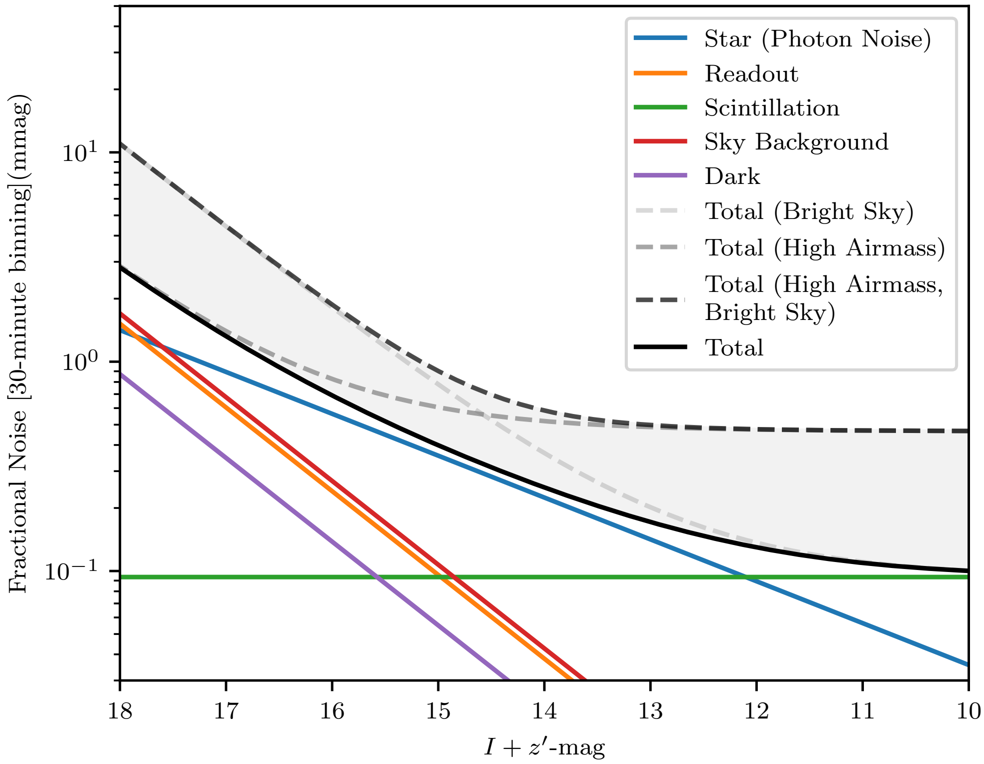

Figure 3.6: Theoretical model for the relative contributions of the major noise sources against stellar magnitude for SPECULOOS observations at Paranal observatory. The binning is 30 minutes. The constant green line is the scintillation, which depends on airmass and exposure time (in this case 40 s) but not stellar brightness. The purple line is the dark current (~0.33 e^- s^{-1} pixel^{-1} at -60°C) and the orange line is the readout noise (6.3 e^-), both of which are intrinsic to the instrumental set-up. The photon noise from the sky background, which varies from night to night depending on cloud cover and proximity to the Moon, is shown in red. The blue line is the fundamental photon noise of the stellar target. The solid black line is the total combination of these five main noise sources for a dark sky (50 ADU pixel^{-1}) and low airmass (X = 1 at zenith). The dotted grey lines show the total noise if the sky was bright (2000 ADU pixel^{-1}) or if the airmass was high (X = 2.5, i.e., when the target is far from zenith), and the dotted black line shows the noise limit if there was both a bright sky and high airmass (such as for observations at the end of the night). The shaded region, therefore, shows the variation of the noise model for a range of observing conditions.

source, whereas S_{sky} is the average sky background per pixel in units of photons $\text{s}^{-1}\text{pixel}^{-1}$. d is the dark current in $\text{e}^{-}\text{s}^{-1}\text{pixel}^{-1}$, and R is the time-independent readout noise in e^{-} pixel^{-1} (equivalent to a standard deviation). If the data is binned in time then the error becomes (with n_{bin} points in each bin):

$$\sigma_{\text{total,bin}} = \frac{\sigma_{\text{total}}}{\sqrt{n_{\text{bin}}}}, \tag{3.7}$$

assuming that the points in each bin are independent of one another. For $I + z'$-magnitudes brighter than 17, this is indeed the case, as σ_{total} is dominated by the star's photon noise and scintillation (see Figure 3.6). However, for fainter magnitudes, the sky background term, which depends on clouds and the proximity of the Moon, will dominate. The sky background noise may not be completely time-independent from one bin to the next, and, in this case, Equation 3.7 would not be applicable. The signal-to-noise ratio follows from Equation 3.6:

$$\begin{aligned}\text{SNR} &= \frac{S_{*}Gt_{\text{exp}}}{\sqrt{S_{*}Gt_{\text{exp}} + S_{\text{sky}}Gt_{\text{exp}}n_{\text{pix}} + (\sigma_{\text{sci}}S_{*}Gt_{\text{exp}})^2 + dt_{\text{exp}}n_{\text{pix}} + R^2 n_{\text{pix}}}} \\ &= \frac{S_{*}\sqrt{Gt_{\text{exp}}}}{\sqrt{S_{*} + n_{\text{pix}}\left(S_{\text{sky}} + \frac{d}{G} + \frac{R^2}{Gt_{\text{exp}}}\right) + (\sigma_{\text{sci}}S_{*})^2 Gt_{\text{exp}}}}.\end{aligned} \tag{3.8}$$

Using this noise model I can predict the theoretical noise in SSO observations, and use this prediction to evaluate the quality of my data reduction. For the atmospheric scintillation, I use Equation 3.5. Each of the SSO telescopes has a diameter of 1 m, Paranal observatory in Chile has an altitude of 2669 m and Osborn et al. (2015) empirically derived the value of C_{Y}=1.56 specifically for Paranal. See Figure 3.6 for a demonstration of the scale of the various contributing noise sources, with values typical for the SSO. By considering that the observing conditions (specifically the airmass and sky background level) can range wildly between nights, and over the course of a single night, it is possible to explore the range of noise limits that we would reach. Even if there was very high airmass (X=2.5) and a bright sky background (2000 ADU pixel^{-1}), the theoretical noise limit should be sub-millimag for targets brighter than $I + z'$-mag = 15, with 30-minute binning.

However, the theoretical noise model has limits. It does not take into account additional noise within the telescope (such as flat-fielding or pointing errors etc.), though it is possible to mitigate these sources through careful data reduction techniques. Also, as I consider brighter and brighter stars I have to take into account the non-linear regime as saturation of the CCD is reached. The other major factor I have not considered here is the variability of the stars themselves. If a target object's brightness varies on short timescales then it will significantly limit our ability to reach SPECULOOS's noise floor.

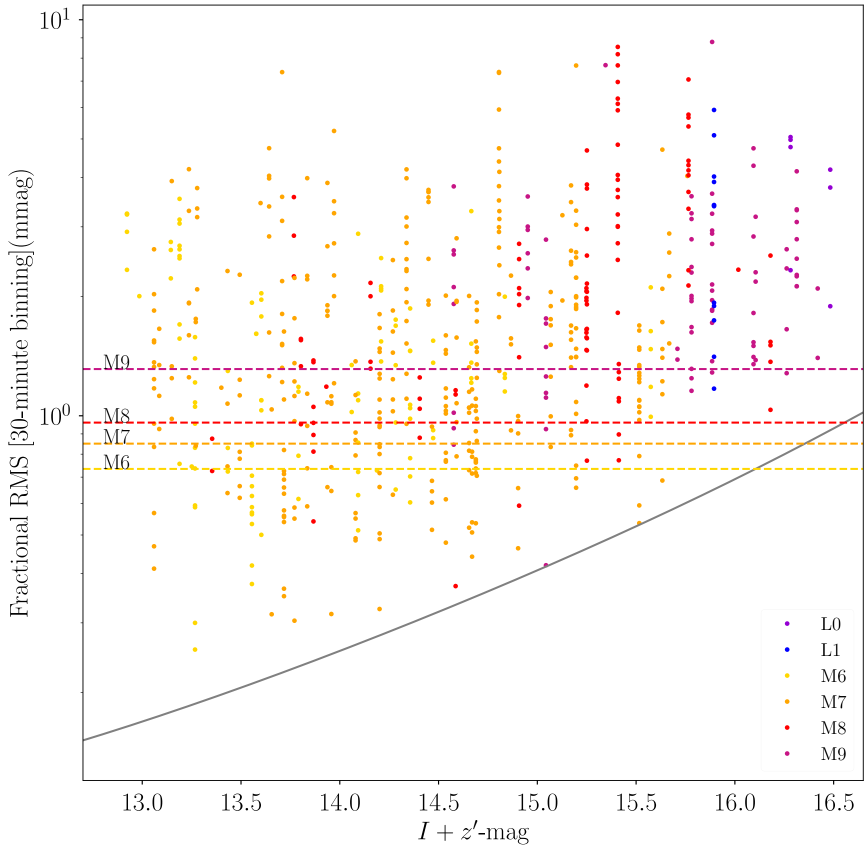

Figure 3.7: Fractional RMS (for 30-min binning) of all the SSO's UCD target lightcurves carried out with $I + z$ filter from 1st January to 18th September 2019 (from Murray et al. 2020). There is a data-point for each target on each night of observation - the vertical spread in RMS correspond to varying fractional RMS on different nights of observation for the same target. The theoretical noise model with the best possible observing conditions is shown in grey. The dashed lines show the minimal level of precision needed to detect a single transit of a TRAPPIST-1b-sized planet (1.127R$_\oplus$, Delrez et al. 2018a) around stars of different spectral types at 9-sigma.

3.2.2 Measured Photometric Precision

In Murray et al. (2020) I presented the typical photometric performances of the SSO facility, and its capability to detect single transits of Earth-size planets, during its first year of operation. In that paper, the dataset accounted for 98 targets and 179 combined nights of observations with multiple SSO telescopes from 1st January to 18th September 2019. However, I have now collected a significantly larger dataset from the SSO, therefore, in this section I will apply the methodology used in that work to the 176 objects in the SSO data sample described in Section 3.1.

I calculate the measured fractional RMS (for 30-min bins) for every night's target lightcurve.

The aperture chosen for each night is determined by the SSO Pipeline automatically and can vary between nights. I include both the plot from Murray et al. (2020) (Figure 3.7) and its updated version in Figure 3.8, to demonstrate that the results remain consistent. To ensure there are at least 5 bins for each lightcurve, I only included lightcurves for which there were more than 150 minutes of total exposure. The binning time-scale I adopted to compute the RMS is set to match the typical transit duration of a short-period planet orbiting an UCD. I remove any timeframes where the airmass is above 1.75, the sky background is above 1000 ADU pixel^{-1}, the FWHM is above 6.5 pixels or is flagged as bad weather. These thresholds are slightly more conservative than those described in Section 3.1.2 as our objective is different. Here, I am considering the typical photometric precisions that the SSO achieves for the majority of Paranal observing conditions, as opposed to previous work, in which I remove the data points particularly afflicted by the atmospheric conditions. All the lightcurves in these two plots have been corrected for the second-order extinction effects of temporally-varying PWV, outlined in Section 2.4.2.

These figures demonstrate that for quiet targets on nights with good observing conditions the SSO is reaching the best possible precision, as determined by the noise model in Section 3.2.1. The targets the SSO observes typically have exposure times from 10–60 s, therefore, I assume the noise model for 60 s exposure, with an overhead of 10.5 s, which gives 25 data-points in each 30-min bin. The noise model illustrated in Figure 3.7 is assumed for an aperture of 11.3 pixels on a "best-case scenario" night, where the airmass is 1 and the background sky level is 49.0 ADU pixel^{-1} (the lowest recorded sky background recorded between January and September 2019). The noise model illustrated in Figure 3.8 is assumed for an aperture of 8 pixels on the best possible night, with an airmass of 1 and a background sky level of 37.8 ADU pixel^{-1} (the lowest recorded sky background recorded between June 2018 and March 2020). The noise model is included only as an indication of the maximum precision we would expect to achieve, however, the actual RMS may reach precisions below the noise model for several reasons. Firstly, there are random measurement errors that could produce extremely small photometric scatter. Secondly, the noise model is a model: it cannot account for all atmospheric conditions, variations and intricacies. Critically, the noise model is also only drawn for one exposure time (60 s) and one aperture size (8 pixels). There are several nights where we may observe with exposure times longer than 60 s (the longest recorded time is $\sim$130 s), or use an aperture smaller than 8 pixels which would reduce the smallest precision to below this model, however, both of these scenarios are not common. The dark current also increased during this observation span as a result of the CCD temperature increase, which means that the RMS on nights before this change could reach slightly lower values. I used the higher dark current of 0.33 e^{-} s^{-1} pixel^{-1} for the model as this is the dark current for the majority of observations in this time span.

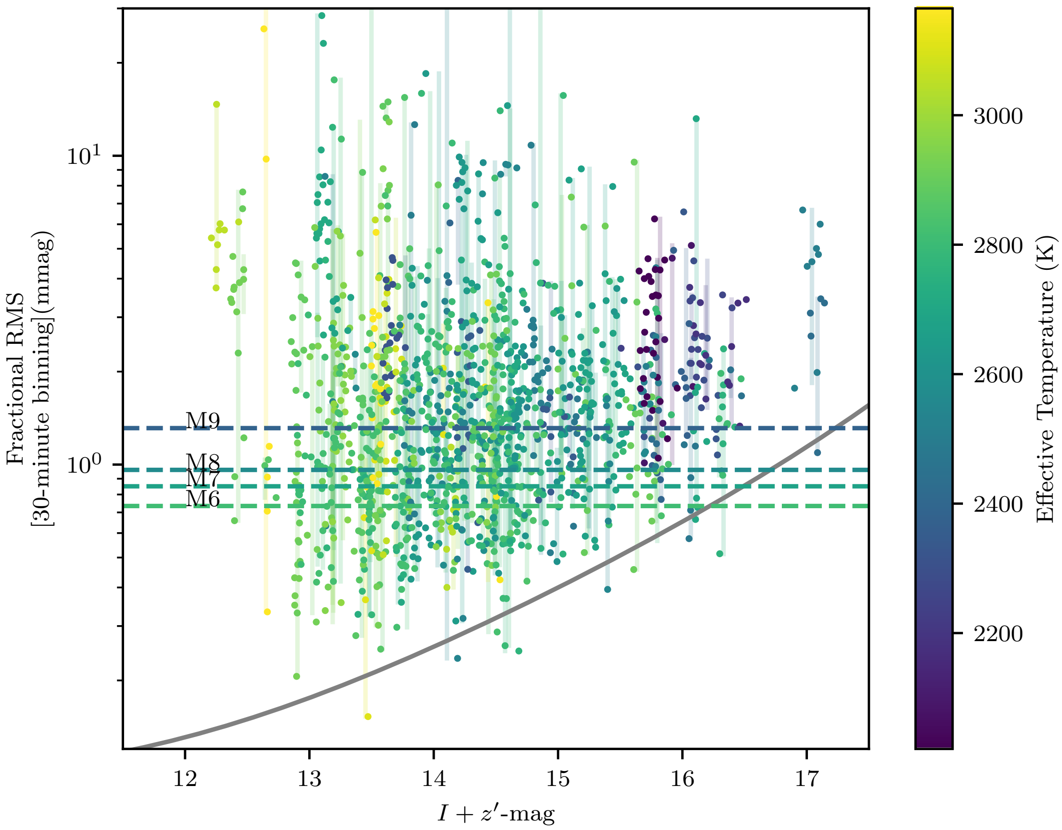

Figure 3.8: Fractional RMS (for 30-min binning) of the SSO's UCD target lightcurves. This plot is the same as Figure 3.7, however, carried out on a larger, updated dataset, from 1st June 2018 to 23rd March 2020. The vertical lines indicate the spread in RMS for each target.

Finally, there is uncertainty in the magnitude measurements, that depend on the accuracy of the photometric transform and the median flux of each star.

There was no correction for photometric variability or selection of the nights with the best observing conditions. This resulted in vertical stripes for each target, corresponding to large spreads in RMS in the lightcurves for different nights. These spreads are related to the wide range of observing conditions and intrinsic stellar variability (that can be irregular or stochastic). The photometric precisions reached for the least active targets in Murray et al. (2020) showed that the SSO is reaching sub-millimag precisions for approximately 30% of nightly lightcurves, with a median precision of ∼1.5 mmag, and up to ∼0.26 mmag for the brightest objects. With the larger dataset, I concluded that the SSO is reaching sub-millimag precisions for around 34% of nightly lightcurves, with ∼1.5 mmag median precision and a maximum precision of ∼0.14 mmag. In this figure, I superimposed an approximation of the minimum photometric precisions required to measure a single transit by a TRAPPIST-1b size planet (1.127 $R_\oplus$, Delrez

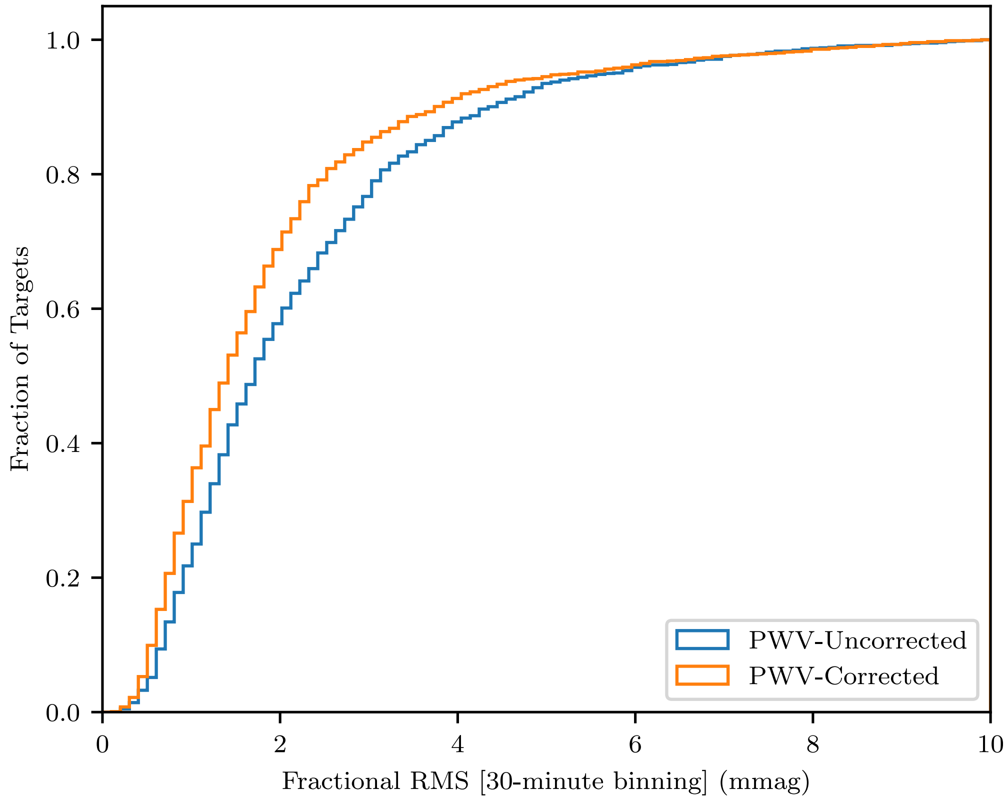

Figure 3.9: A cumulative histogram of the fractional RMS of 30-minute binned SSO lightcurves. The histogram in orange is lightcurves that are corrected for the photometric effects of PWV, and blue is those lightcurves without the PWV correction. Around 24% of uncorrected target lightcurves have sub-millimag precisions, compared to ~34% with the water vapour correction.

et al. 2018a) with a signal-to-noise ratio of 9 for different spectral types, using stellar parameters from Pecaut & Mamajek (2013). The RMS spread due to stellar activity and variable observing conditions can be seen to limit the single transit detection efficiency, thereby demonstrating the need to effectively clean LCs and remove photometric variability. I expect the median precision (and the detection potential) to improve when the stellar variability is properly accounted for. A study of the photometric variability of our objects will be presented in the following chapter.

The Impact of the PWV Correction on Photometric Quality

In the previous chapter, I demonstrated only the impact of the PWV correction on individual lightcurves. Specifically, I showed a night with a large variation in PWV, resulting in a photometric structure that could be mistaken for a transiting planet. Removing this form of contamination is a central motivation for correcting PWV extinction effects. However, the PWV is changing constantly, and so induces continuous photometric variations, not just on nights when features are visibly imprinted. I compare the precisions of all SSO lightcurves before and after PWV correction in Figure 3.9. Without the water vapour correction, I obtained a median precision of ~1.8 mmag for all targets, compared to ~1.5 mmag with the correction. Additionally, ~24% of uncorrected target lightcurves had sub-millimag precisions, compared to ~34% with the water vapour correction. Therefore, properly accounting for telluric effects is a key aspect of reaching the photometric precision goals set out by SPECULOOS.

3.3 Comparison with TESS

NASA's Transiting Exoplanet Satellite Survey (TESS) was launched in April 2018. While TESS is optimised for detecting planets around G to mid M-dwarf stars, their wide bandpass allows them to additionally observe the brightest mid-to-late M-dwarfs with high precision.

Here I present a comparison of a night of simultaneous observation of the M6 star, WOH G 618 (TIC 31381302, J = 10.3 mag, T = 12.5 mag, $I+z'$ = 12.6 mag), by a single SSO telescope to TESS data, in Figure 3.10. For TESS, I include both the publicly available 2-minute cadence data and the final lightcurve from the MIT Quick Look Pipeline (QLP). The output from the QLP was provided by Maximilian Günther and Chelsea Huang, members of the TESS team at MIT. The QLP was developed to extract lightcurves specifically for targets in the 30-minute full frame images (FFIs). It is shown here as an example for FFI photometry, allowing us to gauge the precision that can be achieved for targets which are not part of TESS's 2-minute sample. The QLP and other custom pipelines can be used to extract lightcurves from the FFIs for the majority of mid-to-late M-dwarfs in the TESS fields.

There is excellent agreement between the three data-sets. There remains a periodic structure in the SSO-TESS QLP residuals that correlates with the variability. As this structure does not appear in the SSO-TESS 2-minute residuals (which should be more accurate), I conclude this correlated noise is imprinted by the QLP and not the SSO Pipeline. The SSO lightcurve shows less white noise than TESS, as expected because TESS is not optimised for these very red objects. For fainter and redder UCDs I expect that the quality of the SSO lightcurves will exceed TESS. However, for the brightest SPECULOOS targets the lightcurves will be comparable. This work demonstrates the remarkable precision of both TESS and SSO, especially considering the detection potential when combining simultaneous observations from multiple SSO telescopes and TESS together.

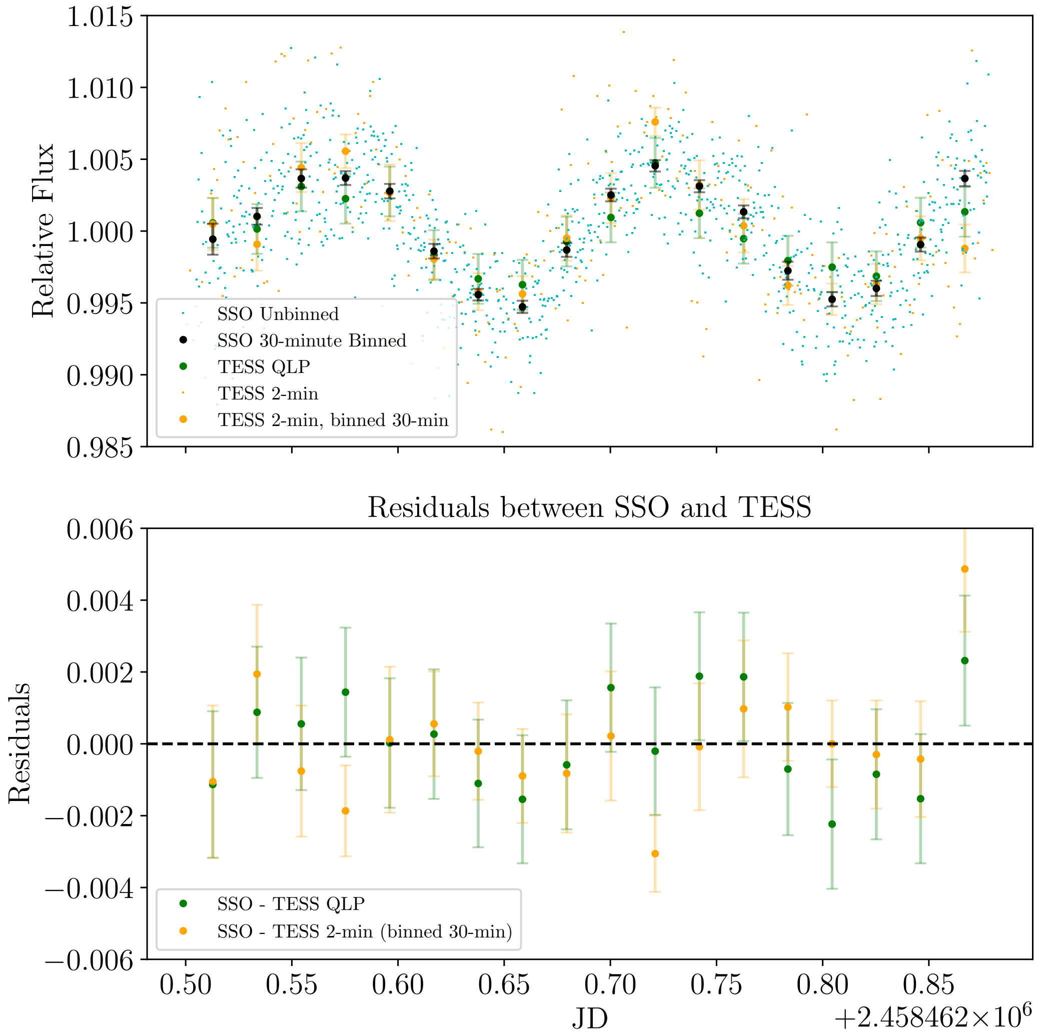

Figure 3.10: **Top**: SSO's differential lightcurve (unbinned points in cyan, binned points in black) compared to the lightcurves from TESS 2-minute cadence data (orange) and MIT QLP 30-minute cadence data (green) for an M6V object (J = 10.3 mag) on 10th December 2018. **Bottom**: The residuals between the TESS and SSO lightcurves.

3.4 Conclusions

An automatic pipeline allows for consistent monitoring of SPECULOOS data products and produces quality checks that can be used to probe for issues with the telescopes. I discuss a number of these quality checks in this chapter, such as dark current, readout noise and photometric transformations.

I define a subset of SSO data to analyse the photometric performance of the SSO's first few years of operation. By implementing careful data reduction and photometric treatment, I show that the SSO can regularly reach sub-millimag precision photometry, for observations of bright, quiet targets. By comparing simultaneous observations taken by SPECULOOS and the space telescope TESS, I demonstrate the exceptional performance of the SSO for ultra-cool objects, and the power of the synergy between the two surveys.

As mentioned in Section 3.2.2, stellar variability can seriously limit our planet detection efficiency. A more tolerant signal-to-noise ratio of 5 would increase these minimum precisions by a factor of 1.8. However, even for a rigorous SNR limit of 9, I show that the SSO is capable of producing lightcurves with precisions high enough to detect single transits from Earth-sized exoplanets orbiting bright, quiet targets, when the observing conditions are favourable. Therefore, this work demonstrates the excellent quality and detection capability of the SPECULOOS-South Observatory. The next chapter will focus on accounting for stellar variability to optimise our photometric precision and to aid the detection of small, rocky exoplanets. By optimising its detection efficiency, SPECULOOS provides a unique opportunity to explore the planetary population around UCDs, matching space-level photometric precisions with an ability to study fainter and redder objects than ever before.

CHAPTER 4

Stellar Variability in the Untapped Ultra-Cool Dwarf Regime

The dynamic relationship between planet and star is of prime importance in evaluating how conducive a planet's environment is for life. Cool stars are known to have heightened levels of activity. However, how this activity extends into the ultra-cool regime and the mechanisms behind it are poorly understood. More clarity into this underlying magnetic activity will allow us to better mitigate the significant problems that photometric variability poses for the detection of small, transiting planets, and to better understand the impact on their habitability.

Upcoming missions aiming to characterise planetary atmospheres in the next few years, such as JWST, will partially focus on planets hosted by red dwarfs due to the many advantages described in Section 1.3.1. However, the activity of host stars will limit the precision of JWST's transmission and emission spectra. This is especially concerning for cool stars where stellar spots can overwhelm or mimic planetary spectral features (Rackham et al. 2018, 2019). Therefore, there is a pressing need to perform long-term photometric monitoring of these objects, not only because stellar variability complicates atmospheric retrievals, but also to ascertain the detrimental and beneficial consequences for life.

This chapter presents a study of flares and rotation of SSO targets observed over almost two years. The work I present in this chapter is based on the work in Murray et al. (2022), where I perform a study of the flares and rotation on a subset of global lightcurves from the SPECULOOS-South observatory. In Section 4.2.1 I combine an automated flare detection algorithm with manual vetting to obtain the final flare sample. I then describe modelling of the flares (Section 4.2.2), calculating the flare energies (Section 4.2.3), and estimating flare rates (Section 4.2.4). I also assess the flare sample's completeness in Section 4.2.5. In Section 4.3 I measure the rotation periods in the SSO sample. The results of the flare and rotation analyses are presented in Section 4.4 along with a comparison to similar flare and rotation studies. The contextualisation of the impact of flares on the potential for life on planets around ultra-cool dwarfs is discussed in Section 4.5.

4.1 The Flares and Rotation of Ultra-Cool Dwarfs

Stellar flares may be a major determinant in whether a planet can initiate and sustain life. While flares can be destructive, through atmospheric erosion (Lammer et al. 2007), ozone depletion (Segura et al. 2010; Tilley et al. 2019), and even extinction events, they can also be an essential power source for life. It is possible that flares could provide the missing energy at the bluer end of the spectrum needed for cool, red dwarfs to initiate prebiotic chemistry (Buccino et al. 2007; Ranjan et al. 2017; Rimmer et al. 2018) and for photosynthesis (Mullan & Bais 2018; Lingam & Loeb 2019a). The additional UV energy may also affect the evolution of a planet's atmospheric chemistry (Segura et al. 2010; Vida et al. 2017). In summary, stellar flares have extreme and far-reaching consequences on the planets hosted by cool stars. Therefore, it is essential to study these stars' flaring activity, and assess how applicable our current understanding of stellar activity is to the ultra-cool regime.

Stellar activity and rotation are also closely connected. Magnetised stellar winds dictate the loss of angular momentum, which slows a star's rotation over time. These stellar winds depend on the structure and properties of the magnetic field. As the rotation slows, this decreases the magnetic activity in a process known as spin down (Skumanich 1972; Noyes et al. 1984). Due to this effect, rapidly rotating stars flare much more frequently than slow rotators (Skumanich 1986; Davenport et al. 2019; Mondrik et al. 2018; Medina et al. 2020). Spots and faculae on the surface of stars and brown dwarfs cause periodic photometric variations as they come in and out of view, on the same timescale as the object's rotation. This allows the rotation period to be deduced directly from the photometry. Therefore, the rotation of a star can provide valuable insights into the magnetic dynamo, the mechanism which generates a star's magnetic field.

This magnetic dynamo is poorly constrained for fully convective low-mass objects (with masses $\leq 0.35\mathrm{M}_{\odot}$, Chabrier & Baraffe 1997), where it is believed to differ significantly from the solar model. Despite this predicted difference, recent work has shown that relationships between activity and rotation remain consistent from partially to fully convective stars (Wright & Drake 2016; Newton et al. 2017; Wright et al. 2018). Spin down is, however, believed to occur on slower timescales for fully convective stars, which accounts for the enhanced activity of mid-to-late M-dwarfs (West et al. 2008; Newton et al. 2017; Jackman et al. 2021). Therefore, rotation, activity, and the relationship between the two, are extremely useful probes of the underlying magnetic behaviour of ultra-cool dwarfs.

Due to the promising nature of M-dwarfs as planetary hosts, there have been several detailed studies of their flaring activity within the past decade. Space telescopes, such as the Kepler/K2 missions, allowed the first insights into the flares of bright M-dwarfs (Davenport et al. 2014; Hawley et al. 2014; Lurie et al. 2015; Silverberg et al. 2016) and of (small numbers of) ultra-cool dwarfs (Paudel et al. 2018; Gizis et al. 2013, 2017). Additionally, the recently launched TESS satellite, with its increased precisions at redder wavelengths, has facilitated studies of the flaring and rotating activity of cool stars (Günther et al. 2020b; Medina et al. 2020; Seli et al. 2021). Several ground-based photometric surveys, such as MEarth, NGTS, the All-Sky Automated Survey for Supernovae (ASAS-SN, Shappee et al. (2014)) and Evryscope (Law et al. 2015) have also carried out detailed flare studies that include M-dwarfs (West et al. 2015; Mondrik et al. 2018; Schmidt et al. 2019; Howard et al. 2019; Martínez et al. 2019; Howard et al. 2020a,b; Jackman et al. 2021).

The most frequent stellar flares are small, fast and challenging to detect above photometric scatter (Lacy et al. 1976). Due to the intrinsic faintness of UCDs at visible wavelengths (~400–700 nm), it is, therefore, challenging to achieve the high photometric precisions necessary to constrain their flaring activity. However, it is possible to perform small, dedicated studies of the much rarer, high-energy flares on UCDs (Gizis et al. 2013; Paudel et al. 2018; Jackman et al. 2019). Therefore, to obtain a sufficient sample size, previous large flare studies have focused on hotter stars, up to mid M-dwarfs, where photometric precisions are much higher. This has resulted in limited flare statistics for the coolest stars, for which we would require a large, high-cadence photometric survey optimised for UCDs, such as SPECULOOS.

For the analysis of the flaring and rotation activity of ultra-cool dwarfs, I used the data sample defined in the previous chapter as my basis (see Section 3.1.1). In this sample, 80% of the objects were observed for less than 77 hours (Figure 3.2), and 50% were observed for less than 51 hours. The short observation times in this sample will inevitably bias me towards detecting flares on targets with high flaring rates in Section 4.2.1, and detecting the clearest rotation periods for fast rotators in Section 4.3.

4.2 A Study of Flares

4.2.1 Flare Detection

As the SSO data sample is constrained to only 176 targets (with 154 observed for more than 20 hours), it is possible to identify flares manually. Therefore, I did not find it necessary to implement a fully automated, complex flare-detection algorithm. Instead, I performed flare detection in two parts: a simple, automatic flare-detection algorithm to extract all the flare candidates, followed by a manual vetting process to confirm them. Both stages of this flare detection process are demonstrated in Figure 4.1 for target LP 655-48 (Gaia DR2 3200303384927512960). I note that this target is one of the most frequently flaring objects detected by the Seli et al. (2021) flare study. This object is a particularly challenging case due to its rapid variability, where flare decays are difficult to separate from photometric modulations.

Automated Flare Search using Lightcurve Gradients

For the first part of the flare detection process, I implemented the automatic flare detection method set out in Lienhard et al. (2020) due to its simplicity, speed, and robustness. This method was developed on lightcurves from the UCDTS survey carried out with the TRAPPIST telescopes. Lienhard et al. (2020) capitalised on the changing gradients in the asymmetric structure of a flare, which involves a sharp increase to a peak followed by a slower exponential decay (Moffett 1974). They evaluated the following two criteria:

$$\frac{2f_j - f_{j-2} - f_{j+3}}{\sigma_j} . \frac{|f_j - f_{j-2}| - |f_{j+1} - f_j|}{\sigma_j} > A_{\mathrm{thresh}} \tag{4.1}$$

and

$$2f_j - f_{j-2} - f_{j+3} > 0, \tag{4.2}$$

where f_j is the flux of the j^{th} (unbinned) data point in the target's lightcurve, and σ_j is the RMS of the section of the lightcurve comprised of the nearest 60 data points. Equation 4.1 assesses the quality of the flare's shape. If I assume that the peak of the flare is at j then the first half of Equation 4.1 confirms that we are at a peak by ensuring the flux at j is greater than the flux a few exposures before and after it. The steeper the peak or larger the flux difference between the peak and surrounding points, the higher its value. The second half of Equation 4.1 hinges on the asymmetry of the flare by requiring that the gradient before the peak, the fast rise, is larger than the gradient after the peak, the slow decay. For this second half of the equation, the greater the asymmetry, the higher its value. By dividing by the local RMS, the aim is to remove flare structures caused by photometric scatter, which is especially problematic for ground-based observing with rapidly changing weather conditions. Equation 4.2 simply prevents the case where both halves of Equation 4.1 are negative.

By considering only small neighbouring sections of the lightcurve in the RMS, and the flux for a few data points on either side of the flare in my criteria, I can detect small flares on rapidly rotating or frequently flaring objects. By comparing the flux differences between consecutive data points we make a small number of assumptions on the timescales of our flares. It is optimal

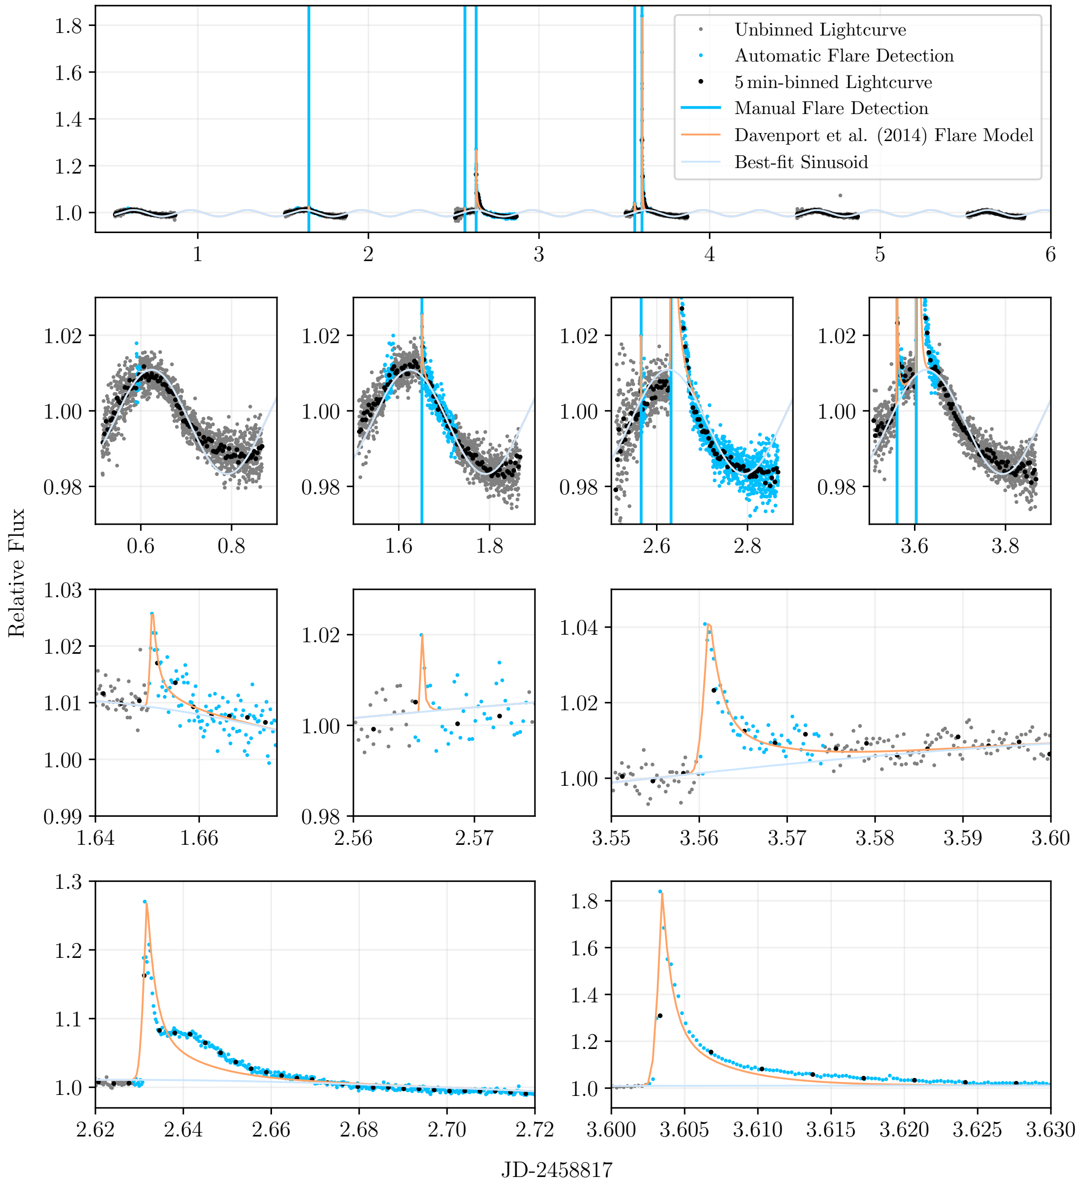

Figure 4.1: Demonstration of the flare detection process on six nights of observations of LP 655-48, an M7 object (J = 10.7 mag), which exhibits both frequent flaring and short-period rotation. The four rows show different views of the same global lightcurve, zoomed in to four individual nights with automatically flagged flares on the second row. Only the five manually-verified flares are plotted on the bottom two rows. The grey points are the unbinned data, while the black points are binned every 5 minutes. I performed a simple least-squares fit of a sine wave with a period of 0.33d (derived in Section 4.3), shown in light blue. It is clear from this fit that the rotation pattern is not perfectly sinusoidal; however, it provides a way to visualise the periodicity. The blue data points are those initial flare regions flagged by the automatic flare detection algorithm described in Section 4.2.1. The vertical, blue lines are the flare candidates that were also confirmed manually. During visual inspection of this lightcurve, I removed two flares due to their small amplitude, which were found too difficult to detach from the photometric scatter. The best-fit Davenport et al. (2014) models for each flare are shown in orange. The bottom left plot demonstrates the case where there is a poor fit to the template, likely due to a combination of overlapping flares.

to avoid making assumptions about flare decay times due to the absence of previous large-scale flare studies on UCDs, and so our understanding of how UCD flares compare to those on other types of stars is still limited. However, in this work, I assume that we have a distinct flare peak and that the flux decay happens at some rate slower than the flux rise. Secondly, I acknowledge that there are time-dependent limitations to this automated flare detection method, meaning that I will systematically detect certain types of flares and miss others. If the flare timescale is so short that the flux increase and decrease are within one or two exposures (a larger risk for longer exposures), we will not detect that flare. Therefore, if the exposure time is comparable to the flux rise time, or longer, I would not expect to detect it. If the flare timescale is very long compared to the exposure time then the differences in flux between adjacent points (in Equation 4.1) will be small. This could mean that the threshold defined in Equation 4.1 is not met, though the dependency of σ_j on exposure time is complex (and a full exploration of this dependency is beyond the scope of this work). Therefore, while I have tried to avoid "hard-coding" a specific temporal relationship for flares into my detection method, the relationship between flare timescale and exposure time will impact the flares that are detected.

In their paper Lienhard et al. (2020) determined $A_{\mathrm{thresh}} = 12$ by inspection of the smallest flares that they intended to remove; however, this value is specified for the photometric precision and typical exposure times of TRAPPIST lightcurves. I would expect the precisions obtained by TRAPPIST's 60-cm telescopes to differ significantly from those achieved by SPECULOOS's 1-m telescopes. Instead, for the SSO lightcurves, I derived a value of $A_{\mathrm{thresh}} = 5$. I found a lower value for A_{thresh} partly due to the increased photometric precision of the SSO and partly because I am choosing to verify the flare candidates that this algorithm detects manually, so I could afford to over-detect at this stage.

Choosing a flare detection method based on the shape of a flare, not outlier detection, may limit the diversity of flare morphologies that I will detect. Since there have been few large-scale studies on ultra-cool dwarf flares, such as those in the SSO sample, any difference in flare structure between earlier and later spectral type objects is still unknown. Therefore, I focused on the flares most resembling the standard flare shape (Moffett 1974) while flagging more unusual lightcurve behaviour in the manual vetting stage (see below).

Small and Large Flares

Both to maximise the detection of small flares and optimise the modelling of high amplitude flares with slow recovery times, I split the flare sample into two categories: *small* and *large* flares. I classified *large* flares as those with at least 2 data points in the flare region more than seven times the running local standard deviation (standard deviation of the surrounding 80 data points) above the running local median (defined similarly as the median of the surrounding 80 data points); otherwise, I classified it a *small* flare. The flare region starts at least two points before the peak. I approximated the end of the flare region by using least-squares to fit the flare decline with the sum of a fast and a slow-decaying exponential, as in Davenport et al. (2014). The decay time of the slower exponential was used to estimate the end of the flare region.

Once a flare was classified as *large* then I used the surrounding 320 points (defined as global), instead of 80 points (local), to calculate the running median and standard deviation.

Validating and Vetting Flares

If the above criteria were met, I ran several additional automatic quality checks to separate flare candidates from flare-like signals replicated by noise or cosmic ray hits. Cosmic rays were flagged after flare detection to avoid removing the peaks of small flares, with only one or two points of significantly increased flux. I removed cosmic rays from the flare candidates by ensuring that there was more than one data point in the flare region after the peak. Additionally, at least 2 points in the flare region must be more than 2 times the running local standard deviation above the running local median of the lightcurve. The local standard deviation and median are as defined above. When calculating the running local median and standard deviation I masked all potential flare candidates in the lightcurve and then interpolated over the flare region. This automated algorithm provides a collection of flare candidates to follow up with manual vetting, to then obtain the final flare sample.

Manual Vetting of Flare Candidates

Once I had a collection of flare candidates, the second part of my flare detection process was to manually inspect these candidates to obtain the final flare sample. I considered each flare candidate in turn and only confirmed those which could be clearly identified as matching the standard flare shape of a sharp flux increase followed by a slower, exponential flux decay (Moffett 1974; Davenport et al. 2014). I also ensured that the flares could not be attributed to rapid changes in atmospheric conditions. As a final validation step, I checked through the global lightcurves in the SSO sample. The only flare-like structures that remained were due to the following events: a flare occurs before the start of the night (where I catch the tail end of the flare) or at the very end of a night, the structures are low amplitude so cannot be distinguished from noise, the structure diverges significantly from the standard flare model, or they are correlated with systematics. Examples of structures identified in the SSO lightcurves that deviated from the standard flare model, and were not clearly overlapping flares, were symmetric flares (with similar rise and decay times) and flares with a linear (non-exponential) decay. During this manual vetting stage I did not add any of these potentially missed flares into the flare sample, as they did not meet the standard flare shape criteria.

Removing Cosmic Rays

When a cosmic ray hits a pixel on the CCD, it causes an instantaneous spike of electron counts in that pixel, which may be mistaken for a short-duration flare. To detect cosmic rays, I flagged points more than three times the running local standard deviation above the running local median of the lightcurve (on which flares had already been masked). Misinterpreting data points in flares as cosmics can result in underestimating flare energies, or even missing flares entirely.

4.2.2 Modelling Flares

To model each flare, I coarsely flattened the lightcurve surrounding that flare by dividing the lightcurve by a median filter, using the local or global criteria as defined above.

I chose to model the flares by fitting the empirical flare template described in Davenport et al. (2014). This paper generates a median flare template from Kepler observations of 885 'classical' flares on the M4 star GJ 1243. They found a sharp flux increase, modelled with a fourth-order polynomial, followed by an initial fast exponential decay and subsequent slower exponential decay. This model requires a time for the peak of the flare, a relative amplitude of the flare peak on the normalised lightcurve, and a full width half maximum ($FWHM_{flare}$), which corresponds to the flare's decay time. I demonstrate fitting this flare model to flares of differing amplitudes and decay times in Figure 4.1.

While this model works well for the more classically shaped flares, it struggles to represent complex flares that do not fit the standard flare morphology (Davenport et al. 2014). This template has yet to be tested on a statistically significant number of low-mass objects, such as UCDs, whose flare profiles may differ significantly. Complex flares include those with multiple peaks (Davenport et al. 2014), oscillations (Anfinogentov et al. 2013), and those which are closely entangled with variability. The overlapping of multiple flares is likely the culprit of more unusual lightcurve structure. During manual vetting, I separated flare regions containing clear multiple peaks, and flagged 27 flares that either did not fit the standard flare model or were not easily separable into multiple flares. This corresponds to 11% of the total flare sample. An example of a possible flare overlap is shown in the bottom left subplot of Figure 4.1.

4.2.3 Extracting Flare Energies

To measure the flare energies, I followed the technique described in Shibayama et al. (2013). I modelled the flare as a blackbody with an effective temperature of 9000±500 K (Kowalski et al.

2013), and I assumed this temperature remains constant. I calculated the luminosity of each star and flare as seen through the SSO's I+z' filter ($L_{\lambda,*}$ and $L_{\lambda,\mathrm{flare}}$ respectively), accounting for the overall system efficiency:

$$L_{\lambda,*} = \pi R_*^2 \int_{I+z'} R_\lambda B_\lambda(T_\mathrm{eff})\,\mathrm{d}\lambda, \tag{4.3}$$

$$L_{\lambda,\mathrm{flare}} = A_\mathrm{flare}(t) \int_{I+z'} R_\lambda B_\lambda(T_\mathrm{flare})\,\mathrm{d}\lambda. \tag{4.4}$$

Here, R_λ is the total SSO response function, which is the product of the transmission in the $I + z'$ filter, the CCD window, the lenses, the quantum efficiency of the CCD, and reflectivity of the mirror coatings. $B_\lambda(T_\mathrm{eff})$ and $B_\lambda(T_\mathrm{flare})$ are the Planck functions evaluated for the star's effective temperature and the flare temperature respectively. Finally, $A_\mathrm{flare}(t)$ is the area of the flare. Since the relative flare amplitude can be found directly from the normalised lightcurve, $F(t) = (\Delta F/F_\mathrm{mean})(t) = L_{\lambda,\mathrm{flare}}(t)/L_{\lambda,*}$, I solved for A_flare as follows:

$$A_\mathrm{flare}(t) = F(t)\,\pi R_*^2 \,\frac{\int_{I+z'} R_\lambda B_\lambda(T_\mathrm{eff})\,\mathrm{d}\lambda}{\int_{I+z'} R_\lambda B_\lambda(T_\mathrm{flare})\,\mathrm{d}\lambda}. \tag{4.5}$$

I note that $L_{\lambda,*}$ and $L_{\lambda,\mathrm{flare}}$ are the luminosities in the observed wavelength band and not the bolometric luminosities. To best estimate the flare amplitude, I used the local or global running median for F_mean. The bolometric flare luminosity, L_flare, was then calculated with the assumption that the star radiates as a blackbody:

$$L_\mathrm{flare} = \sigma_\mathrm{SB} T_\mathrm{flare}^4 A_\mathrm{flare}. \tag{4.6}$$

The total bolometric energy of a flare (E_flare) is then the integral of L_flare over the flare duration:

$$E_\mathrm{flare} = \int_\mathrm{flare} L_\mathrm{flare}\,\mathrm{d}t. \tag{4.7}$$

Using Equation 4.6, and assuming the flare temperature is constant, I could then separate out the time-dependent and independent components:

$$E_\mathrm{flare} = C \int_\mathrm{flare} A_\mathrm{flare}\,\mathrm{d}t, \tag{4.8}$$

where $C = \sigma_\mathrm{SB} T_\mathrm{flare}^4$. Using Equation 4.5 I further separated out the time-independent components that depend on the properties of the star and the flare, and the time-dependent integral:

$$E_\mathrm{flare} = C_\mathrm{star}(T_\mathrm{eff}, R_*) C_\mathrm{flare}(T_\mathrm{flare}) \int_\mathrm{flare} F(t)\,\mathrm{d}t. \tag{4.9}$$

Calculating the energy in this way allowed me to calculate these three components (C_star, C_flare and the integral) separately. When generating a large pool of synthetic flares for injection-recovery tests, as in the following section, this drastically reduces the processing time. Every object in the SSO data sample will have a constant value for C_{star}, which only needs to be calculated once per star. The value for C_{flare} also becomes constant with these assumptions. It is then only the integral of the flare in the normalised lightcurve that remains to be calculated.

For this integral, I decided to use the normalised lightcurve (corrected with the same local or global median filter described in Section 4.2.1) directly for the energy calculation rather than using the flare model. In doing so, I note that the energies I calculated provide a lower bound on the flare energy, as I may have underestimated the flare's peak. I chose not to use the flare model because the flares in the lightcurves are well-sampled, and when there is increased photometric noise in the lightcurve, or more complex flare shapes from overlapping flares (see Section 4.2.2), the fit to the flare model can be unreliable. Integrating the lightcurve directly means that bursts of flares occurring in short succession are counted as one flare of increased energy. While this simplification may have affected the calculation of the flare rates and energies, it should not have significantly affected the results, as I only detected 27 unusually shaped flares that I could not easily separate into multiple flares. While flares are typically approximated as blackbodies with temperatures from 9000–10000 K, flare temperatures have been observed to vary outside this range, both between and within flare events (Howard et al. 2020b). A 5% error of ±500 K (Kowalski et al. 2013) in flare temperature would result in a 20% error in energy (Equation 4.6). However, as we're working with extremely large energies of order $\sim 10^{30}$ ergs in log space, a 20% difference would not significantly change our results (errors in $\log E_{\mathrm{flare}}$ of ~0.3%). Likewise, the typical fractional error in the lightcurves, or $F(t)$, of 0.1% would produce negligible final errors in $\log E_{\mathrm{flare}}$.

4.2.4 Calculating Flare Rates

As flares are stochastic events, the gaps in observations due to the day-night cycle or weather loss do not hinder our ability to calculate flaring rates, which only depend on the total time on-sky. I likely underestimate flaring rates due to the missed flares outlined in Section 4.2.1.

I calculated the flare rates as the number of flares detected per target divided by the observation time (summed across all telescopes). Almost all the targets were observed for less than 200 hours. Therefore, I have a lower limit on individual stars' flaring rates of ~ 0.12 flares d^{-1}.

4.2.5 The Completeness of the Flare Sample

To evaluate the completeness of the SSO flare sample (the minimum energy flares recoverable by the flare detection process for different spectral types), I performed artificial flare injection-recovery tests. Using the global lightcurves from the 154 targets with at least 20 hours of observation, I masked all the flares detected in Section 4.2.1. I used the Davenport et al. (2014) flare model to generate 100,000 artificial flares with amplitudes drawn from a log-normal distribution between relative fluxes of 0.001 and 5, and values for the $\mathrm{FWHM_{flare}}$ drawn from a uniform distribution between 30 seconds (typical exposure time for SPECULOOS) and 1 hour. These distributions were determined by examining real flares in the SSO lightcurves. This amplitude distribution was chosen as amplitudes of 0.001 are comparable with the minimum noise level in SPECULOOS's lightcurves and all flare amplitudes of and above 5 should be recovered easily, therefore there was no advantage to extending the range above this. The $\mathrm{FWHM_{flare}}$ of a flare describes the time it takes for a flare to decay, for which I determined the same range as Davenport (2016). The minimum value of 30 seconds is a typical exposure time of SPECULOOS observations (we cannot recover flares that occur almost entirely between consecutive exposures). The maximum value was chosen to be 1 hour as any flare with a longer $\mathrm{FWHM_{flare}}$ would be too long to recover with SPECULOOS's typical observation durations (4–8 hours). I note that with these distributions I make no assumptions about the underlying flare distribution, frequencies or any relationship between amplitude and $\mathrm{FWHM_{flare}}$ (which has not been well studied for UCDs). The purpose of these injection-recovery tests are to explore the flare recovery within the full parameter space, not to replicate real flaring activity (which will exist within this parameter space). I calculated the energy of the resulting flares for each star, and divided into 6 energy bins in $\log_{10}$ space from 10^{28} to 10^{34} ergs.

I then randomly selected 5 flares from each energy bin and injected them into the SSO sample lightcurves at a random time during observation, ensuring the +/- 5 data points surrounding the time of the flare's peak are within 0.01d (14.4 mins). This prevented flares from occurring too close to the start or end of the night or during gaps in the observations, which would always be missed. I allowed flares to overlap to reflect the scenario seen often in real lightcurves.

I then ran the automatic part of the flare detection method, described in Section 4.2.1, on all injected lightcurves and recorded the recovered flares. I considered a flare to be recovered if the recovered time for its peak flux was within 1 $\mathrm{FWHM_{flare}}$ (of that flare) of the injection time. I allowed for a more flexible recovery time because of the decision to allow flares to overlap, which can lead to more complex structures with multiple peaks. If two or more flares overlap within 1 $\mathrm{FWHM_{flare}}$ they may be detected as only one flare. I assumed if they occur further apart in time then they should be easily separable. This process was then repeated for each energy bin 5 times, making sure to use different flares from my artificial sample in each iteration.

The manual vetting step of the flare detection technique was not feasible for checking the approximately 23,000 flares injected. This may have introduced a bias in my results. During manual vetting, it is more likely that I would conflate low-amplitude flares or flares on high

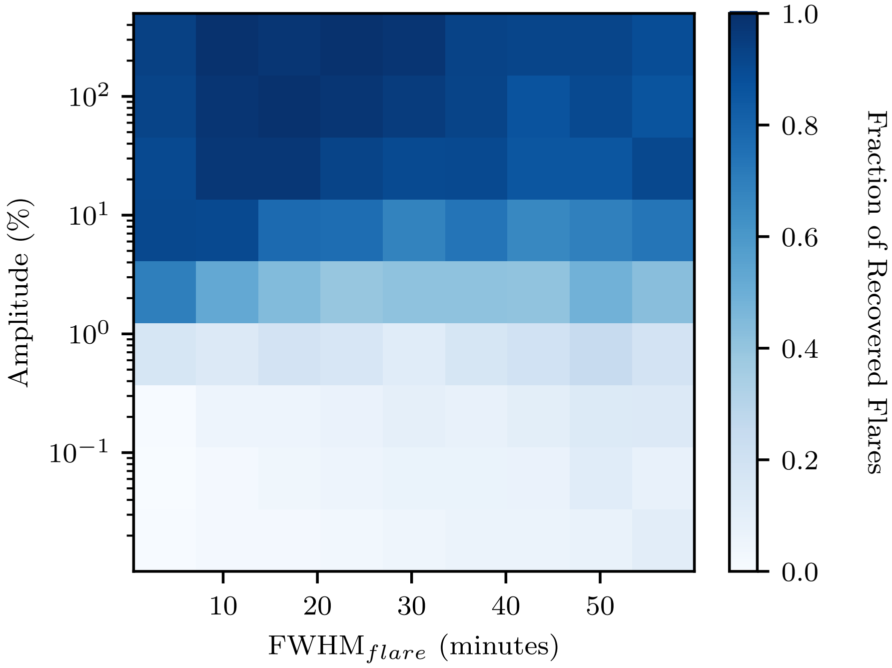

Figure 4.2: Fraction of injected flares recovered as a function of flare amplitude and $\mathrm{FWHM}_{\mathrm{flare}}$ using the automatic flare detection algorithm described in Section 4.2.1.

photometric scatter targets with noise. Therefore, by removing this stage, I have potentially inflated recovery rates for the lowest energy flares, or for the faintest target stars.

Figure 4.2 shows the results of the injection-recovery tests carried out for approximately 23,000 artificial flares. The flare detection algorithm detected a significant proportion of flares with amplitudes above 1%. Due to the short exposure times of SPECULOOS, we have the advantage of being able to detect flares with a short duration (with $\mathrm{FWHM}_{\mathrm{flare}} < 5$ minutes). However, my detection method was more limited for longer duration flares (with $\mathrm{FWHM}_{\mathrm{flare}} \geq$ 60 minutes) due to the day-night cycle, with typical uninterrupted observation windows of 4–8 hours. There will also be some stars which have nights of observation less than 4 hours, due to weather loss or limited visibility. These target's lightcurves will be the most difficult to detect flares on. However, during the manual vetting procedure (Section 4.2.1) I inspected lightcurves with flares removed and did not find any remaining high energy flares in the sample (which were not covered by the previously described exceptions). I also assessed the correlation with the photometric RMS of the lightcurves. When the RMS of a global lightcurve exceeded $\sim 0.7\%$ (for 5-minute binning) I began to see a drop in detection efficiency, resulting in an increase in the minimum detectable amplitude. However, only three targets in the dataset exceeded this RMS.

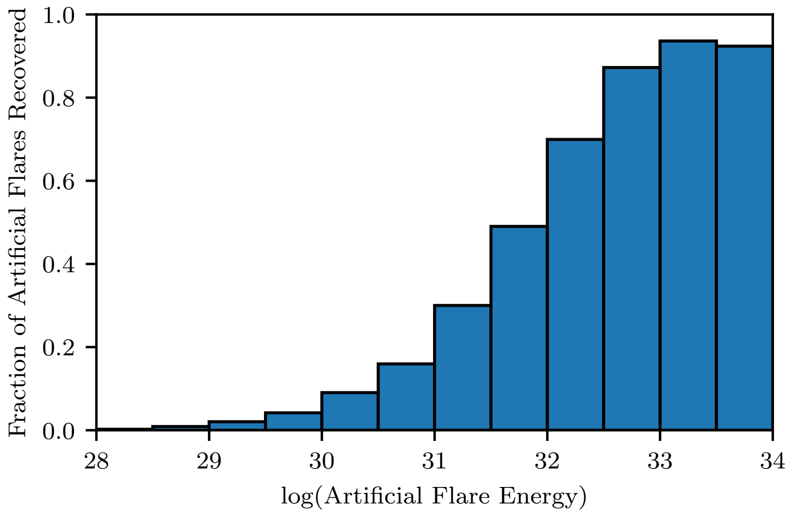

Figure 4.3: Fraction of injected flares recovered as a function of flare energy.

For a given flare amplitude and $FWHM_{flare}$, the flare energy varies depending on the effective temperature and radius of the star (Equation 4.9). The recovery fraction for different flare energies across all stars is shown in Figure 4.3. However, the minimum detectable energy of recovered flares varies with spectral type, as shown in Figure 4.4, which facilitates the detection of lower energy flares on the coolest dwarfs. For low-energy flares, there are two competing biases. The lightcurves for the later, and fainter, M and L dwarfs have higher photometric scatter, which makes it difficult to detect small-amplitude (and lower energy) flares. However, as flares have a strong white light component, there is an increased contrast between flares and the stellar spectra of red dwarfs, which implies lower energy flares should be easier to identify in the coolest stars (Allred et al. 2015; Schmidt et al. 2019). SPECULOOS is also optimised to observe ultra-cool objects. Due to this specificity, I only see a very minimal increase in RMS with later spectral types.

The results from the artificial flare injection-recovery tests (in Section 4.2.5) allowed me to calculate the recovery fraction of the flare detection method for every star, $R(E)$, which can reflect the low recovery rates of low energy flares. The average $R(E)$ for all stars is shown in Figure 4.3, as well as its spectral dependence in Figure 4.4. As the SSO flare sample is already small I was unable to reduce detection-sensitivity effects by simply limiting the sample to flares with energies above a minimum recoverable energy (Paudel et al. 2018), or energies where a sufficient fraction of flares were recovered (Davenport 2016). Instead, I decided to calculate flare occurrence rates in Section 4.5.6 using the method described in Jackman et al. (2021), which models the decline in sensitivity for low energy flares.

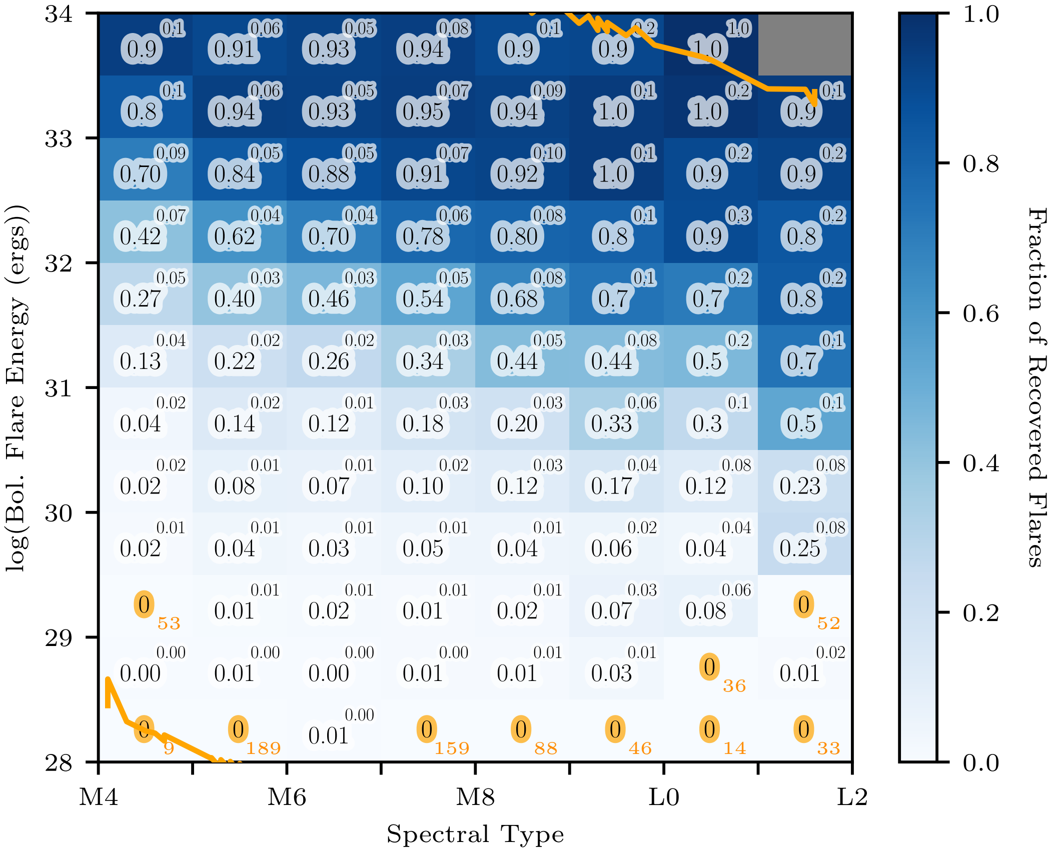

Figure 4.4: Fraction of injected flares recovered as a function of flare energy and the object's spectral type. In the upper right corner of each box, I include the error on the recovery fraction (in that box). If this error ≥ 0.1, then the recovery fraction and error are presented to one significant figure, otherwise two are shown. In the lightcurves of cooler stars, I can detect lower energy flares. The 10 boxes for which no flares were recovered are highlighted in orange. The number of injection-recovery trials are shown in orange in the lower right corner of each of these 10 boxes. The upper and lower flare energy limits, determined by the amplitude and $FWHM_{flare}$ limits of the artificial flare sample and the radius and T_{eff} of each star, are shown in orange. The sample limits only impact the results for spectral types earlier than M5 and later than L0.

4.3 Identifying Stellar Rotation Periods

I searched for rotation periods in the 154 low mass objects with more than 20 hours of observation, with the previously identified flares masked. I applied the Lomb-Scargle periodogram analysis (Lomb 1976; Scargle 1982) to the global lightcurves, binned every 20 minutes. The SSO lightcurves are not uniformly sampled as there are gaps, not only from the day/night cycle, but also from bad weather, masked flares, and changes to SPECULOOS's observation strategy. Therefore, careful treatment of the resulting periodograms is essential to remove aliases.

I searched for periods between the Nyquist limit (twice the bin size) and the entire observation window. While I could not be fully confident in periods greater than half their observation window, I estimated a lower bound on long period rotation by expanding the period range I searched to cover the entire observation window. I then removed all peaks in the periodogram with a false alarm probability below 3σ (0.0027). For any target which had at least one significant Lomb-Scargle peak I visually inspected their periodogram and the corresponding phase-folded lightcurves for all possible rotation periods (i.e., all significant peaks in the periodogram). In addition, by comparing with periodograms of the time stamps, airmass, and FWHM I eliminated signals which arose from the non-uniform sampling, and from ground-based systematics.

I decided to apply a similar classification system as Newton et al. (2016) to the rotating objects in the SSO data sample. Classifying the rotators helps account for the difficulties in observing from the ground through atmospheric systematics and irregular, non-uniform sampling. When examining each lightcurve, I asked myself several questions:

1) Is the period clearly visible in the phase-folded lightcurve?
2) Was the object observed for long enough to span multiple periods?
3) Is the frequency an alias of the "day signal", seen as integer multiples in frequency space (periods of 1, 0.5, 0.33 d etc.)?
4) Is there a correlation with systematics (as seen in the airmass or FWHM periodograms)?
5) Can the period be seen by-eye in the un-phased lightcurve?
6) If this object is observed with more than one telescope, does this period fit them all?
7) Can I easily disentangle the 'real' period from its one-day aliases?
8) Is the amplitude of the periodic signal above the level of photometric scatter in the lightcurve?

If a rotator passed all the above criteria, I classed it as 'A'. If it failed any of the above, but the rotation still seemed likely, I assigned it as a 'B' grade rotator. Most commonly, the 'B' rotators were convincing, but SPECULOOS did not observe multiple cycles, or I could not easily choose between the period and its one-day aliases. Any lightcurves for which I saw

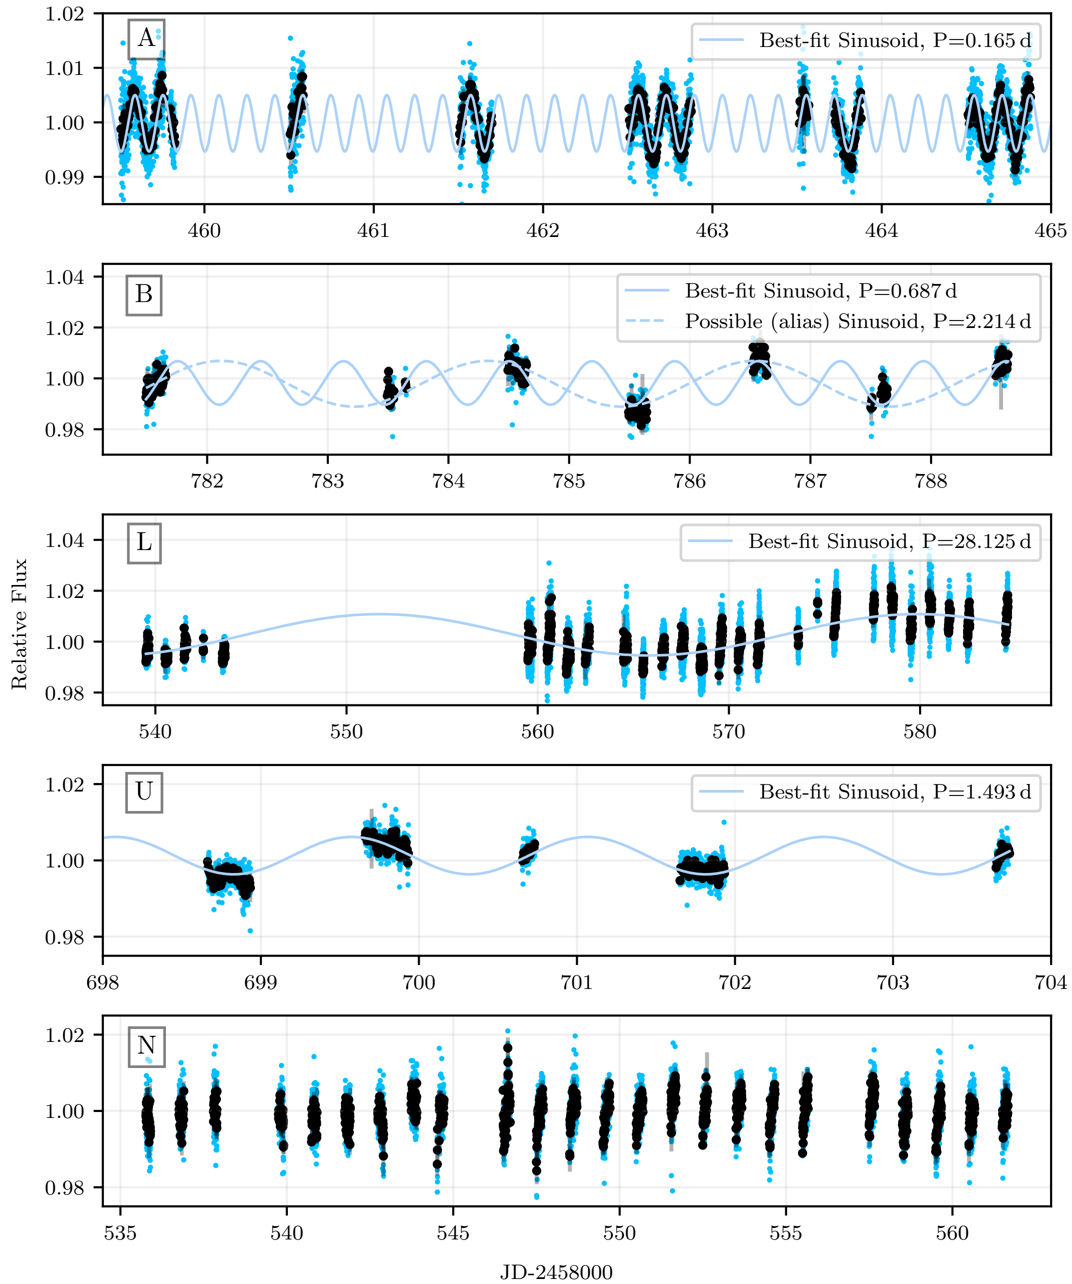

Figure 4.5: An example of each of the five rotation classes from top to bottom: A, B, L, U and N. For the A, B, L and U plots, I include the “best” period. For the rapidly-rotating A rotator, the period of 0.165 d was evident. However, I was uncertain about choosing the true period (amongst its aliases) for the target assigned as a B rotator. I show the P=2.214 d alias of the selected period (P=0.687 d) as well to demonstrate how both rotation periods could produce the observed data. For the L rotator the period of ~28 d is too long to confirm with our observations, as we observe less than two full periods. For the U rotator, there appears to be some small-amplitude rotation with P=1.493 d, however, we do not have enough data to confirm this and the first observation night has some unexplained structure that does not appear to match the sinusoid. Finally, the N rotator does not show any clear rotation pattern above the daily scatter of data points.

Table 4.1: The definitions of the five rotation classes.

Rotation Class	Definition
A	Clear rotation period, passes all 8 criteria.
B	Likely rotation period, passes most of the 8 criteria, but fails at least one.
L	Long period rotation clear but with a period comparable to, or longer than, the observation span.
U	Uncertain rotation period, fails most of the 8 criteria.
N	No detected rotation period.

some periodic structure, but could not easily determine a period, I classed as ‘U’. This may have resulted from broad peaks in the Lomb-Scargle periodogram, or multiple possible periods due to lack of observations, or a very low amplitude periodic signal that was comparable to the noise level. Because of these ambiguities in the period measurements, I did not attempt to place errors on the period estimates. If I did not detect any periodic signal or I could not remove correlations with systematics (such as for very crowded fields), then I classed the object as ‘N’. Finally, in addition to the classes defined in Newton et al. (2016), I added an extra ‘L’ grade. This was for the lightcurves where I saw clear long period rotation, but the period was similar or longer than the time window observed. For these objects the best that could be done was to estimate the lowest possible period measured and acknowledge that there are large uncertainties on these period values. I present an example of each rotation class in Figure 4.5. For reference the definitions of the rotation classes are shown in Table 4.1.

4.4 Combined Results from the SSO Flare and Rotation Study

The SSO data sample, along with the results of the flare and rotation analyses, is presented at the end of this thesis, in Table C.1 in Appendix C.

4.4.1 The SSO Flare Sample

From the SSO dataset described in Section 3.1.1, I identified 234 flares. Of the 154 unique targets, 78 are flaring (50%). These flaring stars span the spectral type range from M4 (T_{eff} = 3160K) to L0 (T_{eff} = 2313K). Figure 4.6 shows the spectral type distribution and proportions of the flaring objects identified in the SSO sample. From this figure, the proportion of flaring stars stays consistently above 60% for objects of spectral type M5–M7. The rate of flaring stars begins to decline around M8 ($\sim$ 30%), with no detected flares for any object beyond L0. The coolest flaring star I detected was a 2313 K, M9.6V object (which is rounded to L0 in Figure 4.6). However, for the L-dwarfs I am limited by the small sample size; therefore, I cannot

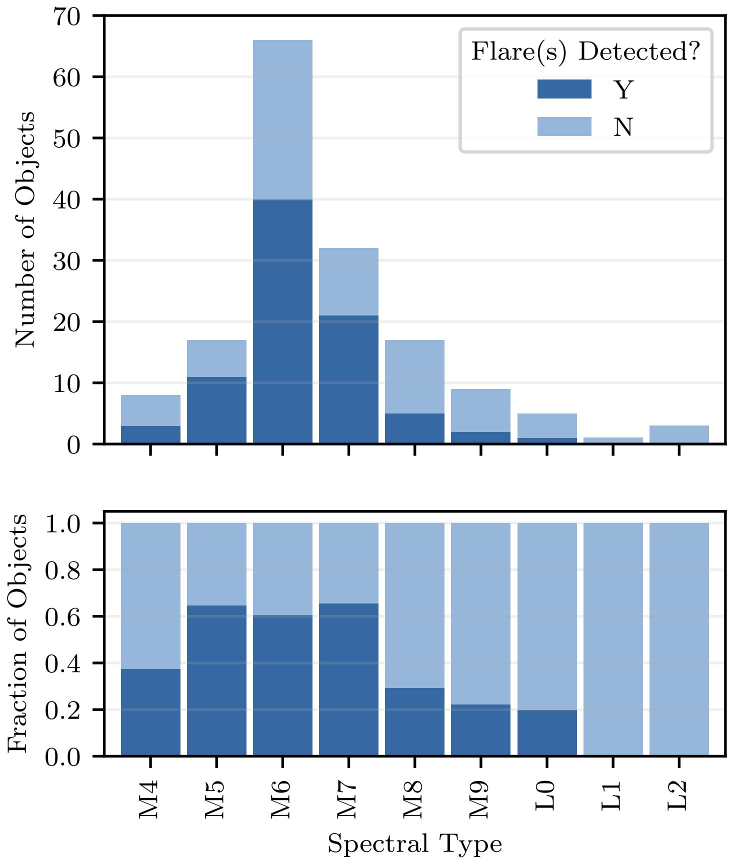

Figure 4.6: **Top:** Histogram of objects in the SSO data sample as a function of spectral type. The flaring objects are marked in dark blue. **Bottom:** Fraction of flaring objects as a function of spectral type.

conclude whether the fraction of flaring objects continues to decrease beyond late M-dwarfs. Likewise, there were no objects earlier than M4; however, these earlier M-dwarfs have been well studied, for example by TESS, Kepler, and MEarth. Despite the small sample of M4 dwarfs, I see an apparent reduction in activity for stars earlier than M5 that aligns with previous results (~30% for M4 with TESS from Günther et al. 2020b, 25% for M5 with ASAS-SN from Martínez et al. 2019). Several authors find a steep rise in flaring fractions around spectral type M4 (West et al. 2004; Yang et al. 2017; Günther et al. 2020b; Martínez et al. 2019). Specifically, West et al. (2004) analysed the Hα emission of cool stars (a different stellar activity indicator based on the strength of chromospheric emission). Hα emission roughly correlates with flaring activity (Yang et al. 2017; Martínez et al. 2019), and West et al. (2004) determined that the fraction of active stars rose monotonically from spectral type M0 to M8 (in agreement with Jeffers et al. 2018), peaked at M8 and declined steadily to L4. My results show an increase

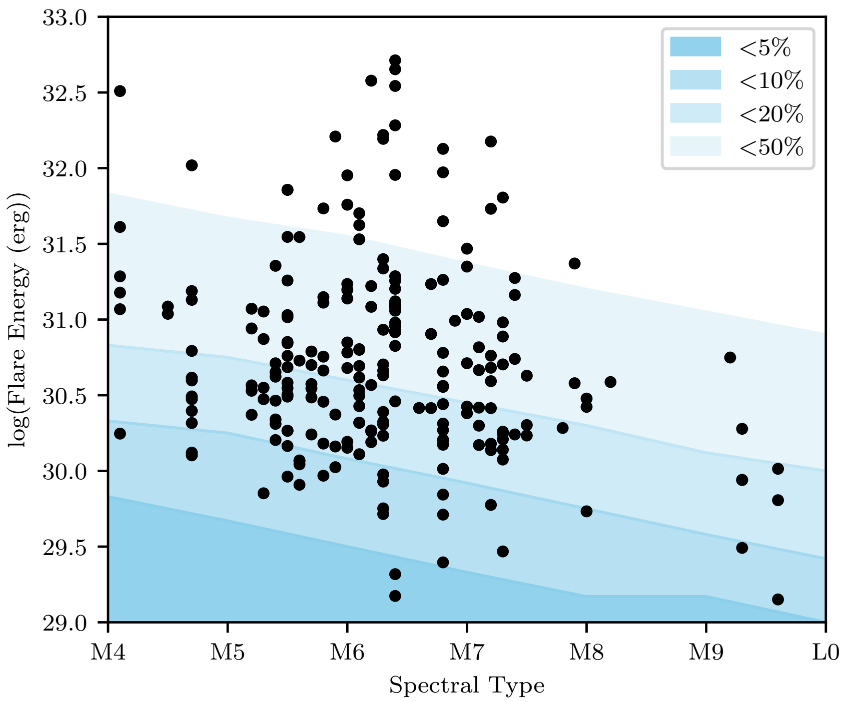

Figure 4.7: Flare energy against spectral type for all detected flares in the SSO sample. By interpolating and smoothing the recovery fractions in Figure 4.4, I include the regions for where the recovery fraction is <5, 10, 20 and 50% (less than this percentage of injected flares were recovered). The drop-off in flares with energies below 10^{30} erg is due to our decreased sensitivity, as described in Section 4.2.5, which more adversely affects the detection of flares on hotter stars. The flares for M9–L0 objects appear to be lower in energy because low energy flares are most frequent and our recovery fraction is higher than for earlier stars. Very few flares are detected on the lowest mass stars, compared to M5–M7 objects which have the highest flaring fractions and of which there are significantly more targets.

from M4 to a broad peak over M5–M7, followed by a decline from M7 to L2.

I probe the parameter space of high flare rate stars with low-to-mid energy flares. Due to the relatively short baselines of SPECULOOS observations, if a target does not flare frequently, it is unlikely I would have identified it as flaring. This also means I am much less likely to have detected the rarer, high-energy flares (Gershberg 1972; Lacy et al. 1976). In conjunction with the difficulties arising from the day-night cycle (typical night observations are 4–8 hours), which complicates the detection of very slowly decaying flares (see Figure 4.2), it is unlikely that superflares with energies > 10^{33} erg (Shibayama et al. 2013) would be detected by the SPECULOOS survey. From studying the completeness of the flare sample, I also conclude it is unlikely that I will detect any flares below $10^{29.5}$ erg for any spectral type.

In this flare sample, I recovered flares ranging in energy from $10^{29.2}$ to $10^{32.7}$ erg, with a median of $10^{30.5}$ erg. I did not detect any superflares, which have energies between 10^{33}–10^{38} erg (Shibayama et al. 2013). The flare detection algorithm's decreasing sensitivity for lower energy flares is responsible for the small numbers of detected flares with energies $<10^{30}$ erg. From the flare injection-recovery analysis, there was also a clear spectral dependence in the flares that could be retrieved. This dependence means that I would struggle to recover flares on the earliest M-dwarfs below an energy of $\sim 10^{31}$ erg. For later spectral type objects, I was able to recover significantly higher fractions of lower energy flares, as shown in Figure 4.4. I present the flare energies of all detected flares as a function of spectral type in Figure 4.7, which demonstrates the impact of the detection sensitivity (and its spectral dependence) on our recovered sample.

Comparison with Seli et al. (2021)

Seli et al. (2021) studied the relationships between age, activity and rotation for a sample of 248 "TRAPPIST-1 analogues" from 30-minute TESS Full Frame Image (FFI) observations. They detected a total of 94 flare events on 21 stars. I compared their target catalogue to the SSO data sample and found 35 objects that appeared in both. Of these 35 objects, I detected 19 as flaring including five of the seven targets that Seli et al. (2021) identified as flaring. The five objects that both surveys identified as flaring stars are Gaia DR2 2331849006126794880, 2349207644734247808, 3200303384927512960 (as shown in Figure 4.1), 4825880783419986432 and 5055805741577757824. The two flaring stars detected from TESS lightcurves, on which I did not detect flares are Gaia DR2 4967628688601251200 and 5637175400984142336. For 14 of the targets in this overlapping sample neither survey detected any flares. The 21 overlapping objects (excluding the 14 that had no flares detected by either study) are presented in Table 4.2.

For 42 stars in their sample Seli et al. (2021) detected rotation periods. Out of the stars that appeared in both UCD samples, Seli et al. (2021) find rotation periods for four objects that also appear in our catalogue. I found rotation periods for 16 of the 35 objects (with rotation classes A, B or L). Therefore, I can also compare both our flaring and rotation results for this handful of targets. For three of these objects I detected an alias of the period extracted from TESS 30-minute FFIs. For each of these three objects I inspected the SSO lightcurves and periodograms. In each case, the rotation period measured by Seli et al. (2021) was the second most promising period, however, the peaks in the periodogram were very close in power, and the final period was chosen to give a better phase coverage. This was the reason for the 'B' grade I assigned to rotators Gaia DR2 2349207644734247808 and 4967628688601251200. Gaia DR2 3200303384927512960 was assigned a grade of 'A', however, it is a unique case. For this target I detected a rotation period of 0.334 d, while Seli et al. (2021) measured a period of

0.50062 d. Rotation periods of 0.33, 0.5 or 1 d are special scenarios where, over short baselines, SPECULOOS only sees the same part of the phase on every observation night. Therefore, due to the limit imposed by the day-night cycle, I am unable to distinguish between these periods. The target for which I did not detect any rotation period, Gaia DR2 4971892010576979840, has an amplitude for periodic photometric variation of 2.1 mmag (measured by Seli et al. (2021)). Upon inspection of this object's SSO and publically-available TESS lightcurves this periodicity was too small to detect above the level of photometric scatter, and residual systematics.

Since TESS is a space-based mission that does not contend with the day-night cycle or ground-based systematics, and they perform 27 day continuous monitoring, this provides a large number of phases for fast-rotating stars, which can counteract lower photometric precisions for red dwarfs. Therefore, I would expect rotation periods derived from TESS lightcurves to be more accurate than those from ground-based observations.

It is, therefore, counter-intuitive that of the 34 overlapping objects, there are 13 for which I detected rotation periods and TESS did not. Five are long period L rotators, four are B rotators, and four are A rotators. The A rotators have periods 3.4, 4.4, 6.5, and 14.7 hrs. The B rotators have periods 0.78, 3.2, 4.3, and 6.8 d. The L rotators have (tentative) periods of 11.5, 11.8, 16.3, 25.7, and 65.5 d. In their analysis Seli et al. (2021) only considered rotation periods shorter than 5 d, so as to probe the fast rotator regime. Therefore the periods of the L rotators (and one of the B rotators) would not be detected by this survey. However, this does not explain their lack of detection of the A and B rotators. In particular I found that the class A rotator Gaia DR2 56252256123908096 has a rotation period of 0.614 d and an amplitude of 29 mmag, well within TESS's capability to detect. Possible explanations are that the 30-minute cadence blurred complex rotation features in the lightcurves, or that the rotation patterns have changed in the time between the SSO observations and TESS sector due to stellar surface evolution. To understand this more clearly, however, I would need to compare with the TESS lightcurves for these objects directly.

For their catalogue of stars, Seli et al. (2021) found that $\sim$8 per cent were flaring. This flaring fraction contrasts this work, where I detected flares on 28 of the 57 M7–M9 objects within the SSO sample (49 per cent). There may be several reasons for this discrepancy. Firstly, Seli et al. (2021) detected flares from lightcurves with a low time resolution of 30 minutes, compared to SPECULOOS which has a time resolution around 20-60 s. Almost half of all the flares (110) detected in the SSO flare sample have flaring regions less than 30 minutes, and even slower-decaying flares would be difficult to detect with only a couple of data points, unless you happened to catch the peak. TESS is also not optimised for very red objects such as UCDs, therefore, the photometric precision declines with decreasing effective temperature. TESS's decrease in precision for very red stars is well-known and is described in Sullivan et al. (2015)

Table 4.2: Comparison of this work and Seli et al. (2021) for 21 overlapping objects with flares detected in either flaring analysis.

Gaia ID	SSO Flaring?	SSO Period (d)	SSO Rotation Class	TESS Flaring?	TESS Period (d)	TESS Ampl. (mmag)
2331849006126794880	Y	N/A	N	Y	N/A	N/A
2349207644734247808	Y	0.687	B	Y	2.13867	20.9
3200303384927512960	Y	0.334	A	Y	0.50062	3.8
4825880783419986432	Y	65.468	L	Y	N/A	N/A
5055805741577757824	Y	25.662	L	Y	N/A	N/A
4967628688601251200	N	0.682	B	Y	0.40391	4.2
5637175400984142336	N	0.573	U	Y	N/A	N/A
3421840993510952192	Y	4.297	B	N	N/A	N/A
3493736924979792768	Y	0.142	A	N	N/A	N/A
3504014060164255104	Y	16.326	L	N	N/A	N/A
4404521333221783680	Y	N/A	N	N	N/A	N/A
4928644747924606848	Y	11.819	L	N	N/A	N/A
4971892010576979840	Y	N/A	N	N	0.70254	2.1
5047423236725995136	Y	N/A	N	N	N/A	N/A
5156623295621846016	Y	6.752	B	N	N/A	N/A
5392287051645815168	Y	0.182	A	N	N/A	N/A
5469802724480366848	Y	N/A	N	N	N/A	N/A
56252256123908096	Y	0.614	A	N	N/A	N/A
5657734928392398976	Y	N/A	N	N	N/A	N/A
6357834388848708224	Y	2.034	U	N	N/A	N/A
6733860940302404864	Y	0.27	A	N	N/A	N/A

and Sebastian et al. (2020) (among others). In Seli et al. (2021) (Figure 2) they reach 3–4 mmag for the brightest targets in their M7–M9 sample. For M7–M9 objects with SPECULOOS, which have magnitudes up to ~14–15, we are able to reach precisions of 0.2–0.3 mmag (see Figure 3.8). For a similar magnitude range (TESS, G_R and SPECULOOS have very similar bandpasses), Seli et al. (2021) obtain precisions of ~ 6 mmag, approximately 20 times as large as SPECULOOS. Both factors inhibit a flare survey on TESS data from detecting low energy flares (which have small amplitudes and short durations) on UCDs. Instead, continuous 27-day monitoring of each TESS field allows Seli et al. (2021) to probe the rarer, high energy (even superflare) flare regime. This is demonstrated by their recovery fraction, which was less than 20 per cent for flares with energies below $10^{31.5}$ ergs, compared to SSO recovery fractions of 50–70 per cent (depending on spectral type) for the same energy. Additionally, later in this chapter I will demonstrate that cooler stars flare less frequently at all energies, therefore, it is likely that a survey focused on high energy flares on cool stars would detect significantly fewer flares. Finally, space-based data is less impacted than ground-based data by time-varying systematics due to the Earth's rapidly-changing atmosphere, and it follows that the risk of false positives is presumably lower.

4.4.2 The Relationship between Rotation and Activity

I recovered 69 (24 A, 22 B, and 23 L) rotators from the SSO data sample, with periods ranging from 2.2 hours to 65 days. I also found 29 U and 60 N class objects. I present a histogram of these rotation classes in Figure 4.8. From here on, I use the term "rotators" to refer only to objects of grades A, B and L.

There were 41 targets for which both flares and rotation were detected. Therefore, of the 78 flaring objects described above, 53% had evident rotation. Alternatively, I detected flares on 59% of the rotators. It appears that while the fraction of rotators across all spectral types stays consistently between 20–50%, there are increasing proportions of slow rotators for the later M9 and L0 stars (see Figure 4.9).

The very fast rotators are much more likely to flare than slow rotators. By comparing the rotation period P with the flaring proportion of stars with rotation periods $\leq P$ (see Figure 4.9), I demonstrated that the likelihood of detecting a star as flaring decreased as the rotation slows. In Figure 4.9, I only included the 46 'A' and 'B' rotators as the uncertainties on the 'L' (long period) rotators are too large. In the SSO sample at least 74% of fast-rotating stars, with $P \leq 2\,\mathrm{d}$, flare. I compare this to flaring proportions of 59% of all rotators, 63% of only 'A' and 'B' rotators, 42% of stars with no detected rotation, and 50% of all stars. I note that the majority (33/46) of 'A' and 'B' rotators have periods $\leq 2\,\mathrm{d}$.

Comparison with MEarth

Within the SSO Sample, 20 targets also appeared in the Newton et al. (2018) MEarth-South rotation sample. Of these 20 objects, our results agreed for eight, and my work suggested periods not found by MEarth for another seven. For the eight objects in agreement: two had long period rotation measured by MEarth, with periods too long to be measured by SPECULOOS (for which I found no rotation); four had no detected rotation with either survey; and two had a similar period measured by both surveys (though one had only a tentative detection with the SSO). For the remaining 12 objects, five had clear, short periods detected by SPECULOOS but not by MEarth, and two had long-period, low-amplitude estimates given by SPECULOOS with no detection in MEarth. I discounted the remaining five objects classified as U/N in both surveys, with missing or disagreeing periods. All the rotation periods measured by the two surveys are shown in Figure 4.10, and their overlapping observations are reported in Table 4.3.

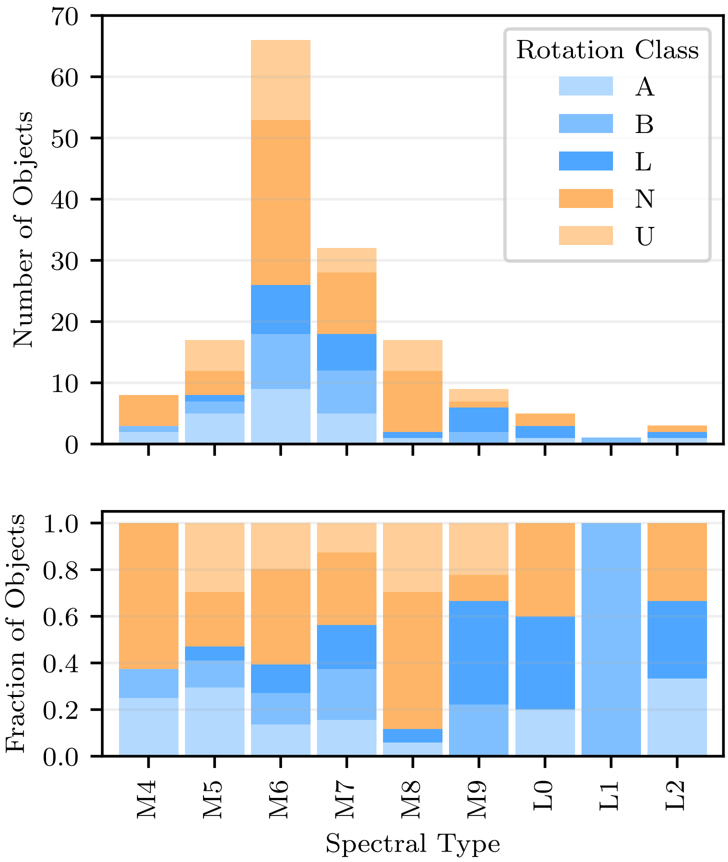

Figure 4.8: **Top:** Histogram of objects in the SSO data sample as a function of spectral type, divided by rotation class. The 'rotators' are shown in varying shades of blue, and the 'non-rotators' are shown in shades of orange. **Bottom:** Fraction of different rotator classes for each spectral type.

The difference in rotation periods measured by the two surveys is likely a result of their different observing strategies. While SPECULOOS continuously monitors every target for 4–8 hours each night over several weeks, MEarth cycles through multiple targets during a night, returning to each at 20–30 minute cadences. SPECULOOS can leverage its observing strategy to detect very short, <5 hour period rotators, whereas MEarth's observing strategy is more suited to measuring very long >50 day rotation periods. I also derive a long-period rotation estimate with SPECULOOS for two objects with small amplitudes that may not be evident in the MEarth data, due to precision limitations on the lowest mass stars. MEarth experiences a drop-off in recovery rate for stars with $M < 0.2M_{\odot}$ (which includes all the objects in the SSO sample), likely due to systematics including their inability to fully correct for the precipitable water effect (Newton et al. 2018). In Newton et al. (2018) they find that their non-detections are biased

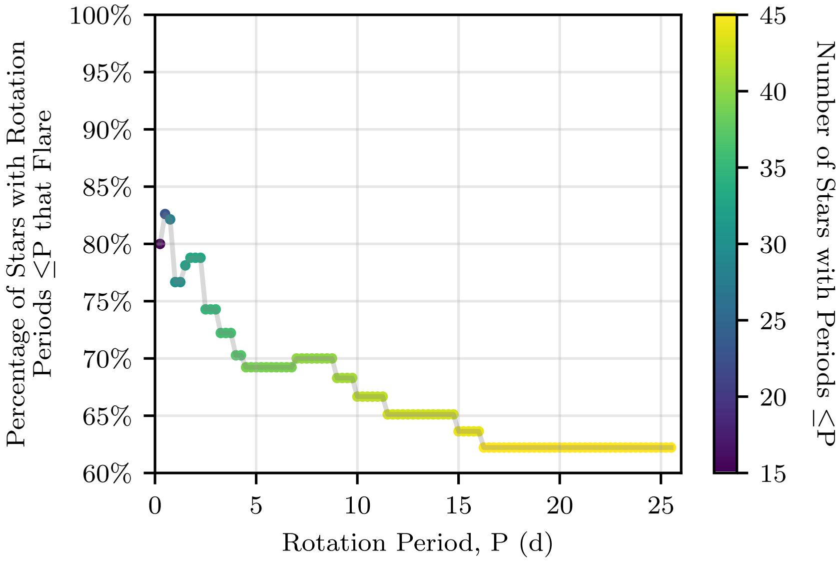

Figure 4.9: The proportion of stars with rotation periods $\leq P$ that have detected flares, against rotation period P, in steps of 0.25 d. I only include the 46 'A' and 'B' rotators, which have periods ranging from 0.09 d to 25.7 d. The colour scale shows the number of stars with rotation period $\leq P$.

towards low-mass stars as 81% of their non-detections have $M < 0.2M_{\odot}$ compared to 60% of their detections. However, the period recovery tests performed in Newton et al. (2018) are only performed for stars with masses down to $0.15M_{\odot}$ (due to this marked decrease in sensitivity), therefore, we do not know their recovery rates for stars similar to the majority of objects in the SPECULOOS sample. Only two of the objects in Table 4.3 have masses above $0.15M_{\odot}$ (6340981796172195584 and 6385548541499112448). Newton et al. (2018) state that these systematics can particularly inhibit the detection of short periods, such as for SPECULOOS's A and B rotators (five of the seven targets in Table 4.3 where MEarth does not detect a period). As for SPECULOOS, MEarth finds that periods close to 1 d are difficult to disentangle from systematics. This challenge of detecting periods around 1 d may be the reason for MEarth's non-detections of targets with Gaia DR2 IDs 5565156633450986752, 6914281796143286784 and 4654435618927743872 in Table 4.3, or indicate that SPECULOOS has detected false positives.

The Newton et al. (2018) sample showed a clear dichotomy of fast rotators, with periods less than 10 days, and slow rotators, with periods greater than 70 days. I was unable to confirm this, as most targets in the SSO sample have been observed for less than a 10-day span. For the clear slow rotators, which are observed for less than a full phase, I could only assign a long-period estimate (L). The apparent gap in rotation periods from the SSO sample between

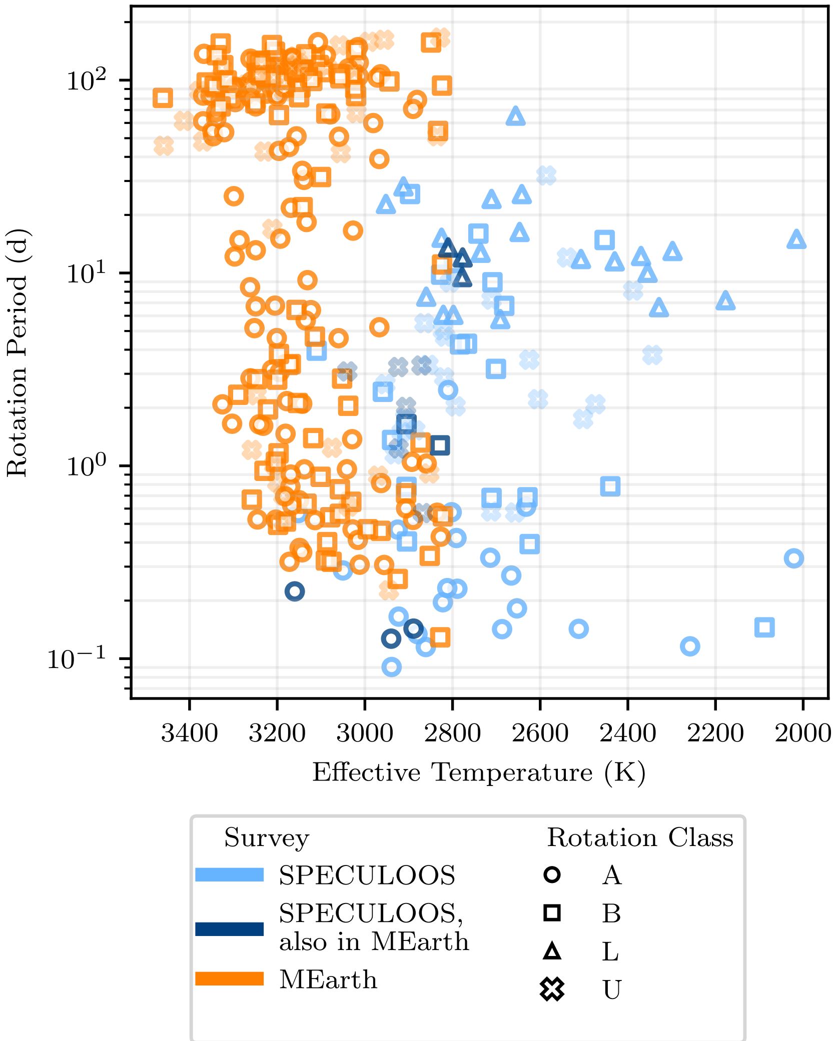

Figure 4.10: A comparison of the rotation periods measured in this SSO data sample with the rotation periods from MEarth (Newton et al. 2018). The objects observed only by SPECULOOS are in light blue, and those observed only by MEarth are in orange. The dark blue points are the measured period in SPECULOOS for objects that have also been observed by MEarth. The different rotation classes are indicated by the marker shapes, with the uncertain (U) rotation periods made transparent. I note that the L rotators, a class unique to SPECULOOS, provide only lower limit estimates of the periods, due to our shorter observing baselines.

Table 4.3: Comparison of SPECULOOS and MEarth rotation periods and classes for overlapping objects in rotation analyses of both surveys (Newton et al. 2018).

Gaia ID	SSO Class	SSO Period (d)	MEarth Class	MEarth Period (d)
4127182375667696128	U	0.57	A	0.571
5072067381112863104	L	12.04	B	11.067
6056881391901174528	N	N/A	B	156.741
6421389047155380352	N	N/A	U	51.202
6340981796172195584	A	0.22	N	N/A
6405457982659103872	A	0.14	N	N/A
2631857350835259392	A	0.13	N	N/A
5565156633450986752	B	1.65	N	N/A
6914281796143286784	B	1.28	N	N/A
6504700451938373760	L	13.56	N	N/A
2393563872239260928	L	9.53	N	N/A
3005440443830195968	U	3.25	U	0.913
4349305645979265920	U	3.31	N	N/A
4654435618927743872	U	1.23	N	N/A
4854708878788267264	U	3.09	N	N/A
6494861747014476288	U	2.02	U	166.165
6385548541499112448	N	N/A	N	N/A
3175523485214138624	N	N/A	N	N/A
3474993275382942208	N	N/A	N	N/A
3562427951852172288	N	N/A	N	N/A

one and two days is likely a bias from the difficulty in disentangling real rotation from 1-day aliases.

4.5 The Consequences for Life on Ultra-Cool Dwarf Planets

The impact of stellar flares on planets hosted by UCDs is multifaceted. On the one hand, flares can cause atmospheric erosion (Lammer et al. 2007), ozone depletion (Segura et al. 2010; Tilley et al. 2019), and in the most extreme case, powerful flares have the potential to cause mass extinction events. On the other hand, UCDs emit strongly in the infrared, with more limited emission in the UV and visible, therefore, several authors find it unlikely that there would be enough incident quiescent UV radiation on planets in the habitable zone for essential photochemical stages of prebiotic chemistry (Ranjan et al. 2017; Rimmer et al. 2018). In addition, these planets do not appear to receive enough visible light for the process of oxygenic photosynthesis, necessary to sustain an Earth-like biosphere (Mullan & Bais 2018; Lehmer et al. 2018; Covone et al. 2021). It is possible that frequent flaring could provide the missing energetic links needed for life to originate (Buccino et al. 2007; Ranjan et al. 2017; Rimmer et al. 2018) and to be sustained (Mullan & Bais 2018; Lingam & Loeb 2019a). Flares can also affect the evolution of a planet's atmosphere (Segura et al. 2010).

Often used to put stellar flares into context, the most energetic solar flare ever directly observed on Earth was the 'Carrington Event' in 1859 (Carrington 1859; Hodgson 1859). This flare released an energy of 10^{32} erg and the associated coronal mass ejection hit the Earth's magnetosphere, resulting in the largest geomagnetic storm on record. This storm disrupted telegraph systems and produced auroras reaching almost as far as the Equator. In the SSO sample I detect 14 flares with energies greater than 10^{32} erg on 10 different stars. Characterising these interactions between planets and their host stars, especially those which have the capability to initiate or destroy life, is vital to our understanding of planetary habitability.

4.5.1 Flare Frequency Distributions

Flare Frequency Distributions (FFDs) allow us to explore how often a star will flare with at least a specific energy. The FFD assumes that the following power law applies (Gershberg 1972; Lacy et al. 1976; Hawley et al. 2014):

$$dN(E) = kE^{-\alpha}dE, \tag{4.10}$$

where N is the flare occurrence rate, E is the flare energy, and k and α are constants. This relation can also be represented as:

$$\log(\nu) = C + \beta\log(\mathrm{E_{flare}}), \tag{4.11}$$

where ν is the cumulative frequency of flares of energies $\geq \mathrm{E_{flare}}$, and the constants C and β are equivalent to $C = \log\frac{k}{1-\alpha}$ and $\beta = 1 - \alpha$ (Hawley et al. 2014). I linearly fit for the two coefficients, C and β, using a least squares optimisation. Physically, β and α represent how often high energy flares and low energy flares occur relative to one another. The larger the value of α, the greater the fraction of the total energy budget is made up by lower energy flares.

To generate FFDs, I extracted only the objects with at least three detected flares to obtain a good linear fit whilst including as many stars as possible. There are 31 objects in the SSO flaring sample (out of 78) with at least three detected flares. I binned the log energies every 0.1 in log space ($10^{28.0}$–$10^{28.1}$ erg, $10^{28.1}$–$10^{28.2}$ erg etc.) and calculated the cumulative frequency of flares greater than that energy. In Figure 4.11 I present these FFDs, along with their best-fit power law from Equation 4.11. I note that this was not corrected for the detection efficiency as in Section 4.5.6, due to the complexity of applying it to the cumulative flare rates for individual stars with low numbers of flares. Instead, in Section 4.5.6, I correct for the recovery rate by combining flares from multiple stars to yield a much larger dataset.

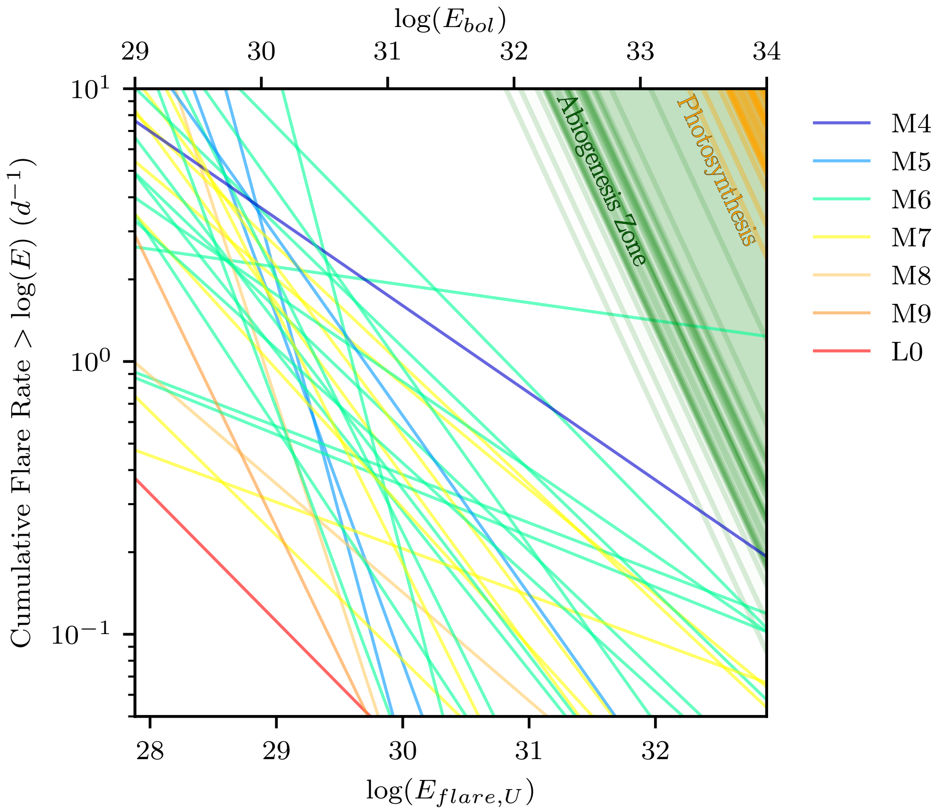

Figure 4.11: Flare frequency distributions for every star with at least three flares. I found the best power-law fit to the cumulative flare rates against flare energy for every star, shown by a straight line. The spectral type of each star is indicated by its colour. The abiogenesis zones are also shown for each star as the regions shaded in green and the photosynthesis thresholds are shown in orange.

Calculating the power-law relationship between flare rate and energy for each target allowed me to predict the flaring rates for high energy flares. By extrapolating this relationship, I could predict the amount of energy that flares would deliver to any planets orbiting those stars. However, extrapolating the power-law relationship from the SSO parameter space (frequent, low energy flares) into the high energy flare regime is dangerously unreliable, as it is common to see breaks in these power-law relationships (Silverberg et al. 2016; Paudel et al. 2018). If I instead treat these linear fits as upper limits then I can give estimates for the maximum frequency of high energy flares. I note that the power laws do not significantly vary with the energy binning chosen for fitting.

Tentatively, I found that as the stars got cooler, the flares detected became less energetic and less frequent. However, there were a large diversity of FFD profiles even within the same spectral type. This could be due to variation in stellar age or metallicity. Since I am limited towards detecting lower energy flares for cooler stars, it is possible that I more significantly underestimated the rates of lower energy flares for the earlier M-stars. This effect may result in steeper power-law slopes, further separating the mid M-dwarfs from the late-M and L dwarfs. However, there was only one M4, M9 and L0 object with at least 3 flares, and within the larger sample of M5–M7 stars, I saw little distinction. The decline in flare rates for cooler spectral types could be partially explained by a decreased detection sensitivity for early-M-stars. If only the infrequent, high energy flares on hotter stars are observable with SPECULOOS, I would have found only the objects that flare most often. Conversely, since my detection method was more sensitive to lower energy flares on late-M and L dwarfs, I was able to identify stars with lower flaring frequencies. This effect explains the lack of low frequency flaring stars of early spectral types, but not the lack of high frequency flaring stars of later spectral types. As my sample of flaring ultra-cool stars is very small (see Figure 4.6), I could not confirm whether they, as a whole, flare less frequently. Paudel et al. (2018) similarly found that the flare rates for low-energy flares decreased as they moved towards later spectral types, with L0 and L1 dwarfs having the lowest flaring rates. Critically, however, they found shallower slopes for cool stars (also in agreement with Gizis et al. 2013 and Mullan & Bais 2018), implying that they had higher occurrence rates for the infrequent high-energy flares, that are mostly inaccessible to SPECULOOS. A high frequency of these flares could be sufficient for a star to enter the abiogenesis zone. However, I am unable to confirm this trend, due to the small parameter space and the large uncertainties on individual power law slopes.

By taking a more conservative threshold of at least five flares to fit the power law, I extracted FFDs for only M5–M7 stars (for which I have a much larger flare sample). For this sub-sample, I found values for α in the range 1.2–2, roughly in agreement with Paudel et al. (2018), who measured 1.3–2 with their sample of 10 UCDs (for a smaller sample, but a similar spectral type range). Paudel et al. (2018) discussed various reasons for the variations in FFD slopes that were not dependent on age or spectral type, such as rotation, stellar spot coverage, and magnetic field topology.

4.5.2 Converting Bolometric Flare Energies to U-band

I calculated the energy in the U-band by integrating the flux density in the U-band spectral response function, as in Günther et al. (2020b). Similarly to Section 4.2.3, I calculated:

$$E_{\mathrm{flare},U} = \int_{\mathrm{flare}} A_{\mathrm{flare}}\,\mathrm{d}t \int_U R_\lambda B_\lambda(T_{\mathrm{flare}})\,\mathrm{d}\lambda. \tag{4.12}$$

Therefore,

$$E_{\mathrm{flare},U} = E_{\mathrm{bol}} \frac{1}{\sigma_{\mathrm{SB}} T^4_{\mathrm{flare}}} \int_U R_\lambda B_\lambda(T_{\mathrm{flare}})\,\mathrm{d}\lambda, \tag{4.13}$$

where R_λ is now the Johnson U-band response function. From this, I estimated that 7.6% of the flare's bolometric energy fell in the U band.

4.5.3 The Abiogenesis Zone

Here I considered the laboratory work of Rimmer et al. (2018) in defining "abiogenesis zones" around each of our potential planet hosts, outside of which it is unlikely a planet would receive enough energy for the synthesis of ribonucleotides as a precursor to ribonucleic acid (RNA), to occur (Patel et al. 2015; Sutherland 2017; Xu et al. 2018). By considering stellar flares as a mechanism for providing this UV energy, I calculated the abiogenesis zones around the stars in the sample from their FFDs. Whether or not this prebiotic photochemistry is possible on a planet does not tell us if life has originated there, but instead, whether this mechanism can allow that planet to generate this first building block for RNA. Conversely, if a planet does not receive the necessary energy for this reaction, it does not rule out alternate prebiotic pathways for the origins of life.

Günther et al. (2020b) determined the necessary flare frequency to power prebiotic chemistry for a planet receiving the same amount of flux from its host as the Earth, by adapting the abiogenesis zone equations from Rimmer et al. (2018) as follows:

$$\nu \geq 25.5 \mathrm{day}^{-1} \left(\frac{10^{34}\,\mathrm{erg}}{E_U} \right) \left(\frac{R_*}{R_\odot} \right)^2 \left(\frac{T_*}{T_\odot} \right)^4 , \tag{4.14}$$

From Section 4.5.1 I calculated the flare energies in the U-band from each flare's bolometric energy, using $E_U = 7.6\% E_{\mathrm{bol}}$. I overplotted this threshold for each star in Figure 4.11.

When the FFD power laws were extrapolated to $E_{\mathrm{bol}} = 10^{34}$ ergs, only one star provided the necessary UV flux to reach the abiogensesis zone (see Figure 4.11). This star is an M6 object. If I extrapolated the FFD power laws instead to $E_{\mathrm{bol}} = 10^{36}$ ergs, then there were 13 objects that reached the abiogenesis zone (with spectral types consisting of 1 M4, 8 M6, 3 M7 and 1 M8). However, as the SSO sample is confined to only probing the low energy flare regime, I must be careful to not over-extrapolate the FFDs to high energies. Therefore, I did not extend the power law predictions further than $E_{\mathrm{bol}} = 10^{34}$ ergs. Almost all the FFDs will eventually reach the abiogenesis zone, but at very high energy the uncertainties in the power law fit would be too large to draw any real conclusions.

4.5.4 Oxygenic Photosynthesis

Lingam & Loeb (2019a) defined a threshold for sustaining a biosphere on an Earth-like planet using flare-driven photosynthesis. By considering "photosynthetically active radiation" in the region of 400-750 nm, they found a similar functional form to Equation 4.14 for the minimum flare rates necessary to receive the same photon flux on a temperate planet as on Earth:

$$\nu \geq 9.4 \times 10^3 \mathrm{day}^{-1} \left(\frac{10^{34}\,\mathrm{erg}}{E_{\mathrm{Bol}}}\right) \left(\frac{R_*}{R_\odot}\right)^2 \left(\frac{T_*}{T_\odot}\right)^4, \tag{4.15}$$

This threshold is a significantly greater inhibitor than that for prebiotic chemistry. I overplotted this threshold for each star in Figure 4.11. None of the stars exceeded this threshold when extrapolating to $E_{\mathrm{bol}} = 10^{34}$ erg. Because the condition for oxygenic photosynthesis is more strict than Equation 4.14, if an object satisfies Equation 4.15, it must also meet the requirements for the abiogenesis zone.

4.5.5 Ozone Depletion

Modelling of the impact of stellar flares on modern Earth-like analogues by Tilley et al. (2019) showed that stars with at least 0.1 flares per day with energies above 10^{34} ergs will strip the ozone layers from any terrestrial planets they host. I did not extend the FFDs to energies above 10^{34} ergs as I am probing the low energy flare regime; therefore, I did not consider ozone depletion in my analysis.

4.5.6 Applying the Sensitivity to the Flare Sample

To compare the average flare occurrence rates for different spectral types, I had to account for the incompleteness of the sample. Due to the difficulty of detecting small-amplitude flares above photometric scatter, I likely underestimated the frequency of low energy flares. This underestimation is often seen as a non-linear 'tail-off' from the expected power law in log-log space (Equation 4.11).

Jackman et al. (2021) derived the following equation (equivalent to their Equation 3) for the number of flares, N, with energies greater than E_{flare}:

$$N(E \geq E_{\mathrm{flare}}) = \frac{k}{\alpha - 1}\left(R(E_{\mathrm{flare}})E_{\mathrm{flare}}^{-\alpha+1} + \int_{E_{\mathrm{flare}}}^{E_{\mathrm{max}}} R'(E)E^{-\alpha+1}\,\mathrm{d}E\right), \tag{4.16}$$

where k and α are the same constants in Equation 4.10, $R(E)$ is the flare recovery fraction, $R'(E)$ is the differentiated flare recovery fraction, and E_{max} is the energy at which the recovery fraction saturates. For high energy flares, where $R = 1$, Equation 4.16 reduces to the power law in Equation 4.10, whereas for low energy flares, where $R = 0$, it will reduce to the constant value of the integral (equivalent to the 'tail-off' effect).

Due to the small population of flares on the earliest and latest stars in the flare sample, I built up flare numbers by combining spectral types into the following five bins: M4–5, M6, M7, M8 and M9–L0. I calculated every star's unique recovery fraction, $R(E)$, based on the results of the injection-recovery tests. Flare energies were binned into 12 logarithmically spaced bins

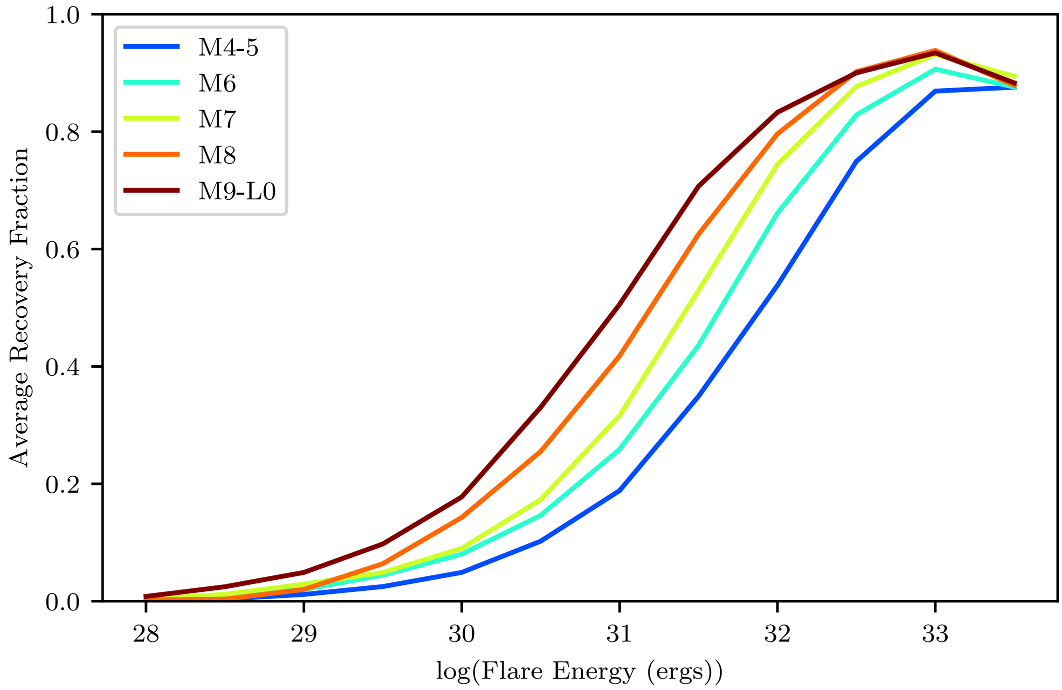

Figure 4.12: The average smoothed recovery fractions against flare energy of stars in each of the five spectral type bins. As expected, there is an increased recovery of all flare energies for the coolest dwarfs.

from 10^{28} to 10^{34} ergs. For every star, I found a value for the recovery fraction in each energy bin by calculating the fraction of flares recovered with energies in that bin. I smoothed the recovery fraction using a Wiener filter of three bins (shown for every star in Appendix D), and averaged the recovery fractions within each spectral type bin to get the recovery fraction of an "average" star. Figure 4.12 shows the average recovery fractions for each spectral type bin. There was, counter-intuitively, a very slight decrease in the recovery fractions for the highest energy flares for all spectral types. However, these flares have the longest decay times, which would be difficult to recover with the length of SPECULOOS's observations on a typical night (4–8 hrs). This theory was supported by the large range in recovery fractions for the highest energy flares within each spectral type bin, shown in Appendix D. Some stars will have nights of observation less than 4 hours due to weather or limited visibility, and these will be the most difficult targets on which to detect slowly decaying flares.

I then extracted the observed flare occurrence rates from the flare sample. First, I isolated only the stars in each spectral type bin. I included all flaring and non-flaring stars in the SSO sample to avoid overestimating the flaring rates. For every energy bin, E, I summed the total number of flares with energies $\geq E$. Then I divided the total number of flares by the sum of

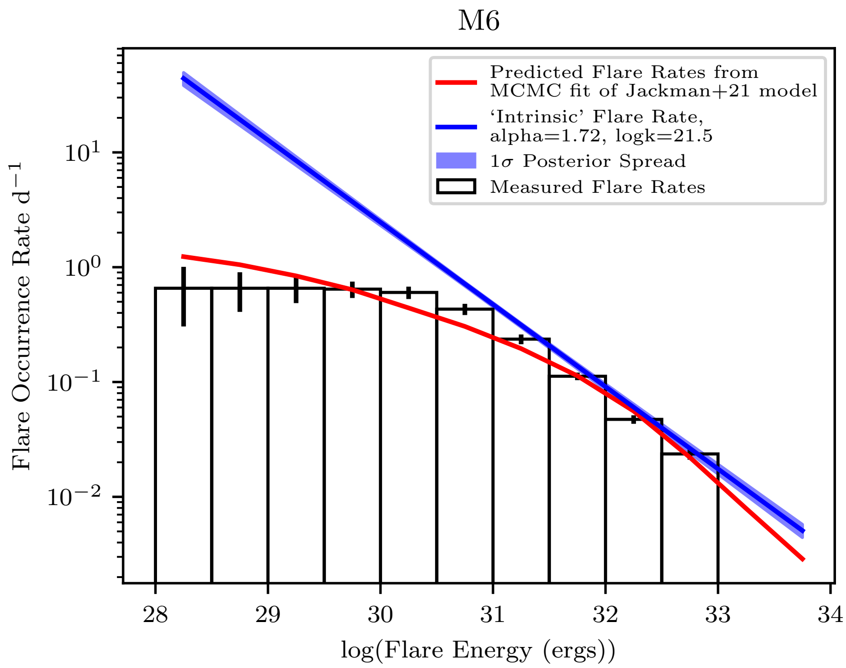

Figure 4.13: The best MCMC fit of Equation 4.16 to the observed flare occurrence rates for the spectral type bin M6. The measured flare rates from the SSO data sample, and corresponding errors, are shown by the bar plot in black. The best fit of Equation 4.16, which accounts for the incompleteness of the flare recovery process, to the observed occurrence rates is shown in red. The 'intrinsic' flare rate (and the 1σ posterior spread) is shown in blue.

the stars' total observation times to produce flaring rates. In doing so, I assumed that 10 stars observed for 10 hours each is equivalent to 1 star observed for 100 hours.

I fit Equation 4.16 to the observed flare occurrence rates and energies using a Markov Chain Monte Carlo (MCMC) parameter optimisation procedure. MCMC methods are a class of algorithms that can sample a probability distribution using Markov Chains to perform Monte Carlo estimates. Markov chains can be visualised as performing a random walk on a chain where the next step is only determined by a probability distribution given the current state (and no previous steps). Monte Carlo methods sample randomly from a probability distribution and use those samples to approximate some value that we would like to find. Practically, in MCMC an ensemble of Markov chains is created, starting from random, widely-dispersed values. These chains are known as "walkers" and follow many successive stochastic steps until we can approximate the true solution. The initial steps are often referred to as “burn-in” steps and are discarded before sampling the posterior distribution.

To implement MCMC I used the EMCEE Python package (Foreman-Mackey et al. 2013). I used 32 walkers for 10,000 steps and the last 2000 steps to sample the posterior distribution.

Table 4.4: The observed flare occurrence rates per day for each spectral type and flare energy bin. The parameter ν_{28} represents the rate of flares per day with $E \geq 10^{28}$ ergs. I did not detect any flares with energies $\geq 10^{33}$ ergs

Spectral Type	ν_{28}	$\nu_{28.5}$	ν_{29}	$\nu_{29.5}$
M4–M5	0.7 ± 0.9	0.7 ± 0.6	0.7 ± 0.4	0.7 ± 0.2
M6	0.7 ± 0.4	0.7 ± 0.2	0.66 ± 0.17	0.64 ± 0.11
M7	0.6 ± 0.5	0.6 ± 0.3	0.58 ± 0.16	0.56 ± 0.12
M8	0.2 ± 0.2	0.17 ± 0.17	0.17 ± 0.07	0.17 ± 0.04
M9–L0	0.14 ± 0.09	0.14 ± 0.05	0.14 ± 0.04	0.10 ± 0.02
Spectral Type	ν_{30}	$\nu_{30.5}$	ν_{31}	$\nu_{31.5}$
M4–M5	0.68 ± 0.17	0.44 ± 0.08	0.24 ± 0.05	0.055 ± 0.014
M6	0.60 ± 0.07	0.43 ± 0.05	0.24 ± 0.02	0.112 ± 0.009
M7	0.53 ± 0.08	0.30 ± 0.04	0.140 ± 0.016	0.060 ± 0.008
M8	0.15 ± 0.03	0.068 ± 0.014	0.017 ± 0.005	–
M9–L0	0.052 ± 0.009	0.017 ± 0.006	–	–
Spectral Type	ν_{32}	$\nu_{32.5}$	ν_{33}	$\nu_{33.5}$
M4–M5	0.037 ± 0.009	0.018 ± 0.007	–	–
M6	0.047 ± 0.004	0.024 ± 0.003	–	–
M7	0.020 ± 0.002	–	–	–
M8	–	–	–	–
M9–L0	–	–	–	–

As in Ilin et al. (2019) and Jackman et al. (2021), I multiplied the errors by $R(E)^{-0.5}$ to account for larger uncertainties in recovering the smallest energy flares. An example of the MCMC fit to the spectral type bin M6 is shown in Figure 4.13, and the MCMC fits for all spectral types are presented in Appendix D. The observed flare occurrence rates for each spectral type bin are presented in Table 4.4 and the results of the best-fit power laws for each spectral type bin are in Table 4.5, and demonstrated in Figure 4.14. For the spectral type bins I obtained similar values of α = 1.88 ± 0.05, 1.72 ± 0.02, 1.82 ± 0.02, 1.89 ± 0.07, and 1.81 ± 0.08 for M4–5, M6, M7, M8, and M9–L0 respectively. This work demonstrates that, on average, the power-law gradient does not change with spectral type in the mid-to-late M-dwarf regime, despite the large variations within each spectral type. The implication is that mid and late M-dwarfs produce similar relative proportions of high and low energy flares; however, the decreasing y-intercept means that the coolest stars have lower rates of flares of all energies.

My results are consistent with three similar studies of the relation between flaring rate and

Table 4.5: The best-fit power law α, for each spectral type bin, as determined using MCMC.

Spectral Type	α
M4–M5	1.88 ± 0.05
M6	1.72 ± 0.02
M7	1.82 ± 0.02
M8	1.89 ± 0.07
M9–L0	1.81 ± 0.08

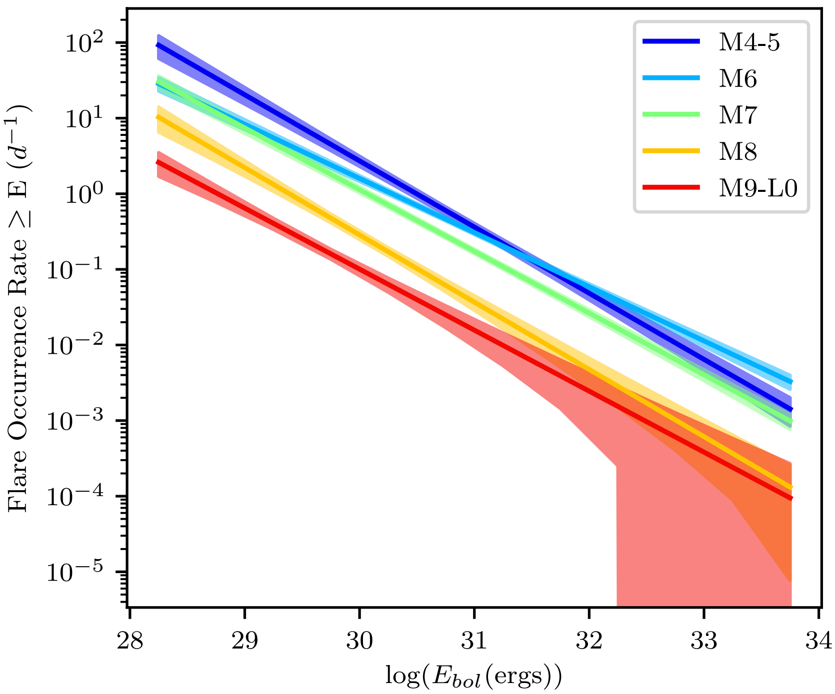

Figure 4.14: The computed flare occurrence rates (Equation 4.10), and their 1σ uncertainties, against bolometric flare energy for an average star of each spectral type bin.

energy for cool stars. Paudel et al. (2018) measured a range of α from 1.3–2.0 for 10 UCDs from K2 lightcurves. Raetz et al. (2020) compiled K2 lightcurves and derived a value for α of 1.83 ± 0.05 for FFDs of the coolest stars in their sample (with spectral types M5–M6). Gizis et al. (2017) found $\alpha = 1.8$ for three UCDs. However, my results have noticeably lower values of α than several other works on mid-to-late M-dwarfs. Both Yang et al. (2017) and Yang & Liu (2019) generated flare catalogues from Kepler lightcurves, and found, on average, fully convective M-dwarfs had power law indices of $\alpha = 2.09 \pm 0.10$ and $\alpha = 2.07 \pm 0.35$ respectively. For a sample of cool dwarfs observed by TESS, Medina et al. (2020) found $\alpha = 1.98 \pm 0.02$. Seli et al. (2021) isolated TESS observations of "TRAPPIST-1 analogues", for which they derived $\alpha = 2.11$. This work also modified the FFD for TRAPPIST-1 generated by Vida et al. (2017) (to include recovery rates), updating their value of α from 1.59 to 2.03 ± 0.02. Lin et al. (2021) used EDEN and K2 data to analyse the flaring activity of the nearby, active M-dwarf Wolf 359. For this target they found $\alpha = 2.13 \pm 0.14$. The lower value of α I find is possibly due to the incompleteness of the sample not fully being described by Equation 4.16. The early tail-off at low energies compared to the Jackman et al. (2021) model, seen for all spectral types (Appendix D), may be evidence of this. Alternatively, previous works may find higher values for α than this work as they have studied higher energy flares on UCD targets using space-based observations. As previously mentioned, I am limited in this study to only frequent, low energy flares due to the observing strategy. Therefore, power laws may steepen at higher flaring energies.

This analysis agrees with the results of Section 4.5.1, which predicted a decline in flaring rates as I moved to the coolest stars. The injection-recovery tests demonstrated that if flares were present on the lowest-mass stars, I would have been able to detect them; therefore, these flares must occur too infrequently to be seen with SPECULOOS's survey strategy.

4.5.7 Habitability of Earth-sized planets around UCDs

From calculating FFDs for the stars in the sample, I placed energetic bounds on whether flares could drive photosynthesis (Section 4.5.4) or provide their planets with enough UV flux for the prebiotic chemistry described in Section 4.5.3 to occur. When the FFD power laws were extrapolated to $E_{\mathrm{bol}} = 10^{34}$ erg, there was only one M6 star that provided the necessary UV flux for abiogenesis, and no stars that passed the threshold for Earth-like photosynthesis. However, there are several caveats to these results. As I saw a non-linear "tail-off" effect for low energy flares due to our detection limits, I likely underestimated low energy flare rates, resulting in shallower FFD power laws in Section 4.5.1. Therefore, without accounting for the completeness of the flare sample, this work could place only upper bounds on the targets whose potential planets can reach the abiogenesis zone or sustain photosynthesis. Towards cooler stars, the frequency of flaring appeared to decrease (see Figure 4.11). However, only one each of the

M4, M9 and L0 stars had at least three flares detected. Due to the low detection efficiency for mid M-dwarfs, I am biased to detect only the most frequently flaring stars; therefore, it is likely that I would have seen an artificial increase in flare rate for M4 objects. In contrast, my flare detection method was more sensitive to flares of all energies on the late-M and L dwarf stars, so this bias does not explain the decrease in flare frequency for these types of objects.

By including the detection sensitivity in Section 4.5.6 and considering an "average" star for each spectral type, I could conclude more general relationships between flare frequency and spectral type. While there was a wide diversity of flare frequencies for individual stars within a given spectral type, on average, I found that cooler stars had lower flaring rates (see Figure 4.14). If this trend holds for a larger sample of UCDs, it may have severe consequences on the search for life around the very lowest mass objects. Similar conclusions have been reached by work on the flaring activity of TRAPPIST-1 (Glazier et al. 2020; Ducrot et al. 2020; Seli et al. 2021) and "TRAPPIST-1 analogues" (Seli et al. 2021).

If most late-M and L dwarfs have flaring rates too low to initiate the synthesis of ribonucleotides, then life would have to originate through a different mechanism. Possible alternatives are surface hydrothermal vents (Rimmer & Shorttle 2019), aerial biospheres (Lingam & Loeb 2019b), impact-shock synthesis (Furukawa et al. 2015) or extraterrestrial delivery (Rimmer & Shorttle 2019; Krijt et al. 2017). I also found that no objects in the sample met the requirement set out by Lingam & Loeb (2019a) to sustain an Earth-like biosphere due to the significantly lower photosynthetically active radiation (with wavelengths of 400–750 nm) incident within UCD habitable zones. These results are consistent with the work of Mullan & Bais (2018) and Lingam & Loeb (2019a). However, I have adopted an Earth-centric viewpoint, and it has been suggested that photosynthesis could proceed using deep-sea hydrothermal vents (Beatty et al. 2005) or at infrared wavelengths, where UCDs are brightest (Scalo et al. 2007; Takizawa et al. 2017; Claudi et al. 2021).

The activity of M-dwarfs is predicted to decrease as they age (West et al. 2008; Paudel et al. 2019). Therefore, it is possible that earlier in these stars' life cycles, they would have produced more energetic and frequent flares, with enough UV flux to trigger prebiotic chemistry on their planets or even for photosynthesis. The results in this chapter represent only a snapshot of each object at its current point in time. I did not consider age in this work due to the considerable uncertainty on ultra-cool objects' stellar age estimates.

In this chapter, I have discussed only two aspects of the impact of flares on the habitability of planets around UCDs. In reality, I would need to apply a more holistic approach to understand whether these planets can support life. There is a multitude of other factors (such as orbital dynamics, stellar variability, tidal locking, atmospheric composition, stellar age, etc.) that influence the environments on planets hosted by UCDs, and other possible pathways for life that have not been considered here.

4.6 The Evolution of Stellar Activity from Mid-M to Ultra-cool Dwarfs

The breadth of the SSO flare sample allows me to compare mid M-stars' flaring and rotating behaviour with ultra-cool dwarfs. I predominantly detected flares on M5–M7 stars (~60-70 per cent flare) with decreasing proportions of flaring stars for both earlier and later M-dwarfs. Since there was a spectral dependency to the flare detection algorithm (Figure 4.4), it was likely that I missed the lower energy flares on mid M-dwarfs and underestimated the proportion of flaring stars. On the other hand, these higher mass objects do produce rarer, high-energy flares more frequently than late M-dwarfs (Lacy et al. 1976). In theory, these high-energy flares should be easier to detect. My results also agree with other surveys, with around ~ 30 per cent of M4 stars flaring. However, these surveys were potentially limited in different ways for mid M-stars that exist at the redder end of their samples, compared to the bluer end of ours.

Lacy et al. (1976) showed that the flares on mid M-dwarfs were more frequent and lower energy than on early M-dwarfs. Hilton et al. (2010) found that this trend continued into the late-M-dwarf regime (M6–M8). From my findings, I saw little difference in flaring frequency between spectral types up to M7. However, my results extended beyond M8, where flares of all energies became less frequent, and I found very little change in α across all spectral types (implying that UCDs produce the same ratio of low to high energy flares as mid M-dwarfs). From the results of injection-recovery tests, I concluded that if higher energy flares were present on the reddest stars in this sample, then I would have detected them; therefore, the lack of high-energy flares for late-M and L dwarfs on the sampled timescales of 0.8 to 8.7 days is likely genuine.

I saw a relationship between flaring activity and rotation, especially for the very fastest rotators. The stars with rapid rotation periods of <5 days, were much more likely to be flagged as flaring than the rest of the SSO sample. At least 74% of the stars with rotation periods shorter than two days flared, compared to around 59% of all rotating stars and 42% of non-rotating stars. Therefore, I was more likely to detect flares on stars for which I could determine a rotation period than stars where I could not. Possible explanations are that fast rotators flare much more frequently, more energetically, or are more likely to flare than stars with no detectable rotation period. There are several reasons we would be unable to detect a rotation period. A star may have a rotation period too long to be measured by the SPECULOOS survey. The photometric variations on the rotation timescale may be too small to be seen, possibly due to small spot covering fractions, a low spot-photosphere contrast, or a particular orientation of spots (e.g., distributed at the poles). This contrast between rotators and non-rotators was further enhanced for the very fastest periods ($P \leq 2$ d). Assessing the spectral distribution of the fastest rotators showed that almost all the M4–M7 stars with periods less than two days flare. Comparatively, I found very few objects with spectral types later than M8 flared, fast-rotating or otherwise. Stars with fast rotation will have more magnetic energy available than slow rotators; therefore, these results are consistent with faster rotators exhibiting more dynamic behaviours (Davenport 2016; Allred et al. 2015; Newton et al. 2017; Yang & Liu 2019; Günther et al. 2020b; Martínez et al. 2019). Specific activity-rotation studies on cool stars, such as those carried out by West et al. (2015), Stelzer et al. (2016), Medina et al. (2020), and Raetz et al. (2020), also found that fast-rotating stars flare more frequently than slow rotators. Seli et al. (2021) found that ∼50 per cent of their fast-rotating ($P < 5$ days) late-M stars flared, compared to ∼70 per cent of our sample with the same periods.

In this chapter, I compared results for SPECULOOS targets that also appeared in the rotation study performed by Newton et al. (2018) on MEarth data and in the activity-rotation study from Seli et al. (2021) on TESS observations. SPECULOOS and MEarth are two ground-based surveys focused on cool stars but with vastly different observing strategies. SPECULOOS is a fast-cadence survey that targets ultra-cool dwarfs for relatively short baselines of 4–8 days. MEarth, however, revisits its M-dwarf targets only every 20–30 minutes, over several months. With SPECULOOS's lightcurves, therefore, we can study the short-term photometric variability of an object, such as its flaring and fast rotation (<10 d). In contrast, with MEarth's lightcurves, we can monitor an object's long-term variability, such as any slow rotation (>50 d) and stellar surface evolution. SPECULOOS and MEarth, however, are both ground-based surveys that have to contest with the challenges of observing through the Earth's atmosphere and the loss of data from the day-night cycle. TESS, a space telescope, performs continuous 27 d monitoring of each field; however, ultra-cool dwarfs are not its ideal targets. Therefore, TESS lightcurves of UCDs are mostly only available from the 30-minute FFIs and generally have lower precisions. SPECULOOS has the short exposure times and high precisions necessary to identify frequent low energy flares; however, for rotation, the irregular data sampling and large gaps (limited by weather conditions and the observing cycle) make it challenging to identify the correct period above its aliases. For this reason, I recommend caution when interpreting the rotation results from this survey and have provided 'rotation classes' to help identify targets with rotation period uncertainties. TESS can only detect the rarer high energy flares (with large amplitudes and slow decay times). However, its ability to provide continuous phase coverage of fast-rotating objects, and observe up to hundreds of cycles, makes it well suited for short-period rotation studies. Comparisons, such as these, between ground and space-based surveys and between fast and slow-cadence observing strategies, highlight that they are highly complementary.

Leveraging the differences and the overlaps between different surveys allows us to fully explore the photometric activity parameter space of UCDs, and unearth more information about the subsequent impact on any potential planets they might host.

4.7 Conclusions

By analysing the high-precision photometry produced by SPECULOOS, I monitored the activity and rotation of late-M and ultra-cool dwarfs. This work allowed me to probe further into the ultra-cool regime than previous photometric surveys and expanded our understanding of the relationships between activity and rotation into a new parameter space. I demonstrated that this work could be used in conjunction with previous and ongoing activity-rotation studies on earlier M-dwarfs, such as those carried out on TESS, K2, and MEarth observations, to assess the diversity of M-dwarfs and how well we could extrapolate the results of these studies to lower mass objects. In this chapter, I presented the benefit of a large-scale flare study of UCDs, such as this, to extend existing flare catalogues.

Using a flare sample containing a wide range of spectral types, I made general predictions about how flaring activity changes as we consider very low mass stars. I also addressed two different scenarios in which flares may assist the origin and sustenance of life on planets they host. Firstly, I found that the quiescent UV flux of cool stars was too low for the synthesis of ribonucleotides, a major step in prebiotic chemistry. Secondly, there would not be enough visible light on planets around cool stars for Earth-like photosynthesis. In both cases, while I found it was possible that flares on highly active, warmer M-dwarfs could provide the necessary extra radiation, the frequency of flares of all energies appeared to decrease with spectral type. The ratio of low to high energy flares remained consistent across all spectral types. These two considerations alone suggest that UCD systems may not be favourable sites for abiogenesis, though other factors—known and unknown—certainly impact the likelihood of finding life in these systems as well.

Chapter 5

The Frequency and Diversity of Temperate, Earth-sized Planets around UCDs

The final stage of this thesis is to consider the planets orbiting ultra-cool dwarfs. The detection of the TRAPPIST-1 planetary system among only 50 UCD targets, established an optimistic start for the SPECULOOS project. However, to date, SPECULOOS has not detected any planets since the initial TRAPPIST-1 system. To find planets, the SPECULOOS team visually inspects each night's lightcurves to flag transit-like features. However, the presence of stellar activity and atmosphere-induced systematics can severely limit our ability to detect planets manually. Additionally, assessing single transits on a night-by-night basis ignores the improved detection capability when stacking multiple transits.

In this chapter I develop a transit-search algorithm. This algorithm begins with an automatic detrending method for correlated systematics and stellar variability, using Gaussian Process regression. It then searches the detrended lightcurves for transit-like structures using an automated box-fitting software. By applying this search algorithm to the SSO lightcurves I look for transiting planets that may have been missed by manual inspection.

To then derive information about the underlying planetary populations around ultra-cool dwarfs, it is essential to first understand the detection efficiency. For this purpose, I performed transit injection-recovery tests on the SSO lightcurves. These tests quantify the retrieval of planets based on their size and orbital period. Finally, using these results, I make predictions about how common planets similar to TRAPPIST-1b are. The content of this chapter will appear in a publication, currently in preparation, about the frequency and diversity of ultra-cool dwarf planetary systems.

5.1 Gaussian Processes as a Tool for Removing Variability

5.1.1 What is a Gaussian Process?

Gaussian processes (GPs) are a method of deriving the functional relationship between variables. Often, in regression analysis, we know the function, which allows us to find a probability distribution over the parameters that is consistent with the data (such as with linear regression). With GPs, however, we find a probability distribution over the infinite possible functions that fit the data. I will not go into mathematical detail here, but for a more exhaustive overview of Gaussian Processes I direct the reader to the book by Rasmussen & Williams (2006).

Similar to how a Gaussian distribution is defined by a mean and a variance, a Gaussian Process is described by a mean function μ and a covariance function (defining the covariance matrix) k. The diagonal elements of the covariance matrix, σ_i^2, describe the variance of variable i and the off-diagonal elements, σ_{ij}, represent the covariance (i.e., similarities) between variables i and j. When applying a GP to a dataset we can specify a prior mean and covariance, known as a kernel. Then the hyperparameters of the kernel can be optimised by applying the GP to training data and maximising the log-likelihood function.

5.1.2 GPs Applied to Exoplanet Problems

Applying Gaussian Processes within the framework of exoplanets is well-established due to their flexibility and ease of use. GPs can be applied to remove systematic noise (e.g., spatial correlations) and stellar variability from photometry, while preserving transit signals (Aigrain et al. 2016; Luger et al. 2017; Serrano et al. 2018; Sestovic & Demory 2020; Lienhard et al. 2020; Barros et al. 2020; Smith et al. 2021a). GP regressions have also been used successfully to account for activity-related effects in radial velocity measurements (Aigrain et al. 2012; Haywood et al. 2014; Rajpaul et al. 2015; Faria et al. 2016) and for instrumental systematics in transmission spectroscopy (Gibson et al. 2012), amongst many other applications (Ilin et al. 2021; Borsato et al. 2021; Luger et al. 2021).

In the exoplanet literature, the *squared exponential kernel* (Rasmussen & Williams 2006) is frequently used for correlated noise (Aigrain et al. 2016; Luger et al. 2017; Sestovic & Demory 2020), as is the *quasi-periodic kernel* (Rasmussen & Williams 2006) for stellar variability (e.g., Haywood et al. 2014; Grunblatt et al. 2015; Rajpaul et al. 2015; Aigrain et al. 2016; Faria et al. 2016; Luger et al. 2017; Sestovic & Demory 2020). I will further define both kernels in the next section. For activity-induced signals, the *quasi-periodic kernel* is often preferred over a strictly periodic kernel, as it allows for the evolution of periodicity, for example, if stellar surface features change over time. CELERITE (Foreman-Mackey et al. 2017) and GEORGE (Foreman-Mackey et al. 2014; Ambikasaran et al. 2016) are two commonly used implementations of GP regressions, designed for astrophysical applications.

5.1.3 GP-Detrending of the SSO Variability

Transit-finding tools often require a 'flat' lightcurve to work optimally. Therefore, to model and remove the stellar variability (and any remaining systematics) exhibited by the SSO lightcurves I used a Gaussian Process approach (Rasmussen & Williams 2006). The central challenge was to find a function that modelled the long-term behaviour of the data (such as periodic brightness modulations) and nightly atmospheric variations, without modelling transit features. In this scenario, the covariance matrix determines the similarity between two flux measurements separated in time. Therefore, the kernel can capture periodic flux variations as well as smooth gradients caused by changing atmospheric conditions over the course of a night.

I used the package GEORGE for GP regression, due to the large number of kernels it makes available. Whilst GEORGE is less restrictive in kernel choice than CELERITE, it is computationally slower (Foreman-Mackey et al. 2017). From the previous chapter I have gained information about each target's rotational behaviour, which was used to inform the choice of kernel. For targets which had no detected periodicity, I used a *squared exponential* kernel (Rasmussen & Williams 2006). This kernel dictates that the correlation between data points exponentially decreases as more time passes between them. It is defined as:

$$k(t_i, t_j) = A \exp\left(-\frac{(t_j - t_i)^2}{l^2}\right) \tag{5.1}$$

where i and j are two different data points in the lightcurve (taken at times t_i and t_j), A is the amplitude of correlation and l is the length-scale over which correlations decay.

If instead, a target does have detected periodicity I used a *quasi-periodic* kernel (Rasmussen & Williams 2006):

$$k(t_i, t_j) = A \exp\left(-\frac{(t_j - t_i)^2}{l^2} - \Gamma \sin^2 \frac{\pi(t_j - t_i)}{P}\right) \tag{5.2}$$

where Γ describes the scale of the correlations, and P is the log of the period. The quasi-periodic kernel defines a periodic correlation that weakens over time. While it is unlikely that I would observe stellar spot evolution within the short timescales for which most SPECULOOS targets are observed (<200 hours), I included the aperiodic component (rather than using a strictly periodic kernel) as we may revisit objects years later. Therefore, in the case of a non-periodic lightcurve there are two hyperparameters to optimise (A, l), compared to four for the periodic case (A, l, Γ, P). GEORGE can also fit the white noise in the lightcurve, by adding a constant value along the diagonal of the covariance matrix. I did not have training data to optimise the hyperparameters of the kernel, instead I used the lightcurve data directly to determine their values. I optimised for the hyperparameters non-linearly by using the L-BFGS-B routine[a] (Byrd et al. 1995; Zhu et al. 1997) implemented in SCIPY (Virtanen et al. 2020) to maximise the log-likelihood of each lightcurve. The mean of the GP then determines the best-fit model used to detrend each target.

To generate the SSO lightcurves for this chapter, I stitched together each target's series of nightly, normalised lightcurves. These lightcurves are better optimised for the atmospheric conditions of the night and, as such, are less affected by systematics than the global lightcurve. As a GP will detrend these lightcurves, any offsets between nights will be accounted for. Before performing GP regression, these lightcurves were binned every 10 minutes and sigma-clipped twice to mask any residual flares and large transits. In essence, binning and sigma-clipping have the same effect; they remove short-duration structures to prevent over-fitting with the GP. The first sigma-clip was a basic 3σ clip of the entire lightcurve, to remove the most extreme outliers. I also included a second, "running" sigma-clip due to a number of targets that exhibited rapid photometric variability. Here, I binned the lightcurve into hours, longer than a typical transit duration, to produce a running median (a series of points representing the median of each hour bin) and a running standard deviation (defined similarly). I then took the median running standard deviation, σ, to reduce the significance of those containing transits or remaining flares. This value of σ is then used to clip the values more than 3σ above or below the running median. This method provides a reliable sigma-clip method for variable stars. All the nightly lightcurves were then combined together, rather than GP-detrending each night individually. In doing so, I give the GP more information about the periodicity (or long-term behaviour) and lower the significance of any infrequent, short features. Each step of this method is shown in Figure 5.1. The power of this method is also demonstrated in Figure 5.10, where I recovered a small transit injected into the lightcurve of a rapidly varying target.

5.2 Searching for Transits Hidden in the SSO Lightcurves

The target differential lightcurves produced by the SSO Pipeline for each night are viewed on an online interface, called the SPECULOOS Portal (designed by P. P. Pedersen). The Portal is used by the entire SPECULOOS team to visually inspect and interact with the data, and

[a]L-BFGS-B is based on the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Fletcher 1987). BFGS uses descent direction to iterate towards a local minimum of a function.

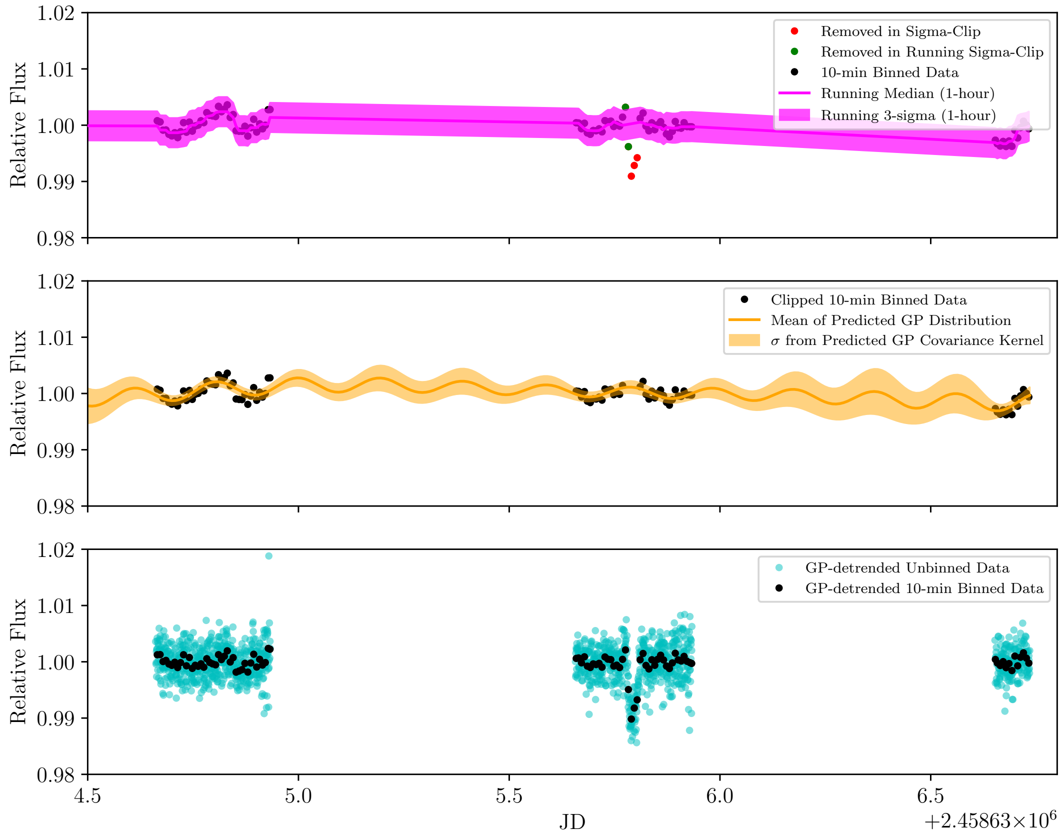

Figure 5.1: Figure showing the main stages of the GP-detrending process on a subsection of a lightcurve with a planet injected. **Top:** Black data points are the 10-min binned lightcurve data. Red points are those removed during a basic 3-sigma clip. The pink line shows the running median (using 1-hour bins) and the shaded pink region demonstrates the 3-sigma threshold for the running sigma-clip. The green points are those removed during this second (running) sigma-clip. **Middle:** The sigma-clipped binned data for the GP-fit is shown in black. The mean of the predicted GP distribution is shown by the orange line. The standard deviation, σ, from the diagonal of the predicted GP covariance kernel is represented by the shaded orange region. **Bottom:** The final GP-detrended unbinned (cyan) and 10-min binned data (black).

to manually flag interesting features and potential transits. In this way, all the differential lightcurves produced by the SSO Pipeline should have already been vetted for transit shapes. However, due to the large number of lightcurves (a combined total from all the SPECULOOS facilities of 6-12 each night), underlying photometric variability, residual systematics and lack of a clear global overview, some transits may have been missed.

To check for any potential planets hiding in the data I ran an automatic transit search through each SSO lightcurve (from the SSO data sample defined in Section 3.1.1). This transit-search pipeline involves GP-detrending clean ('clean' as defined in Section 3.1.2) lightcurves and applying a transit-finding algorithm. For this algorithm I used the ASTROPY implementation of Box Least Squares, or BLS (Kovács et al. 2002), described in the following section.

5.2.1 Box Least Squares

Box Least Squares (BLS) is an algorithm that searches for periodic box-like signals in a data series. It does so by folding the time series on a range of periods (determined by the user), binning and fitting a box. A box is defined as a periodic, discrete change in value between two levels (equivalent to a step function), with most of the data points at the higher level. Data points in the folded lightcurve are down-weighted based on their variance. By maximising the log-likelihood of the transit model over a grid in transit duration and epoch (for every period that is checked), BLS produces a spectrum of log-likelihood against period. The most likely period determined by BLS is the peak of this spectrum.

This simple model does not consider the rounding of a transit from limb-darkening or gradual ingress/egress phases. However, these phases are short for short-period planets around UCDs and will not seriously affect the detection. These planetary factors should be taken into account when modelling the specifics of the transit; however, for transit detection, a box shape is sufficient. It is well-known that approximating a transit by a box shape with BLS will lead to underestimating the transit depth. This can be remedied by using a transit-like shape, such as with Transit Least Squares (Hippke & Heller 2019), but with some computational cost.

5.2.2 Applying BLS to SPECULOOS

BLS is a tool designed to detect transits based on their periodicity. While BLS can detect single transits of sufficient SNR (in this case, the folded lightcurve would have limited data in-transit and significantly more data out-of-transit), its main purpose is to identify low-SNR signals that would not be detected with a single transit. Therefore, given the observation durations within the SSO sample (most targets observed for less than 80 hours) it might seem preferable to use an algorithm designed to catch single transits. However, I note that it is only

Program 1 of SPECULOOS's target list (365 targets apart from TRAPPIST-1) where we focus on finding planets within the habitable zone. These targets would have habitable zones of 8–10 days (Sebastian et al. 2021). The observing strategy for the remaining 1292 targets has been optimised to detect short period planets, such as TRAPPIST-1b, where we would detect multiple transits, and for which BLS would be well suited.

The ASTROPY implementation of BLS calculates a period grid based on user inputs of the minimum and maximum period to search. For the SPECULOOS lightcurves, I limited BLS to search for 10,000 periods between 0.5 and 10 days. I specified a fixed transit duration of 0.02 d. Although BLS allows for multiple transit durations to be searched, this decision significantly improved the speed of the BLS algorithm. I chose 0.02 d because Earth-sized (0.5–2$R_\oplus$) planets with periods of 1–4 days have typical transit durations of ~0.01–0.03 d. BLS then folds the lightcurve on the 10,000 periods in the resulting period grid and searches for a transit on each phase-folded lightcurve, returning the corresponding log-likelihoods. BLS also provides the transit depth and phase (where in the lightcurve the transit occurs) for the most likely period.

Results

I applied the transit-search pipeline to the 154 target lightcurves in the SSO data sample described in Section 3.1.1, binned every 5 minutes. With no constraints on the output, except a minimum log-likelihood of 3.5, BLS flagged 73 targets as having at least one transit-like feature. However, upon visual inspection it was clear that most of these "transits" had depths similar to the photometric scatter. I, therefore, included a threshold on signal-to-noise ratio, to isolate only the most promising candidates:

$$\mathrm{SNR} = \frac{\delta\sqrt{N}}{\sigma} \tag{5.3}$$

where δ is the transit depth measured by BLS, N is the number of data points within the transit (also provided by BLS) and σ is the average running standard deviation of the binned lightcurves, defined in Section 5.1.3. Signals with SNR < 6.5 were removed, as below this value I could not identify the transit above the noise level in the lightcurve. It is worth mentioning that using the transit depth provided by BLS is not optimal, as described above. It is likely that I will underestimate the SNRs and miss transits that mistakenly fall below the threshold. The inclusion of this threshold left 19 planetary candidates. However, the majority (11) of these targets only had partial transit features at the start or end of nights of observation. At these times the atmospheric conditions are often the least favourable; therefore, I decided not to trust any potential signals for which we had observed less than half of each possible transit. After removing the partial transits, there were 8 planet candidates remaining.

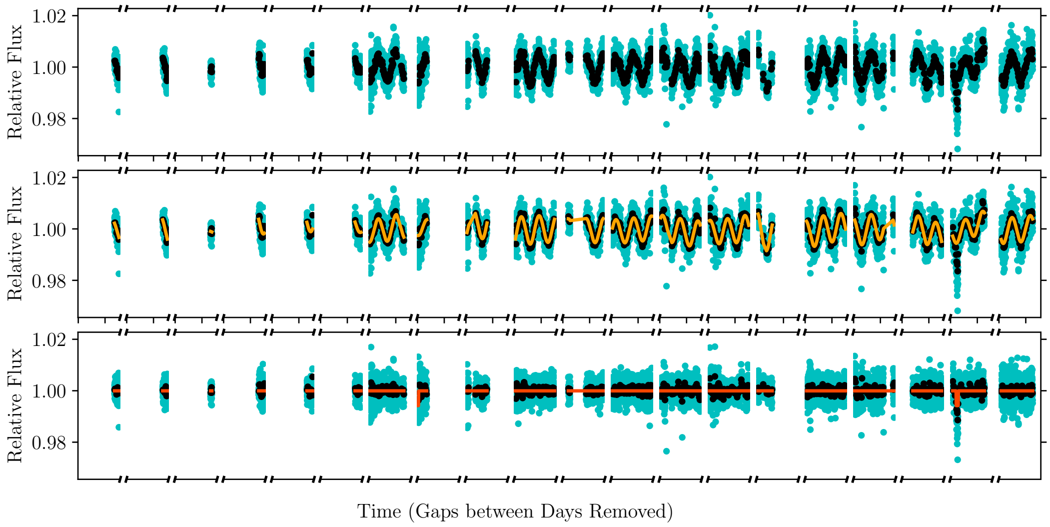

Figure 5.2: In this figure, I show the first of two transit candidates detected in the SSO Data Sample. For each row, the panels represent different days, with the time between them removed for clarity. I show the lightcurve in the top row, unbinned in cyan and binned every 5 minutes in black. I present the same data with the mean GP model, represented by the orange line, in the middle row. The bottom row contains the lightcurve, now detrended by the GP model, to which I have applied the BLS transit-search algorithm, with the BLS model shown in red. The transit candidate detected by BLS is shown in the second last panel, clearly below the level of photometric variability.

With SPECULOOS, the data sample is small enough that there can be manual vetting of transit features. Visual inspection and flagging of nightly lightcurves for transit-like structures is the current preferred method for planet detection in the SPECULOOS consortium. I considered each of these 8 lightcurves in turn. One appeared to be the result of a bad water vapour correction, from spurious LHATPRO data. Three more had structures that correlated strongly with the FWHM (possibly from a crowded field, defocused observations or another nearby star). A fifth lightcurve was on the boundary to be flagged as saturated, so was most likely in the non-linear regime (saturation leads to rapid drops and changes in flux). The sixth lightcurve was extremely active with irregular variability, and so was poorly modelled by the GP. Therefore, I determined that these six candidates were not convincing. Finally, only two signals remained, both consistent with changing aperture size.

While these signals are promising, it is important to note that neither of these potential transits look like those of the TRAPPIST-1 planets. In other words, neither of the two structures is a symmetric, deep, box-like transit. There are also concerns to note about both features. The first feature (Figures 5.2 and 5.3) occurs in a lightcurve that exhibits rapid photometric variability. While the transit structure is clearly outside the level of this variability (which is

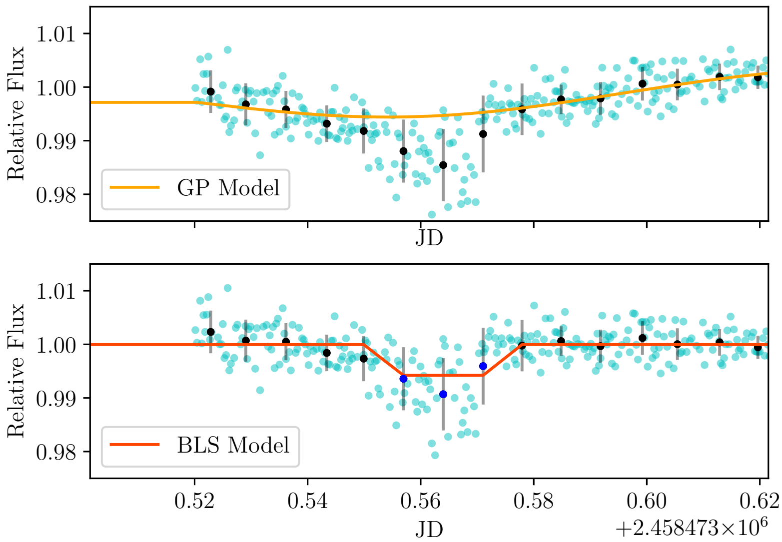

Figure 5.3: The first transit candidate recovered by BLS in the SSO data sample. The cyan points are the unbinned data, and the black points are binned every 5 minutes. In the top plot, I show the differential lightcurve with the GP model in orange. The lightcurve is GP-detrended in the bottom plot, and the BLS transit model is over-plotted in red. The dark blue points are those that BLS determines are in-transit. The depth of this transit structure is significant, however, at this time the photometric scatter is also enhanced.

well-modelled by the GP, as seen in Figure 5.2), there is an increase in photometric scatter at the same time (as seen in Figure 5.3). Despite these caveats, this feature remains the most promising candidate transit in the SSO data sample. The second feature (Figures 5.4 and 5.5) occurs on a night where DONUTS, our auto-guiding software, was not running, therefore we have positional drifts over the course of the night. Therefore, this shallow feature, which also occurs close to the end of a night, could be the result of non-homogeneous pixel response as the target moves position on the CCD window. For this target a second "transit" was detected (seen in Figure 5.5), however it was most likely the result of a bad water vapour correction, from spurious LHATPRO data. Therefore, in both cases, only one clear full transit-like feature was detected. This means that we have essentially no information on their potential periods. And so, although both of these structures were also flagged as transits during manual transit searches by members of the SPECULOOS team, there is no clear consensus as to their nature from only a single event.

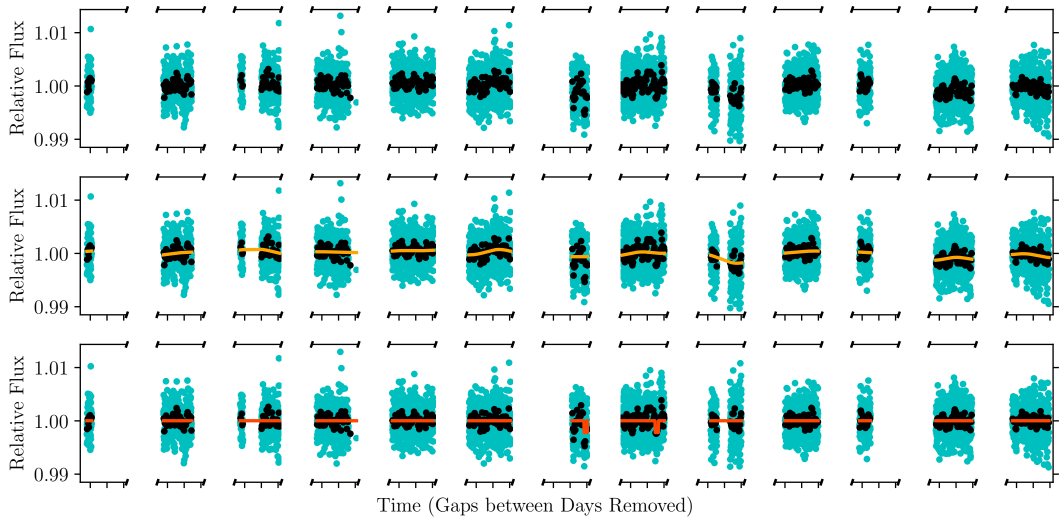

Figure 5.4: Similar to Figure 5.2, I present the second of two transit candidates in the SSO data sample. In this figure, we see two transit structures detected by BLS (on the 7th and 8th days of observation). However, the second (8th panel from left) is caused by the PWV-correction using spurious LHATPRO data and disappears when the correction is removed.

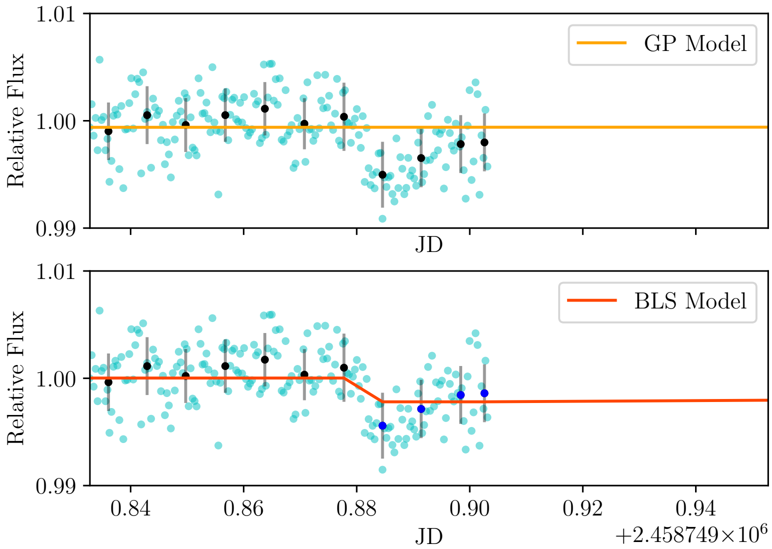

Figure 5.5: Similar to Figure 5.3, I present the second transit candidate recovered by BLS in the SSO data sample. This transit structure is not well-modelled by BLS due to the decision to use a fixed transit duration of 0.02 d.

Testing the Transit-Search Pipeline on Transits from Real Planets

Though not included in the SSO data sample, there are a couple of targets which were observed for TESS follow-up that have known transiting planets. These objects make excellent test cases for the transit-search pipeline. All known planets in the systems around these TESS follow-up targets were successfully recovered.

The first TESS target is LP 791-18, or TOI-736. Crossfield et al. (2019) discovered a transiting super-Earth (on a 0.95 day orbit) and sub-Neptune (on a 5 day orbit) around this red dwarf with spectral type M6. Using BLS I recovered two transits (shown in Figure 5.6a and b) of the larger, $2.3R_{\oplus}$-sized sub-Neptune planet, with a period of 2.5 d, an alias of the real period of 5 d, and a transit depth of 1.6% (measured transit depth with TESS was 1.7%). By masking the detected transits, I re-ran the transit-search pipeline to see if I could recover the smaller, super-Earth planet. Here, I only applied the transit-search method iteratively to demonstrate its capability to detect small planets; it is not a part of the automatic algorithm. On a second run I detected the smaller planet with a period of 0.95 d (which agrees with the measured P=0.95 d) and 0.27% transit depth (measured 0.4%), shown in Figure 5.6c. I detected no other significant transits.

The second TESS target was LHS 3844, which was discovered to host a "hot Earth"; a $1.3R_{\oplus}$ planet with an ultra-short period of 11 hours (Vanderspek et al. 2019). This planet was one of the first to be discovered by TESS. This M5-type target was observed with three SSO telescopes, and transits from LHS 3844b were detected in all three lightcurves. BLS recovered six transits of this planet, three of which are demonstrated in Figure 5.7.

I did not include TRAPPIST-1 in the original sample (even though it is part of the SPECULOOS target list) as the observations with SPECULOOS are scheduled very differently to our normal survey mode. The purpose of observing TRAPPIST-1 with the SSO is to monitor the transit depths of the TRAPPIST-1 planets over time. Therefore, we are predominantly observing in-transit data, which risks being detrended out by the GP. Despite this potential difficulty, I tested the transit-search algorithm on the lightcurves I had for each of the four telescopes. With Callisto, we observed two nights, however, on one night there was a blended transit of TRAPPIST-1b, g and d. The transit-search algorithm recovered planet b on both nights. For Europa's lightcurve I recovered one transit of planet b. Ganymede had one night of observation of planet g, which was also successfully recovered.

The most interesting case, however, was the lightcurve for Io. This lightcurve contained 79 nights of observation, and transit observations of all of the TRAPPIST-1 planets. Despite the challenges of applying GP regression to this particular lightcurve, I was able to recover 28 transits from planet b with a period of 1.5109 d and a transit depth of 0.72%. From Ducrot

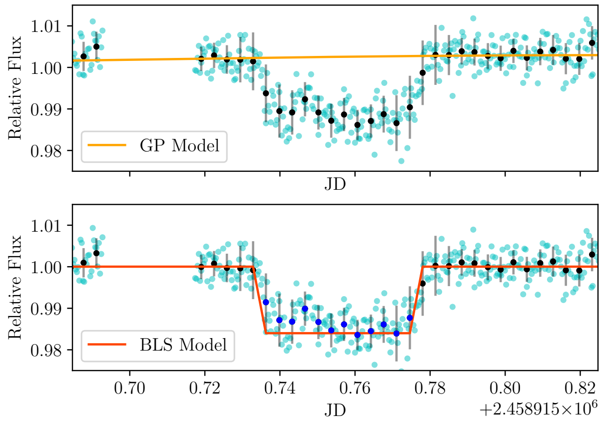

(a)

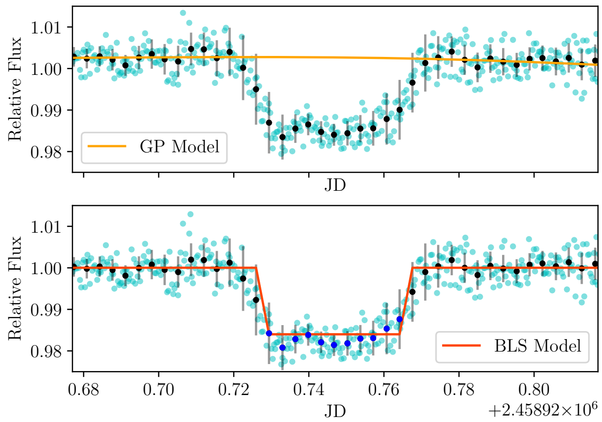

(b)

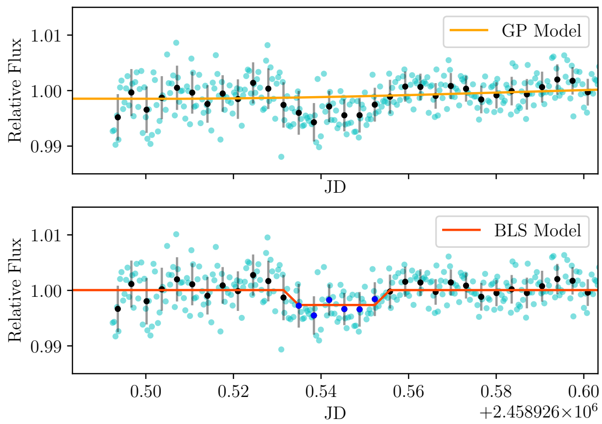

(c)

Figure 5.6: Two transits, (a) and (b), from LP 791-18c, and one from LP 791-18b (c), observed with SPECULOOS-South and recovered using BLS. BLS recovered a period of 2.5 d for planet b and 0.95 d for c. The cyan points are the unbinned data and the black points are binned every 5 minutes. The top plots of (a), (b) and (c) show the differential lightcurves with the GP model in orange. The bottom plots show the GP-detrended lightcurve with the BLS transit model in red. The dark blue points are those that BLS determines are in-transit.

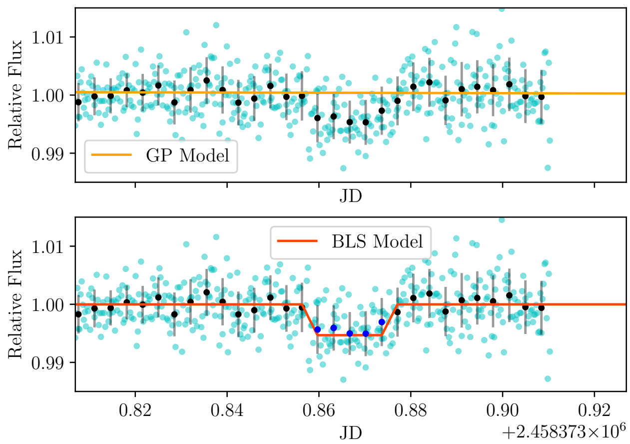

(a)

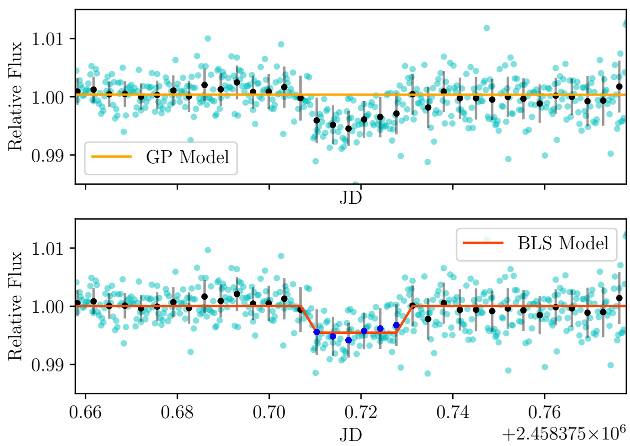

(b)

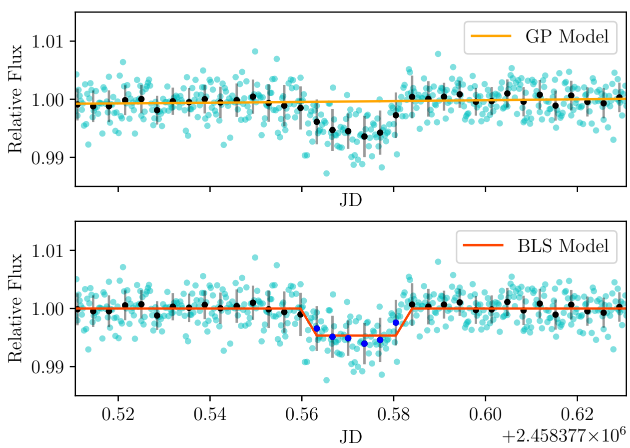

(c)

Figure 5.7: Similar to Figure 5.6, this figure shows three transits from LHS 3844b, all taken with different telescopes: (a) Callisto, (b) Io, and (c) Europa.

et al. (2020) the period of TRAPPIST-1b was measured as 1.5109 d and its transit depth is 0.72%. Iteratively masking the detected transits and re-running my transit-search algorithm, I recovered (in decreasing BLS likelihood order):

1) TRAPPIST-1c, from 8 transits, with P=8.476 d and a depth of 0.77%. The measured period is 2.422 d and transit depth is 0.7% (Ducrot et al. 2020).
2) A "planet" with P=1.557 d, however this combined one transit each from planets c, e and g.
3) TRAPPIST-1f with a period of 4.603 d, half its measured period of 9.207 d (Ducrot et al. 2020).

Continuing to perform iterative masking of transits and transit-searching recovered only more false "planets" that combined transits from multiple TRAPPIST-1 planets. In conclusion, the automatic transit detection algorithm recovered multiple transits from planets b, c and f successfully. This code also picked up individual transits (and attempts to combine them incorrectly) from planets d, e and g, but not h. It is not surprising that I did not recover planet h as it has the smallest transit depth and longest period of all the TRAPPIST-1 planets. It is worth mentioning that the recovered depths (and corresponding SNR values) are not representative of these type of planets in lightcurves that were observed in survey mode, as GP detrending was seen to warp the transits.

5.3 Transit Detection Efficiency

5.3.1 Injection-Recovery Tests

To draw any conclusions about the planetary occurrence rates around ultra-cool dwarfs we must first understand the transit detection efficiency of SPECULOOS. The detection efficiency tells us: if there was a transiting planet around a particular target how likely would we be to detect it, given its parameters? The most intuitive (and accurate) way to determine our transit detection efficiency is through injecting synthetic transits into real lightcurves from the SSO and assessing the recovery using the transit-search pipeline (as has been done in Giacobbe et al. 2012; Petigura et al. 2013; Demory et al. 2016b; Sestovic & Demory 2020; Lienhard et al. 2020 etc.).

I performed transit-injection tests on 154 SPECULOOS targets from the SSO data sample. Of these targets, 33 have been observed with more than one telescope (non-simultaneously). I did not combine separate lightcurves together prior to GP-detrending, as different telescopes experience different systematics. Therefore, I injected 1000 planets into each telescope's lightcurve and recovered them independently. This decision may have resulted in 2000, 3000 or 4000 planets injected for these 33 targets, however, it will not impact the detection efficiency. I generated transits for the 1000 artificial planets for each lightcurve using PyTransit (Parviainen 2015), a transit lightcurve modelling package that implements the Mandel & Agol (2002) model. I supplied the limb-darkening coefficients from TRAPPIST-1b (Ducrot et al. 2020). For each of the planets injected, their parameters (radius R_{p}, period P, and inclination i) were drawn from the following uniform distributions:

$$\begin{aligned} \frac{R_{\mathrm{p}}}{R_{\oplus}} &\sim U(0.5, 6.0) \\ \cos i &\sim U(\cos i_{\mathrm{min}}, \cos i_{\mathrm{max}}) \\ \log \frac{P}{\mathrm{days}} &\sim U(\log 0.5, \log 10.0) \end{aligned} \tag{5.4}$$

where $U(a, b)$ represents a value drawn from a uniform distribution between a and b. i_{min} and i_{max} are the minimum and maximum inclinations for a transiting planet. From Equation 1.9 we know the condition on the inclination for a transit is dependent (in the case of a circular orbit where $R_{\mathrm{p}} << R_*$) on the radius of the host object and the semi-major axis. From Kepler's third law, the semi-major axis itself is a function of the host mass and the orbital period (Equation 1.21). The host mass and radius were assumed constant, taken from the SPECULOOS target list (Sebastian et al. 2021). However, the inclination limits depend on the orbital period; therefore, when drawing the planetary parameters, I drew individual inclinations from a range set by each period. I only considered circular orbits (with e=0) as I do not know the underlying eccentricity distribution. Close-in planets are also expected to experience significant circularisation of their orbits (Luger et al. 2017). The time at which the first transit was injected was also randomly drawn from a uniform distribution, $\phi \sim U(0, 1)$, where ϕ is the phase of the period, such that the first transit was injected at ϕP from the start of observations.

I injected artificial planets in turn into each SPECULOOS target's differential lightcurve(s) and ran the transit-search pipeline. The planets were injected prior to the GP-detrending to allow for the possibility that the GP could over-model and remove, or warp, the transit signals. If I were to inject planets after GP-detrending this could artificially inflate the recovery results, as it is easier to detect a transit in a 'flat' lightcurve than in a lightcurve exhibiting time-dependent structure. From BLS I only considered the most likely transit, I did not look at other peaks in the periodogram. I determined that a transit was recovered successfully if at least one transit was detected by BLS above the SNR threshold (6.5, described in Section 5.2.2), with the detected transit region overlapping with at least half of the real transit. I note that due to the decision to use one transit duration of 0.02 d, the recovery of planets with transit durations > 0.04 d will likely be underestimated. These durations correspond to planets with long periods, which are inherently much less likely to be recovered by a ground-based survey. I chose not to include conditions on recovering the period of the injected transit (as was done in Giacobbe et al. 2012;

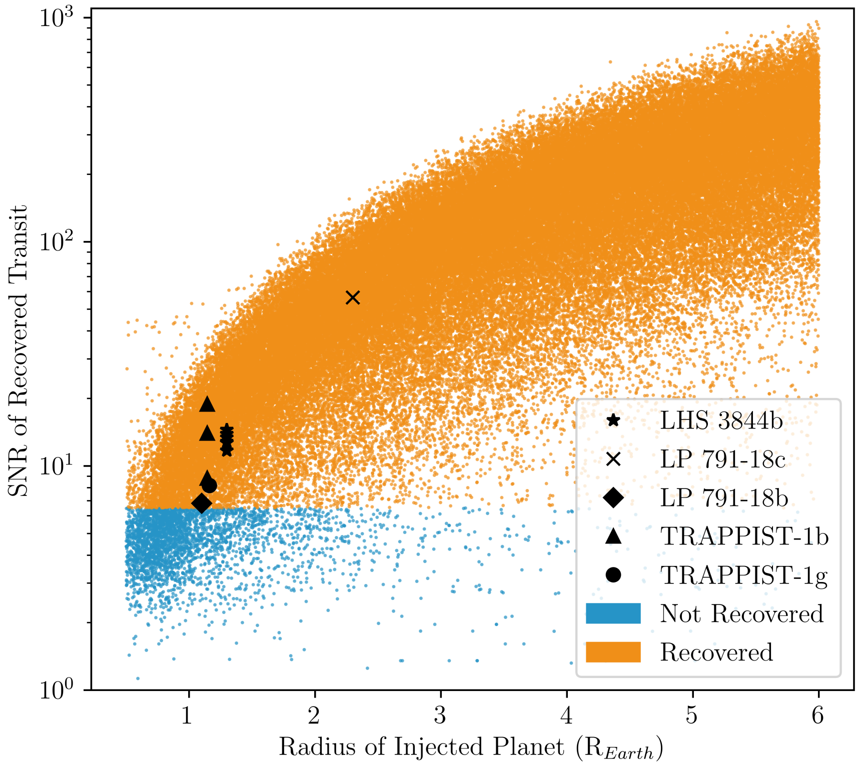

Figure 5.8: The measured signal-to-noise ratios (SNR) of recovered planets, highlighting the SNR threshold of 6.5. The planets with SNR ≥ 6.5 (and so recovered) are shown in orange, and the transits with SNR < 6.5 (therefore not recovered) are shown in blue. Over-plotted are the measured SNR values for real transit detections of LHS 3844b, LP 791-18b and c, TRAPPIST-1b (measured by three different telescopes) and TRAPPIST-1g. If they had existed in the SSO sample, then all five of these planets would have been successfully recovered, according to the criteria set in the injection-recovery tests.

Petigura et al. 2013; Demory et al. 2016b; Sestovic & Demory 2020). Due to the nature of ground-based observing there are large gaps in the time series, often meaning that we recover an alias of the true period. Additionally, I allowed for the detection of single transits, which have very poor constraints on their period and would certainly be followed up with targeted observations. I included the SNR > 6.5 threshold, the justification for which is discussed in Section 5.2.2, as a requirement for recovery, the impact of which is shown in Figure 5.8. Several other works have also used a detection significance above a certain threshold for single transit recovery (7.5 in Berta et al. 2013; 3.5 in He et al. 2017). I decided not to implement a transit depth threshold (Sestovic & Demory 2020; Lienhard et al. 2020). Most recovered transits have depths at least half of the true transit depth. However, because I have limited the duration to 0.02 days, a box of this length does not always model well the transit signals with much longer or shorter durations (>5 d or <1 d periods), leading to an incorrect depth measurement. I did not consider this a concern, as my aim for these injection-recovery tests was to flag a transit signal, not to model transits accurately. As SPECULOOS is a targeted survey, all transit signals will be followed up by manual validation, and at this later stage the transit signal can be modelled. This manual process was not included in the tests, due to the sheer number of injected transits. I note that this may have artificially increased our detection efficiency, particularly for small planets whose shallow transit depths are more difficult to visually identify above the photometric scatter.

The injection of a synthetic transit and application of the transit-search pipeline is demonstrated in Figures 5.9 and 5.10. In particular, the three injections in Figure 5.10 highlight the strength of the double sigma-clip method (described in Section 5.1.3) for targets with rapidly varying brightness. In both cases, the planet was successfully recovered, even where only one transit was observed. For the 195,000 planets injected, I binned them in R_{p} (0.5–6 $\mathrm{R}_\oplus$ in 0.25 $\mathrm{R}_\oplus$ bins) and $\log P$ (0.5–0.91 d, 0.91–1.66 d, 1.66–3.02 d, 3.02–5.49 d, 5.49–10 d) space and extracted average recovery fractions for each bin. This allowed me to see general trends in detection efficiency as I considered planets of various sizes and orbital periods. The recovery fractions, as a function of planet radius and period, are shown in the first plot in Figure 5.11.

5.3.2 The Probability of Detecting a Planet

There are three factors to consider when determining the probability of detecting a planet. The first is the probability that the planet exists, $\mathbb{P}_{\mathrm{exist}}$. The second is the probability that a planet transits, as seen from Earth, $\mathbb{P}_{\mathrm{transit}}$. If the geometry is not favourable for a transit, then a transit survey would never be able to detect it. The third is the probability that a transiting planet is detected in a lightcurve, $\mathbb{P}_{\mathrm{detect}}$. This last probability follows directly from the detection efficiency in the previous Section 5.3.1. With these last two factors I define the completeness, $H = \mathbb{P}_{\mathrm{transit}} \cdot \mathbb{P}_{\mathrm{detect}}$, which determines how likely I would be to find a planet if it exists at a random alignment. $\mathbb{P}_{\mathrm{detect}}$ and the completeness are both presented in Figure 5.11. The total probability of finding a planet (with radius R_{p} and period P) for a particular host object (with radius R_* and mass M_*) is, therefore:

$$\mathbb{P}_{\mathrm{find}} = \mathbb{P}_{\mathrm{exist}}(R_{\mathrm{p}}, P) \cdot \mathbb{P}_{\mathrm{transit}}(R_*, M_*, P) \cdot \mathbb{P}_{\mathrm{detect}}(R_*, M_*, R_{\mathrm{p}}, P) \tag{5.5}$$

$\mathbb{P}_{\mathrm{transit}}$ and $\mathbb{P}_{\mathrm{detect}}$ are dependant on the host object, whereas, crucially, the probability that the planet exists in the first place depends on the planet occurrence rates for the types of hosts we are considering. If I assume that this occurrence rate is constant among the hosts in my sample, then $\mathbb{P}_{\mathrm{exist}}$ is independent of the stellar parameters. From Equation 1.11 we know that the probability for a system to transit is $\sim R_*/a$.

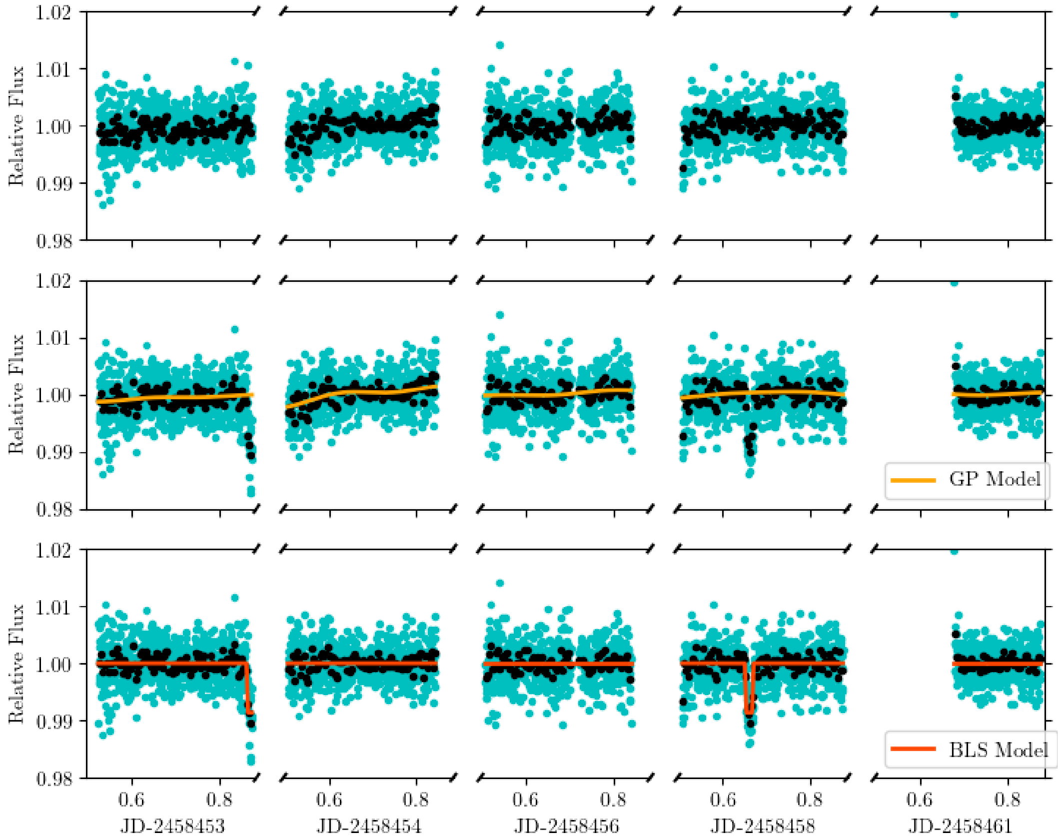

Figure 5.9: **Top**: Cleaned lightcurve of five nights of observation for an M5-type object with no detected periodicity. The cyan points are unbinned data and the black points are binned every 5 minutes. The time between nights of observation has been removed to improve visibility. **Middle**: The same lightcurve, however, with an injected transiting planet of R_p=1.3$R_\oplus$ and P=1.2 d. Two transits are visible for this planet, a partial transit at the end of the first night (SNR=15.3) and a full transit on the fourth night (SNR=19.8). The orange line is the best-fit GP using a *squared exponential* kernel. The GP has been fit with binned, sigma-clipped data to reduce the chance of modelling the transit. **Bottom**: The lightcurve detrended with the GP model. BLS was then run on the flattened lightcurve and the red line indicates the most likely planet it found, with a recovered period of 1.2 d (=injected period).

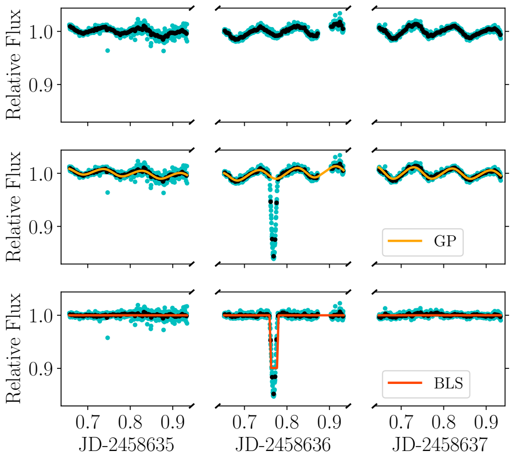

(a) R=5.98R$_\oplus$, P=0.61 d

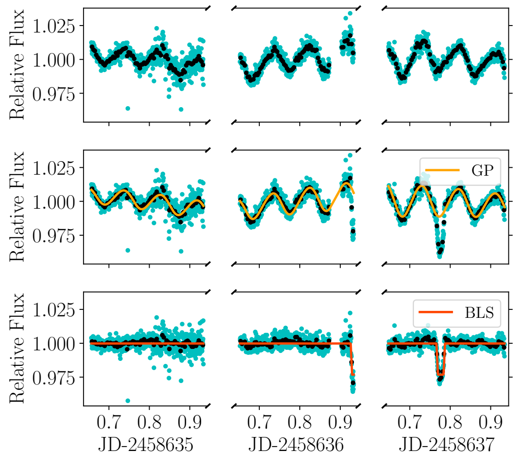

(b) R=2.12R$_\oplus$, P=0.48 d

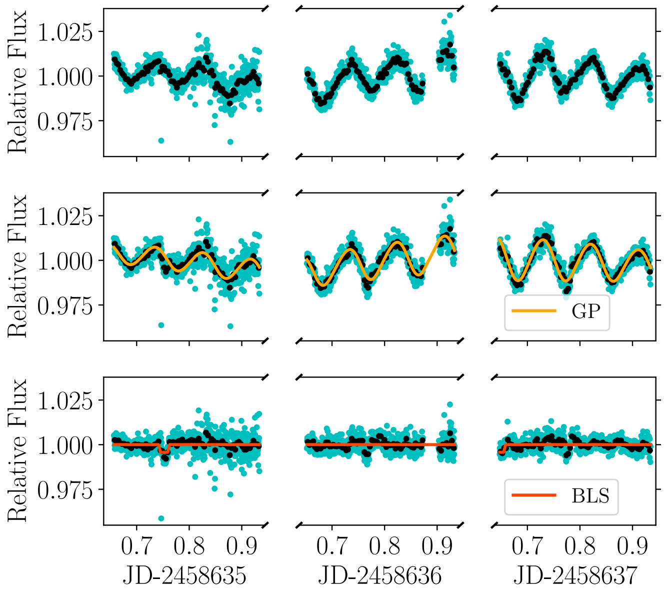

(c) R=1.39R$_\oplus$, P=1.01 d

Figure 5.10: Similar to Figure 5.9. Lightcurve of three nights of observation for an object with a spectral type of M5, demonstrating photometric variability with a rotation period of ~0.09 d. Three different planet injection scenarios are shown in (a), (b) and (c) for planets with radii 5.98, 2.12 and 1.39 R$_\oplus$ and periods 0.61, 0.48 and 1.01 d respectively. The orange line in each middle plot is the best-fit GP using a *quasi-periodic* kernel. The GP has been fit with binned, sigma-clipped data to reduce the chance of over-modelling transits with depths (a) much larger than, (b) comparable to, and (c) much smaller than the amplitude of photometric variability. The bottom plot in each case shows the lightcurve detrended with the GP model and searched for transits using BLS. The transits were successfully recovered in all three cases with SNR (a) 170, (b) 41, and (c) 7.2. However, in (c) only one transit was detected (at JD 2458635.75), two were missed (at JD 2458636.76 and JD 2458637.77), and a false positive transit was detected on the third day due to a small flare (though it would be rejected with a SNR of 5.6). In (c) we are very close to the detection limit for this target.

5.3.3 Results of the Injection-Recovery Tests

Assuming there is no chance of false positive detections, then there were three possible outcomes from the injection-recovery tests:

- the injected planet transits during a night of observation and was recovered,
- the injected planet transits during a night of observation and was not recovered,
- or the injected planet did not transit during a night of observation and could not be recovered.

Of all the injected planets: 48.5% were detected, 7.3% were not detected, and 44.2% fell into the final category and were not observed. A planet was 'not observed' if less than half of one transit occurred during a night of observation. The ratio between the first two outcomes reflects the quality of the transit-search pipeline. 87% of the planets that had at least half of one transit observed by SPECULOOS were recovered. If a transit was detrended by the GP, or was very shallow, then, even when it was observed, the planet may not have been recovered. Therefore, as a sense check, I examined the recovery rates of only the planets which were observed. I found that I recovered ~95–97% of large planets ($>4R_\oplus$), which dropped rapidly for small planets ($<1R_\oplus$), due to the difficulty in recovering them above the noise level. There was also no decrease in detection efficiency with increasing period which justified the choice of BLS even for single transit cases (see the discussion at the start of Section 5.2.2). The final outcome informs us about the limitations imposed by gaps in our data. This can be very useful in assessing our observing strategy moving forward.

The results of the injection-recovery tests are shown explicitly in Figure 5.12 (binned in Figure 5.11a and smoothed in Figure 5.13). The recovery fraction decreases as orbital period increases and as the radius of the injected planet decreases (particularly at radii $< 1.2\,R_\oplus$). These results show that I am unlikely to recover any planets smaller than $\sim 0.8\,R_\oplus$ or with periods longer than ~5 d (see Figure 5.13). Long period planets, which transit less frequently, are more likely to be missed by a ground-based survey with large data gaps. The geometric probability of transit also drops sharply for planets with longer periods. Small planets, with shallow transit depths, are more difficult to detect above the level of photometric scatter within the lightcurves. As expected, there were lines of non-detections for periods of integer numbers of days (1, 2, 3 days etc.) or for $(n + 0.5)$ days (1.5, 2.5, 3.5 days etc.), where planet transits consistently fell into the gaps between nights of observation (Figures 5.12 and 5.13).

5.4 The Occurrence Rate of Planets Like TRAPPIST-1b

For a given target in the SSO sample I assumed that the probability of detecting a planet could be described by a Poisson-Binomial distribution (as in Lienhard et al. 2020). The Poisson-Binomial

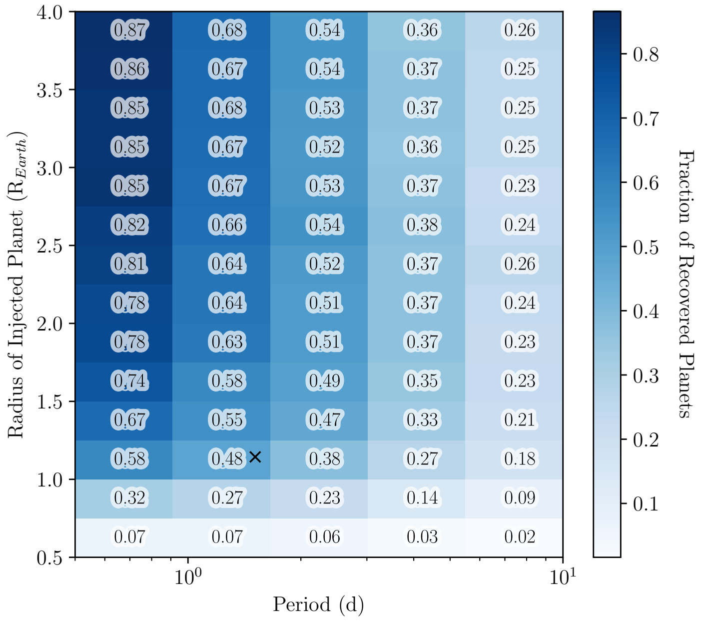

(a) Detection Sensitivity

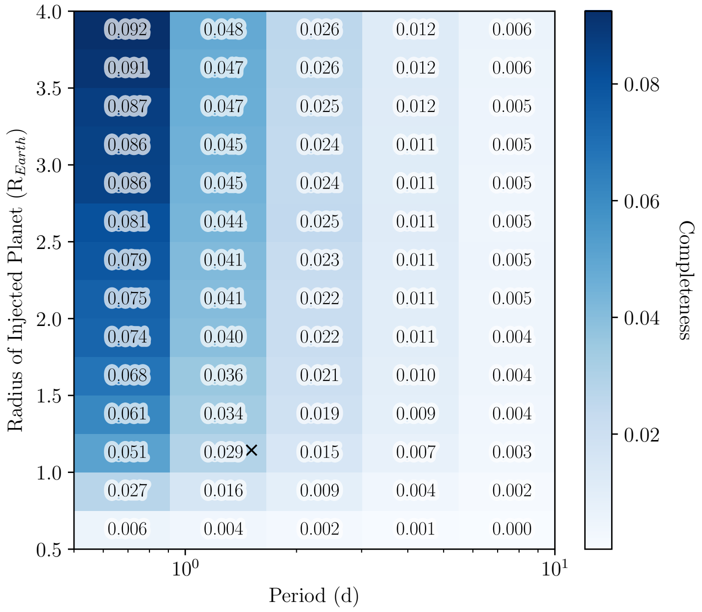

(b) Completeness

Figure 5.11: (a) shows the fraction of recovered planets, assuming that a planet exists and is transiting. The errors in the recovery fraction range from 0.01–0.03. (b) illustrates the completeness, which is the chance of detecting a planet with any orbital geometry. The errors in the completeness range from 0.001–0.007. The parameters of TRAPPIST-1b are marked by a black cross in both plots. If TRAPPIST-1b (R_p=1.144 $R_\oplus$, P=1.5109 d, Ducrot et al. 2020) were orbiting an object in the SSO sample, there is a 2.9% probability that I would have detected it with a random alignment, and a 48% probability that I would have detected it if it were transiting. Both plots are only shown up to radii of 4$R_\oplus$, as for larger *R* the results do not significantly change (the recovery fraction saturates at large radii). The recovery fraction for large, short-period planets never reaches 1 because there are periods where no transits will occur during observations, in some or all lightcurves, due to observation gaps (such as P=0.5 d).

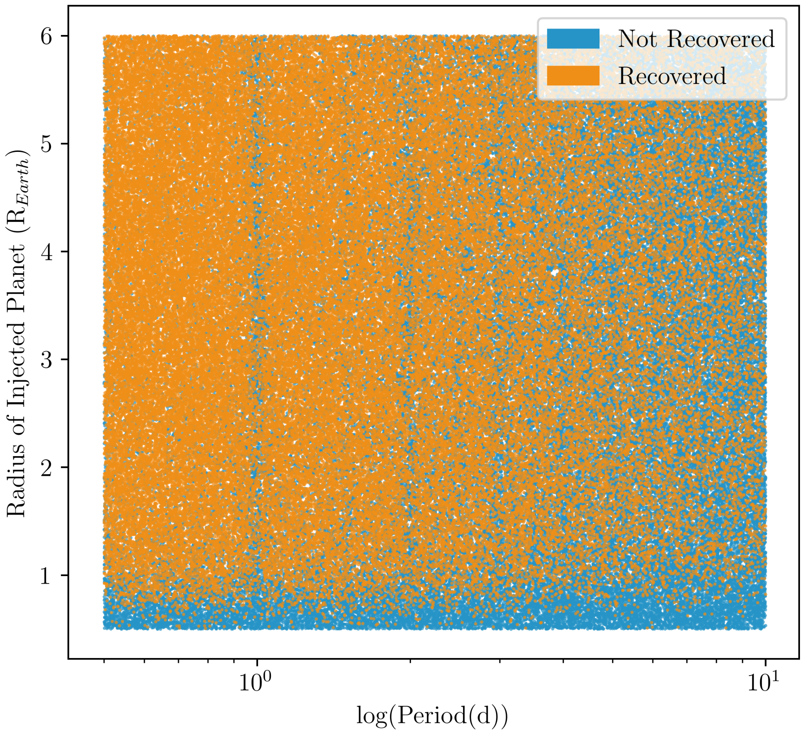

Figure 5.12: The transit-injection results of every planet inserted into the SSO lightcurves. The recovered transits (criteria for recovery described in Section 5.3.1) are shown in orange, and the transits not recovered shown in blue. This is a useful companion to Figure 5.11a, showing finer detail, such as a difficulty to detect planets with periods n or $(n + 0.5)$ days, for $n = 1, 2, 3...$

distribution describes a sum of independent (non-identical) Bernoulli trials, or experiments with a binary outcome. For the detection of planets, I assumed that each lightcurve was Bernoulli trial where the outcome is either a planet was found or was not found. The probability of finding a planet for each lightcurve is different (which separates this distribution from Binomial), given by $\mathbb{P}_{\text{find}}$. The Poisson-Binomial distribution has a mean $\mathbb{E}_{\text{planets}}$, equal to the number of planets we would expect to have detected in the sample. The expected number of detected planets in the data sample is defined as:

$$\mathbb{E}_{\text{planets}}(R_{\text{p}}, P) = \sum_{i=1}^{N_*} \mathbb{P}_{\text{find}}(R_{*,i}, M_{*,i}, R_{\text{p}}, P) \tag{5.6}$$

where N_* is the total number of trials, or target objects. From the transit injection-recovery tests I computed $\mathbb{P}_{\text{detect}}$, and from geometry $\mathbb{P}_{\text{transit}}$ is known for each star-planet combination. Therefore, the completeness could also be derived. By assuming that every star hosts a planet in each of the parameter bins, I calculated $\mathbb{E}_{\text{planets}}$, as shown in Figure 5.14. In Figure 5.14 I have made the assumption that we have one planet around every star in the finite parameter bins I have chosen. In actuality we would have a smooth underlying distribution of occurrence

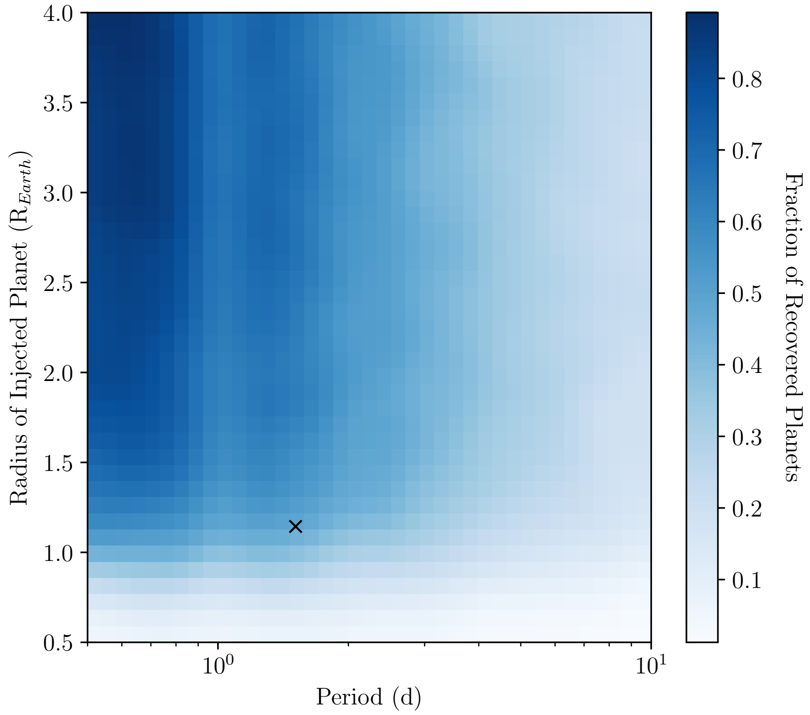

Figure 5.13: Equivalent to Figure 5.12, but binned and smoothed using a Gaussian filter. The recovery fraction of the transit-search drops steeply for planets with radii below $0.8\,R_\oplus$ and periods longer than 3–4 days. TRAPPIST-1b is marked by a black cross.

rates rather than a fixed value. Therefore, the Figure does not represent, for example, that there should be 4.56 planets in the radius-period bin with TRAPPIST-1b and also 8.01 planets of the same radius with shorter periods. This figure is meant to indicate that for a given parameter bin if the underlying occurrence rate was, on average, 1, then we would have found 4.56 planets, and tell us nothing about the surrounding bins. Therefore I recommend the reader not focus on the large steps between adjacent bins (which is clearly unphysical) and consider each parameter bin in isolation. Clearly, it is also unlikely that the occurrence rates for most of these types of planets would be 100%, as then I would have found significant numbers of them (~14 large planets with periods less than 1 d). Therefore, using the actual detections and non-detections in the SSO sample I can place constraints on the occurrence rate, $\mathbb{P}_{\text{exist}}$, of various planet types.

The probability of detecting a particular planet orbiting an object is given by $\mathbb{P}_{\text{find}}$, meaning that the opposite, the chance of not detecting that planet is $1-\mathbb{P}_{\text{find}}$. Using the Poisson-Binomial distribution then it becomes possible, from Wang (1993), to find the probability of a number of successes n (or planet detections) out of a total number of trials (or stars) N_*:

$$\mathbb{P}(X = n) = \sum_{k\in S_n} \prod_{i\in k} p_i \prod_{j\in k^c} (1 - p_j) \tag{5.7}$$

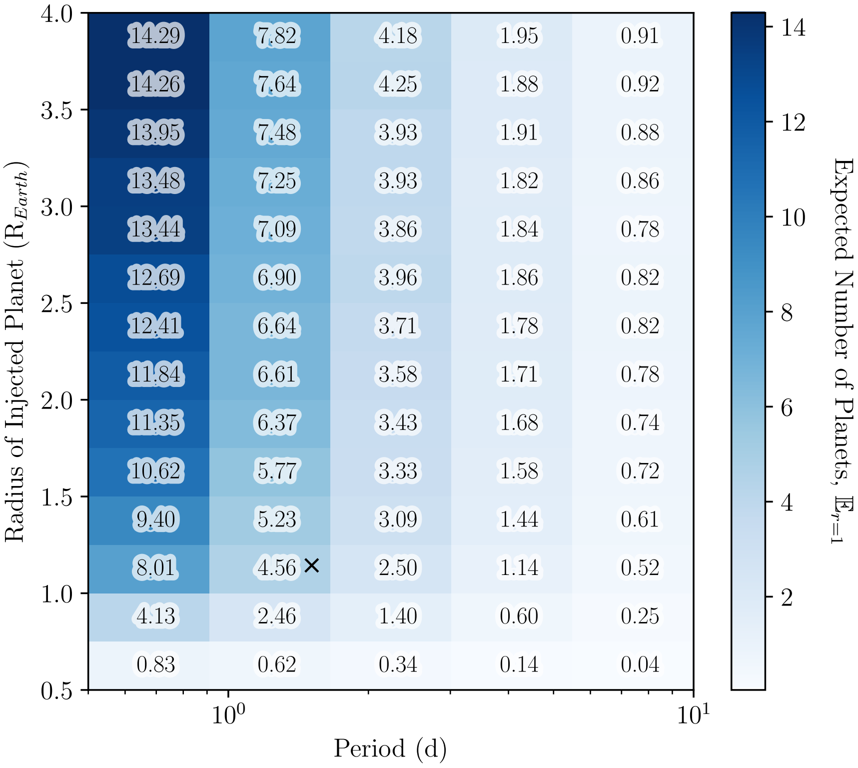

Figure 5.14: The expected number of planets I would have found in the SSO sample, if every star hosted one of the planets in every parameter bin. The expected number is equivalent to the mean of the Poisson-Binomial distribution. TRAPPIST-1b is marked with a black cross. If there were a planet like TRAPPIST-1b (in the same parameter bin) orbiting every observed object, I would have found 4.56 of these planets. Therefore, since I found none it is likely that the percentage of stars hosting these planets is markedly lower than 100%.

where S_n is the set of all subsets of n integers that can be selected from $\{1, 2, 3, ..., N_*\}$ stars and k^c is the complement of k. For example, if $N_* = 3$ and $n = 2$ then $S_n = \{\{1,2\},\{2,3\},\{1,3\}\}$. This is equivalent to the sum over possible outcomes. For each lightcurve $p_i = \mathbb{P}_{\text{find}}$. If the probabilities are all equal ($p_1 = p_2 = ... = p_i$) then this reduces simply to the binomial distribution. Assuming that all the individual probabilities are less than one, we can use the following recursive formula (Chen et al. 1994; Shah 1973):

$$\mathbb{P}(X = n) = \begin{cases} \prod_{i=1}^{N_*}(1 - p_i) & \text{if } n = 0 \\ \frac{1}{n}\sum_{i=1}^{n}(-1)^{i-1} \cdot \mathbb{P}(X = n - i) \cdot \sum_{j=1}^{N_*}\left(\frac{p}{1-p}\right)^i & \text{if } n > 0 \end{cases} \tag{5.8}$$

In the specific case that no planets are found, $n = 0$:

$$\mathbb{P}(X = 0) = \prod_{i=1}^{N_*}(1 - p_i). \tag{5.9}$$

Similarly, for the detection of one planet, $n = 1$:

$$\mathbb{P}(X = 1) = \mathbb{P}(X = 0) \sum_{j=1}^{N_*} \frac{p_j}{(1 - p_j)} \tag{5.10}$$

and for the detection of two planets, $n = 2$:

$$\mathbb{P}(X = 2) = \frac{1}{2} \left(\mathbb{P}(X = 1) \sum_{j=1}^{N_*} \frac{p_j}{(1 - p_j)} - \mathbb{P}(X = 0) \sum_{j=1}^{N_*} \left(\frac{p_j}{(1 - p_j)} \right)^2 \right), \tag{5.11}$$

and so on. This recursive formula quickly becomes unwieldy for large numbers of trials or successes.

5.4.1 The Probability of Finding TRAPPIST-1b

For this section I will consider only the synthetic planets with radii 1-1.25 $R_\oplus$ and periods 0.91–1.66 d. This is equivalent to the parameter bin in Figures 5.11 which contains TRAPPIST-1b (R=1.144$R_\oplus$, P=1.5109 d, Ducrot et al. 2020). I will refer to these planets as "planets like TRAPPIST-1b" or "TRAPPIST-1b analogues", however, this is only based on radius and period, and no other considerations. While I am extracting a specific radius and period bin, the actual recovery fraction, completeness and, ultimately, the number of planets I would expect to find, is likely a smooth distribution (as demonstrated in Figure 5.13). By binning the results I can build a larger dataset for statistical robustness, but also I lose any finer structure within that bin. This trade-off in increasing the number of artificial planets vs. representing a realistic occurrence rate distribution is why I consider only a small parameter space containing TRAPPIST-1b.

My reasoning in isolating this narrow parameter space is to derive the probability of detecting no TRAPPIST-1b analogues in the SSO dataset, assuming a particular occurrence rate. Although there were a couple of promising transit events detected in the SSO lightcurves, if I consider the specific case of TRAPPIST-1b analogues, then I found no similar conclusive planets in the dataset. If every object in the sample hosted a planet like TRAPPIST-1b ($r = 1$), then, using Equation 5.6, I would expect to have detected at least four planets. Therefore, from Equation 5.9, I would have a $\sim$1% chance of not finding a single one of these planets. If, instead, the occurrence rates of TRAPPIST-1b were 0.5, or 0.1, then the probability I would not find one would be 10%, or 63%, respectively.

Within the SSO data sample I found no planets like TRAPPIST-1b. Not detecting any planets only allows me to place upper bounds on the planetary occurrence rates of that type of object (see Equation 5.9). However, the next challenge is whether to include TRAPPIST-1b as a detection. Not only is TRAPPIST-1 a part of the SPECULOOS target list, but also, due to its spectral type and brightness, it was a high priority target for JWST. Therefore, it is highly likely TRAPPIST-1 would have been one of the first 50 targets observed by SPECULOOS in survey mode, had the TRAPPIST UCDTS mini-survey not happened. However, without performing transit injection-recovery tests on the TRAPPIST-1 lightcurve, I cannot calculate its completeness. Therefore, I would need to assume the completeness of TRAPPIST-1 can be taken as the average of all objects in the data sample. Which poses the question: does including or not including TRAPPIST-1 introduce a larger bias? I concluded that TRAPPIST-1 would have been in the SSO data sample had its planetary system not been found. Therefore, it would be contrived to deliberately exclude TRAPPIST-1 in a statistical analysis of planets in the SSO data sample.

Including the detection of TRAPPIST-1b, I detected one planet with size 1.0–1.25 $R_\oplus$ and period 0.91–1.66 d. The recovery fraction is ~48% (for the ~1700 injected planets in that bin) which gives a completeness of 2.9% by including the geometric probability of a transit. From this result, Equation 5.10 was calculated for a range of $\mathbb{P}_{\text{exist}}$, to estimate the most likely occurrence rate of planets like TRAPPIST-1b. These probabilities are shown in Figure 5.15. The value of $\mathbb{P}_{\text{exist}}$ with the highest probability is 0.22. I note that here I have assumed all planets in this parameter bin have the same occurrence rate, and that all targets in my sample have the same chance of hosting one of these planets. I also compared the use of Equation 5.10 against a binomial distribution $\mathbb{P}_{1\text{detection}|\text{r}} = N_* p(1-p)^{N_*-1}$ (where p is the occurrence rate multiplied by the average completeness). Both probability functions produce similar peaks for the most likely occurrence rate. If we continued to observe more targets without another planet detection then this occurrence rate would decrease, and the constraint would tighten. Conversely, if another planet was found (or multiple) soon, then this would increase the occurrence. Therefore 0.22 planets per star only provides the most likely occurrence rate for the current outcome of the SSO sample, and a tentative estimate for UCDs. With $r = 0.22$ the number of expected planets in the TRAPPIST-1b parameter bin is 1.00, which agrees with the one detection we have, of TRAPPIST-1b. However, I cannot place 95% confidence limits on any occurrence rate <1 with only 154 targets. We will need to wait for several more years of observations with SPECULOOS to make more concrete conclusions on the frequency of temperate, terrestrial planets around red dwarfs.

5.4.2 Predictions about Future Detections

By extrapolating from the current results, I can make predictions about what the implications on the planetary occurrence rate would be if we observed 100, 200 or 1000 more stars without detecting another planet. I took the average completeness of our survey and performed the same analysis as above but including an extra 100, 200 or 1000 objects. The results are presented in Figure 5.15. The same result was achieved by instead including 100 completeness values from the current sample (essentially doubling 100 of our targets). Here, I have assumed that future

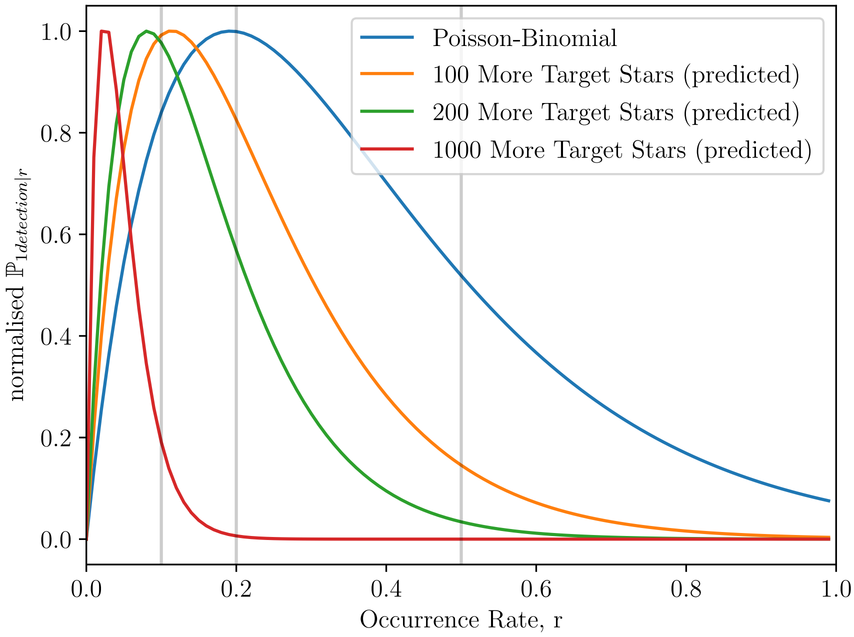

Figure 5.15: The normalised probability (Poisson-Binomial) distribution for the occurrence rate r of planets similar to TRAPPIST-1b given that one (and only one) planet has been detected in the SSO sample, $\mathbb{P}_{\text{1detection}|\text{r}}$. The probability for the current state of the SSO sample is calculated in blue, which assumes one success out of all trials (or that the planet could be orbiting any object in the sample). This distribution is consistent with a peak of $r = 0.22$. The extrapolated results for an extra 100, 200 and 1000 lightcurves without another planet detection are shown in orange, green and red, respectively. Vertical grey lines are plotted for r = 0.1, 0.2 and 0.5.

targets will be very similar to those we have observed so far, and we will continue to observe them in the same way. Due to the priority ranking of SPECULOOS's target list (Sebastian et al. 2020, 2021), this may not be a valid assumption, as we have observed more of our bright targets (with high SNR) early on. As there is no simple way to predict the change in recovery fraction for future targets, or any future planet detections, this is only a speculative estimation of how tightly the survey will constrain the occurrence rate of this type of planet over the next few years, or (by considering 1000 extra targets) its lifetime.

In a similar way, I also attempted to predict how many targets we would need to observe to detect another planet like TRAPPIST-1b, if the occurrence rate was indeed 0.22. Equation 5.11 tells us the probability of two successes given N_* objects. I varied the number of stars N_* by including only N_* of the targets in the sample. Once I reached the actual number of targets in the sample, I repeated the sequence. The results are shown in Figure 5.16a. Again, I assumed that the future target stars would be very similar to those already observed. Therefore, by varying N_*,

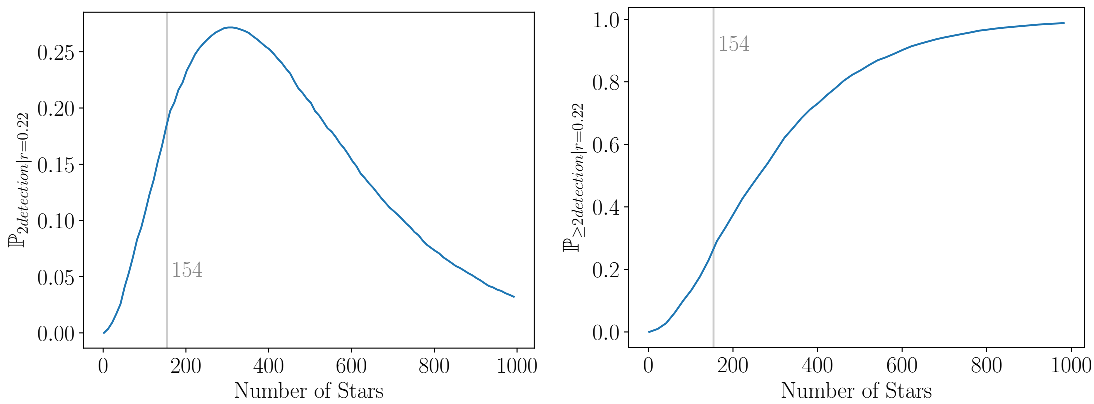

(a) Detection Sensitivity

(b) Completeness

Figure 5.16: (a) The probability, $\mathbb{P}_{\text{2detections}|\text{r}=0.22}$, of detecting exactly two TRAPPIST-1b-like planets in a sample of n stars, assuming an occurrence rate of 22%. I extract the peak of ~310 target objects. The probability of detecting exactly 2 planets decreases with large n as the chance of detecting more planets increases. (b) The probability, $\mathbb{P}_{\geq\text{2detections}|\text{r}=0.22}$, of detecting at least two TRAPPIST-1b-like planets. I mark a line for 154 objects, equivalent to the SSO sample.

I determined that we would have to observe 310 objects to have the greatest chance of detecting exactly 2 TRAPPIST-1b-like planets. I then calculated the probability of detecting 2 or more of these planets (shown in Figure 5.16b) by deducing that $\mathbb{P}(X \geq 2) = 1 - (\mathbb{P}(X = 0) + \mathbb{P}(X = 1))$, and using Equations 5.9 and 5.10. Therefore, given that we have observed 154 targets in the SSO sample, there is approximately a 30% chance that we would have detected at least another TRAPPIST-1b type planet.

5.5 In Perspective

In this chapter I presented a transit-search algorithm for SPECULOOS lightcurves, consisting of a GP-detrending method followed by a BLS-based search tool. I searched for transit structures in SSO lightcurves, finding only two candidate signals. These signals, however, are not conclusive, and in each case, there is only one. I validated the transit-search algorithm by applying it to three known planetary systems, LP 791-18, LHS 3844 and TRAPPIST-1. The algorithm recovered transits from all known planets in the three systems, except TRAPPIST-1h. I injected transits from a range of synthetic planets into the SSO lightcurves and applied my transit-search algorithm to derive its sensitivity and limitations. From orbital geometry, I know the likelihood of a planet to transit. Therefore, the final piece of the puzzle is how many of our targets host planets.

SPECULOOS's survey strategy is optimised for very short period, Earth-sized planets, such as those around TRAPPIST-1. In deriving occurrence rates, therefore, I focused on planets similar to TRAPPIST-1b (R=1–1.25 $R_\oplus$ and P=0.91–1.66 d), with the aim of discovering whether TRAPPIST-1-like red dwarf planetary systems, were in fact exceptional discoveries. Over the course of almost two years of observations I found no planet like TRAPPIST-1b in period and radius. If every object in the SSO sample hosted one of these planets, then we should have found at least 4 of them. The probability of finding none of these planets is 1%. However, if I count TRAPPIST-1b as a discovery within the SSO sample we do have one detection. Given 154 non-detections in my sample, and assuming a Poisson-Binomial distribution, I found it most likely that 22% of ultra-cool dwarfs would host a TRAPPIST-1b analogue. However, it is difficult to draw concrete conclusions, as it is not possible to exclude any possible occurrence rate < 100% with 3σ certainty. The more targets that the SPECULOOS survey observes, the better this value will be constrained.

If 22% of UCDs are host to short period, terrestrial planets like TRAPPIST-1b, then I would expect to have found $\sim$1 planet. This result is consistent with our current state of detections (including the TRAPPIST-1 discovery). From the number of observed objects in the SSO sample, I predicted a $\sim$30% probability that we would have found another TRAPPIST-1b. Instead, we would need to double the SSO sample size to have an approximately 60% chance of detecting another such planet. I conclude that planetary systems like TRAPPIST-1 are not ubiquitous.

Thanks to the Kepler and K2 missions, planets around early and mid M-dwarfs are much better sampled than for UCDs, therefore we can expect them to have more robust occurrence rates. Several of these studies performed on Kepler M-dwarfs (Dressing & Charbonneau 2013; Morton & Swift 2014; Gaidos et al. 2014; Dressing & Charbonneau 2015; Mulders et al. 2015; Gaidos et al. 2016; Hsu et al. 2020; Cloutier & Menou 2020) find similar results: $\sim$2–2.7 sub-Neptune planets[b] around each early M-dwarf (with R=1–4 $R_\oplus$ and $P < 200$ d), twice as many as G-type stars like our Sun, and that M-dwarfs host mostly terrestrial-sized planets with periods longer than 10 d. Dressing & Charbonneau (2015) and Hsu et al. (2020), in particular, find 0.57 and 0.43 super-Earths per early M-dwarf, respectively. Additionally, they conclude that planets larger than 3$R_\oplus$ are very rare around this type of stellar host. Hardegree-Ullman et al. (2019) break down the planetary occurrence rates of Kepler mid M-dwarfs as a function of spectral type. They derive occurrence rates of 0.86, 1.36 and 3.07 planets per star for M3, M4 and M5 dwarfs. In addition, they find almost double the occurrence of short-period, rocky planets (0.5–1.5 $R_\oplus$) for their sample as Dressing & Charbonneau (2015) do for early M-dwarfs,

[b]To clarify, I define *sub-Neptunes* as all planets smaller than Neptune ($< 4\,R_\oplus$), *mini-Neptunes* as planets with sizes 2–4 $R_\oplus$, and *super-Earths* with radii 1–2 $R_\oplus$.

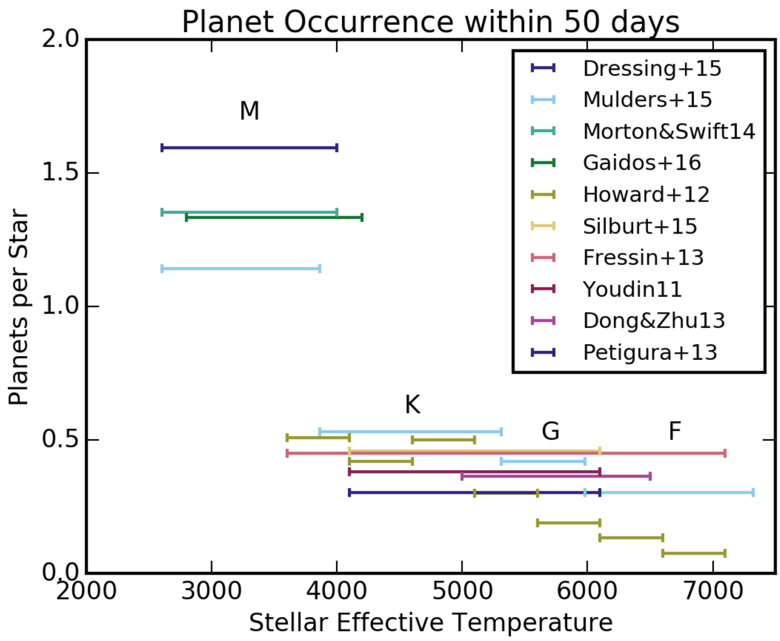

Figure 5.17: Figure 6 from Mulders (2018), which directly compares the occurrence rates of sub-Neptune-sized planets (1–4 $R_\oplus$), with periods less than 50 d, measured by 10 different studies (Dressing & Charbonneau 2015; Mulders et al. 2015; Morton & Swift 2014; Gaidos et al. 2016; Howard et al. 2012; Silburt et al. 2015; Fressin et al. 2013; Youdin 2011; Dong & Zhu 2013; Petigura et al. 2013), as a function of effective temperature.

and 0.27 super-Earths (in this case, 1.5–2.5 $R_\oplus$) per star. However, it is worth mentioning that this analysis was based on a significantly smaller number of planets than most of the earlier M-dwarf studies, and likely have larger uncertainties. In a review of the occurrence rate literature, Mulders (2018) similarly agrees that the frequency of close-in, terrestrial planets increases as we move to cooler stars (see Figure 5.17, or Figure 6 in Mulders 2018).

With a redder observing bandpass than Kepler, TESS will be able to unlock key insights into the planetary populations of cool stars. Since its launch, TESS has already discovered 40 confirmed planets, and over 180 planet candidates, around M-dwarfs (e.g. Vanderspek et al. 2019; Crossfield et al. 2019; Ment et al. 2020), and there have been several new yield estimations for M-dwarf planets (Barclay et al. 2018; Ballard 2019; Feliz et al. 2021). Feliz et al. (2021) analysed 30,000 TESS full-frame images (30-minute cadence) from nearby mid M-dwarfs to estimate occurrence rates. The team find an average occurrence of 3.62 sub-Neptunes (with periods 1–9 d) per mid M-dwarf, in agreement with the increase in occurrence as we move to cooler stars.

Radial velocity surveys tell a similar story. The HARPS team found that ~40% of early and mid M-dwarfs have a super-Earth (with minimum masses $1 \leq M \sin i \leq 10_{\oplus}$) in the habitable zone (Bonfils et al. 2013). The CARMENES team recently predicted an abundance of low-mass planets (also with minimum masses $1 \leq M \sin i \leq 10_{\oplus}$) around a sub-sample of 71 early-to-mid M-dwarfs (M0–M6), with every star predicted to host 1.32 planets, up to periods of 100 d (Sabotta et al. 2021). Just as for transit surveys, these results support an increase in planets around early and mid M-dwarfs compared to hotter, Sun-like stars. However, how any of these findings might extend to ultra-cool dwarfs has yet to be determined.

Due to such a small catalog of UCD planets, it becomes difficult to perform a statistical census of their planetary populations. From the TRAPPIST UCDTS survey, Lienhard et al. (2020) estimate a lower limit on the occurrence of planets like TRAPPIST-1b of 7–14%. This result is based on the detection of TRAPPIST-1 out of a sample of only 40 UCDs. Due to the small sample size, however, this prevents the authors making strong conclusions about the planetary occurrence rates of UCDs. Sestovic & Demory (2020) analysed UCDs observed by K2 to bound the occurrence rate of super-Earths and giant planets as less than 1.14 and 0.29 respectively. With only one detection of a mini-Neptune out of around 700 objects, they are only able to provide upper limits for these two types of planets, and estimate the occurrence of mini-Neptunes to be $0.2^{+0.16}_{-0.11}$. They also conclude that planetary systems like TRAPPIST-1 must not be ubiquitous, as they did not find any in their sample. Sestovic & Demory (2020) find an upper limit of r=0.36 for super-Earths (1.5–2.5 $R_{\oplus}$ with periods 1–10 d) with 95% confidence, in good agreement with the results on earlier M-dwarfs (Hardegree-Ullman et al. 2019). Independently, Sagear et al. (2020) used a very similar K2 sample, with no detections, to place an upper bound of 0.19–0.57 (depending on the planet size) on the occurrence of sub-Neptune planets with periods 1–1.58 d, and 0.25–0.58 on the occurrence of hot mini-Neptunes with periods 1.0-4.07 d. He et al. (2017) observed 44 L and T brown dwarfs with the Spitzer Space Telescope for an average of approximately 20 hours each. For extremely short-periods (<1.28 d), they placed an upper bound on the occurrence rate of 0.67±0.01 for sub-Neptunes, and 0.87±0.03 for rocky planets (0.75–1.25 $R_{\oplus}$). From the discoveries of the Kepler-42b and TRAPPIST-1b systems they derived an average occurrence rate of 0.27 for Earth-sized planets with periods 0.55–6 d, and concluded that the occurrence was likely 0.2–0.3. Ultimately, none of these results rule out an increasing occurrence of terrestrial planets in the ultra-cool regime.

All these studies, as well as my own, rely heavily on *not* detecting planets, or detecting very few. My results from Figure 5.14 imply that hot planets larger than $2R_{\oplus}$ are rare around UCDs, since I did not recover any in the SSO sample. This result aligns with several previous studies (Demory et al. 2016b; He et al. 2017; Lienhard et al. 2020; Sestovic & Demory 2020; Sagear et al. 2020) and with the results for earlier M-dwarfs (Berta et al. 2013; Dressing & Charbonneau 2015). This may be due to an actual lack of these types of planets, or an observational bias if this size of planets tend to have longer orbital periods, where our completeness is low. I can directly compare the occurrence rate results derived by Sestovic & Demory (2020) and Sagear et al. (2020) in the radius-period bin that contains TRAPPIST-1b. Sestovic & Demory (2020) find a 95% credible upper bound of 36% for planets of 1–1.5 $R_\oplus$ with periods 0.9–1.6 d. Sagear et al. (2020) find an upper limit of 57% for planets with radii 1–1.5 $R_\oplus$ and periods 1.0-1.58 d with 95% confidence. I find the most likely occurrence rate of these type of planets to be 22%, consistent with the bounds determined by He et al. (2017), Lienhard et al. (2020), Sestovic & Demory (2020) and Sagear et al. (2020).

I find that, with SPECULOOS's current observing strategy, we are much less sensitive to rocky planets with periods longer than 3 days. Kepler tells us that these cooler planets are bountiful around earlier M dwarfs (Dressing & Charbonneau 2013), and that extremely short period planets are rare. So far it is unclear whether this applies to UCDs as well, though perhaps it explains the low numbers of detections by SPECULOOS and MEarth (Berta et al. 2013), whose observing strategies both favour very close-in planets. To probe longer period planets, that orbit further away from their host, we would need to spend longer observing each target. However, this would come at the cost of the number of targets we could observe, and would, therefore, reduce our chance of detecting another coveted, short-period, TRAPPIST-1b-like planet. I note that in the SSO sample very few targets were observed for their full 100–200 hours; 80% have been observed less than ~80 hours, therefore, our detection efficiency will be higher for longer period planets once those targets have been fully observed.

The SPECULOOS strategy is designed to find short-period, habitable zone, terrestrial planets, just like the TRAPPIST-1 system (Sebastian et al. 2021). But how optimal is this observing strategy? If these planets are scarce, and TRAPPIST-1 was a lucky find, then there is a risk that SPECULOOS may yield few, or no more, detections in its lifetime. However, these planets are exceptionally interesting to study, and, for now at least, remain our best candidates on which to detect life. This means that while they may be more rare than we had hoped, finding another TRAPPIST-1 could be worth the wait. With the SPECULOOS project ongoing for at least several more years, it remains to be seen whether this strategy will pay off.

5.6 Future Directions

There are also several potential changes that I would implement to my current method. I have made a couple of large assumptions in this chapter. Firstly, I have not considered errors for the mass and radii measurements in the SPECULOOS target list. Secondly, while a Poisson-Binomial distribution provides a simple approximation of the state of detections/non-detections, it may not be the best way to describe planetary occurrence. In reality, we do not know the likelihood function, and I will explore a range of alternative methods in the coming months. Hierarchical Bayesian models (HBM), where data uncertainties are represented by distributions for model parameters, have been used to successfully represent planetary occurrence rates (Foreman-Mackey et al. 2014; Sestovic & Demory 2020; Hsu et al. 2020). Approximate Bayesian computation (ABC) can bypass evaluation of the likelihood function by generating simulated data sets (Hsu et al. 2020; Kunimoto & Bryson 2020), though it is extremely computationally expensive. Also, by repeating this work, but including several possible transit durations in the BLS transit-search, from 0.01–0.05 d, this should improve the detection efficiency. I will also test the use of Transit Least Squares, a similar algorithm to BLS but which searches for transit shapes with limb-darkening included. Not only will more accurate transit durations provide better SNR estimates, but it will prevent transits being rejected if the BLS duration spans less than half of the real transit duration. Though it is computationally expensive, it may allow me to place a threshold on the transit depth for successful recovery.

In repeating this work, it would also be interesting to track the Gaussian Process hyperparameters in the non-periodic cases, in case they illuminate some insights about the nature of remaining systematics in SPECULOOS lightcurves. If I expanded on the GP-modelling aspect of this work to consider more than just time-varying systematics, we may also find that there are underlying correlations with instrumental (e.g. position of the target on the CCD) or atmospheric parameters (e.g. airmass) that we had not fully considered and could be easily detrend. Exploring the GP hyperparameters and multi-parameter modelling of our lightcurves was out of the scope of this PhD work, however, it could be extremely useful to the SPECULOOS team moving forward.

I will continue to work on improving these constraints, with the plan to publish my results early next year. I will expand the work in this chapter to the entire period-radius space, to make predictions more generally about the occurrence of terrestrial, temperate planets around ultracool objects. Deriving these results as a function of spectral type would be a necessary addition to this work: do mid M-dwarfs host similar numbers of planets as late M-dwarfs? Or as L-type dwarfs? Previous results have shown a strong variation in occurrence rates between early and mid M-dwarfs (Hardegree-Ullman et al. 2019), and it remains to be seen if this also applies to later-type stars. In this chapter, I have also not considered the presence of multi-planet systems. Compact multiple systems are predicted to be common around mid M-dwarfs (Muirhead et al. 2015; Hardegree-Ullman et al. 2019). These multi-planetary systems may be similarly frequent around UCDs, considering detections of those around TRAPPIST-1 and Teegarden's Star. The chance of detecting at least one transit from one of the planets in a multi-planetary system is greater than for a single planet (demonstrated on TRAPPIST-1 in Section 5.2.2). Therefore, by including transits from multiple planets in lightcurves, I would expect our detection efficiency to increase substantially. Doing so would allow me to make predictions about the architecture of ultra-cool planetary systems. As TRAPPIST-1b was recovered successfully, and was the most significant signal in the TRAPPIST-1 lightcurve, I assume that it would have been recovered regardless of the presence of other planets. A lack of any further significant planet detections with SPECULOOS can also allow me to assess SPECULOOS's observing strategy. Will observing each target object for longer, rather than observing more target objects, increase our detection potential? The next few years of the SPECULOOS project will further clarify these results, and help optimise our strategy, by providing hundreds more non-detections, and hopefully, detections.

CHAPTER 6

CONCLUSIONS AND THE FUTURE

In this concluding chapter, I summarise the four main projects that comprise this thesis. I highlight the main results and suggest potential directions for future work to expand on the ideas presented here.

6.1 Summary

The central aim of my PhD has been to optimise the detection of temperate, terrestrial planets around ultra-cool dwarfs for the SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars) project. This project's main goal is to provide the most promising candidates for future atmospheric characterisation by the James Webb Space Telescope (JWST). I have focused on the data from the largest facility of the SPECULOOS network, the SPECULOOS Southern Observatory (SSO) at ESO Paranal Observatory in Chile, which comprises four 1-m, robotic telescopes. To find terrestrial planets orbiting in their host objects' Habitable Zones the most promising place to look is around ultra-cool dwarfs. Not only are UCDs abundant in our local stellar neighbourhood, but Earth-sized, temperate planets are easier to detect around UCDs than any other type of star. Additionally, UCDs are predicted to host plentiful numbers of small planets, and these planets remain the best candidates for the future detection of spectroscopic signatures in their atmospheres. However, in the past, detecting small planets around UCDs was deemed unfeasible due to the low photometric precisions that could be achieved at visible wavelengths. Following from this historic lack of exoplanet surveys focused on UCDs, we know almost nothing about their planets. To remedy this, the past decade has delivered a new generation of 1-m class, high precision, near-infrared optimised telescopes, bringing these red objects into reach. The SPECULOOS project (among others, Gillon et al. 2013; García-Mejía et al. 2020; Tamburo & Muirhead 2019; Gibbs et al. 2020) seeks to reveal the hidden planetary populations around UCDs by discovering transiting planets and placing constraints on their frequency and diversity.

For my PhD, I maximised the detection potential of temperate, Earth-sized planets through four main avenues of work: extracting precise photometry by mitigating ground-based systematics, evaluating the photometric performance of the SSO, analysing stellar variability to constrain its photometric impact, and finally, establishing bounds on the unexplored planetary populations around UCDs.

To have the best chance of finding Earth-sized transiting planets we must extract extremely precise photometry from the SSO. In large part this has meant dealing with the photometric challenges of observing from the ground, as observing very red targets in the near-infrared can result in complex systematics. These systematics are related to the temporally varying, wavelength-dependent opacity of the Earth's atmosphere. To mitigate these effects, I led the development of the SSO Pipeline in PYTHON, a data reduction pipeline specific for SPECULOOS, and the extremely red targets that it observes (Murray et al. 2020). This automated pipeline performs reduction and calibration of images, crossmatches with known star catalogues, measures precise aperture photometry, and generates differential lightcurves. The entire SPECULOOS team then searches these lightcurves for transiting planets. The SSO pipeline also introduced a novel method for correcting the effects of telluric water vapour on ground-based observations from first principles. This pipeline has now been adopted to also process the observations from our SPECULOOS Northern Observatory.

To optimise the SSO Pipeline, I assessed the precision of its photometric products over the first few years of SSO operations. The quality checks and feedback loops implemented into the pipeline are essential for monitoring the health of the observatory, as well as to evaluate whether SPECULOOS is reaching its survey goals. This work highlighted the excellent precision of the SPECULOOS telescopes, but also the need to account for stellar variability. This variability dominates the photometric contamination limiting the detection of terrestrial-sized planets. I compared simultaneous observations taken by SPECULOOS and the space telescope TESS, to demonstrate the impressive performance of the SSO for ultra-cool objects, and the power of the synergy between the two surveys. Analysis of the SSO's photometry shows how, with specialised photometric treatment, the SPECULOOS network can reach the required sub-millimag level precisions to detect terrestrial, habitable-zone planets around bright, quiet UCDs.

A key feature that severely affects our planetary detection potential is the photometric activity of the host stars. While UCDs are seen to be very active objects, with both large-amplitude photometric variability and flares, the source and consequences of their behaviour are not well understood. For many of our targets the limiting factor for precisions high enough to detect Earth-sized planets, is stellar activity, therefore I implemented a flare detection algorithm using a hybrid of manual and automatic methods, to find and mask flares. I also measured the rotation periods of our targets to detrend out periodic brightness modulations using Gaussian Processes. In parallel with this work on mitigating stellar variability in our lightcurves, I performed a large-scale study of the flaring and rotational activity of our ultra-cool targets (Murray et al. 2022).

The phrase "know thy star, know thy planet" is regularly coined in the exoplanet community, however, it is particularly relevant when considering ultra-cool dwarfs. Despite being the most encouraging targets for which to detect rocky planets, UCDs demonstrate extreme flaring activity, and photometric variability, that appear to be major concerns for the presence of life. However, the emission of white-light radiation caused by flares may help, not hinder, life, by providing the missing energy at the bluer end of the spectrum, that is scarce around these very red objects. I studied the flare rates and energies from the large sample of flares in the SSO's lightcurves, to probe the relationships between activity and rotation (Murray et al. 2022). I compared these results to similar studies using MEarth and TESS data to underline their highly complementary nature. I directly addressed two different scenarios in which flares may assist the origin and sustenance of life on planets they host. On average, I found that the UV flux of flaring cool stars was too low for the synthesis of ribonucleotides, a major step in prebiotic chemistry, and there was not enough visible light incident on planets around cool stars to sustain an Earth-like biosphere. In both scenarios, I found that the flaring activity of the most active, earlier type M-dwarfs could provide this necessary radiation. However, the frequency and energy of flares decreased with spectral type. Considering only these two Earth-centric pathways for life suggests that planets around UCDs may not be favourable sites for abiogenesis or photosynthesis, though there are a huge number of other factors, not considered here, that will affect their habitability.

Finally, after removing the contamination of our ground-based photometry from both non-astrophysical and astrophysical sources, it became possible to quantify the detection potential of the SSO for Earth-sized, temperate planets (Murray et al., in prep.). For this purpose, I performed transit injection-recovery tests to determine the detection sensitivity of my transit-search algorithm and the survey's completeness limits. These factors dictate how likely I would be to detect a planet if it exists. From the results of a transit search through the SSO's lightcurves, I found a handful of potential signals. However, I found no transits similar to the TRAPPIST-1 planets. This result is key to constraining the occurrence of short period, Earth-sized planets, on which SPECULOOS's strategy is based. With only one detection, of TRAPPIST-1b, I found it very unlikely that every UCD would host a similar type of planet. In fact, given the current survey result, I derived that the most likely occurrence rate of these types of planets is 22%. If less than 22% of UCDs hosted a TRAPPIST-1b analogue, then I predicted that, to guarantee discovery of at least one more similar planet, we would need to observe a great number more stars. I plan to expand on these results, to more generally place limits on the existence, or occurrence rates, of temperate, terrestrial planets around low-mass objects. Tighter constraints on planetary populations will lead to more accurate yield estimates for ground and space-based missions and tell us whether red dwarf planetary systems are common.

As a member of the SPECULOOS team, my work has involved not only the results presented in this thesis, but also opportunities to engage in all facets of the project. I travelled twice to Paranal, Chile, once to actively take part in the commissioning of the Callisto telescope, and a second time for a maintenance mission. I regularly operate the SSO telescopes remotely as part of a monthly observing schedule. Working with other members of the SPECULOOS team has also allowed me to collaborate on the scheduling of targets, the creation of the interactive Portal, and compare data reduction pipelines between the SSO, Saint-Ex and TRAPPIST.

6.2 Looking Forward

Even in the duration of my PhD there have been giant leaps in the exoplanet field. In particular, the launch of the TESS mission in 2018 has led to over 700[a] publications, including 340 this year alone. Despite not being specifically aimed at late-type stars, TESS's relatively red bandpass has allowed for more in-depth study of low-mass, cool stars (up to late M-dwarfs) than its space mission predecessors. Notwithstanding these significant advances made in exoplanetary science, we still know very little about the objects at the coolest end of the main sequence, UCDs, and the planets they may host. The SPECULOOS project, which began official scientific operations in 2019, marks the first large-scale photometric transit survey specifically targeting ultra-cool objects. This project aims not only to search for temperate, Earth-sized planets amenable for future atmospheric characterisation, but also to study over one thousand of these extremely red UCDs, in more detail that ever before. SPECULOOS, and other UCD surveys (such as EDEN and the CARMENES radial velocity project), will help answer some of the biggest questions surrounding their magnetic activity, the frequency and diversity of their potential planetary systems, and the opportunities for life within those systems. With several more UCD studies planned for the future (García-Mejía et al. 2020; Tamburo & Muirhead

[a] https://heasarc.gsfc.nasa.gov/docs/tess/tpub.html

2019), the next decade looks extremely promising for advancing our understanding of these mysterious objects.

By finding the best candidates capable of initiating and supporting life, I aim to help steer the search for life outside the Solar System. I will continue expanding and improving on the work presented in Chapter 5, with the aim to publish it in early 2022. Using the largest photometric sample of these loss-mass objects to date I will tighten the constraints on their mostly unknown planetary populations. I will make predictions about the ubiquity of Earth-sized, temperate planets, to improve yield estimates for current and future transit surveys focused on ultra-cool dwarfs. Finally, I hope to address whether red dwarf planetary systems, such as TRAPPIST-1, are in fact exceptional discoveries.

More generally, the future of exoplanetary science looks increasingly bright. The SPECULOOS project aims to provide first-class planetary candidates for atmospheric follow-up with JWST, which will launch later this year. Then for at least the next decade, JWST will lead the search for life, delivering powerful new measurements of those exoplanets' atmospheres. Fully characterising the host objects, and the environments that they create for their planets is one of the keys to interpreting these measurements, and to finding the most promising planets to host life. Stellar contamination from active host stars will be a significant limitation to the scientific potential of this program, and so constraining stellar activity (for example, stellar spot coverage, flaring activity etc.) will be extremely important. This may prove a major issue for studying the atmospheres of planets orbiting highly active, and poorly studied, M-dwarfs and UCDs, that are excellent candidates for JWST. Activity studies such as those performed in this thesis (Murray et al. 2022) and elsewhere (Newton et al. 2016; Seli et al. 2021) will therefore be extremely helpful moving forward, especially those that exploit the synergy between ground- and space-based surveys. The launches of PLATO, ARIEL and LUVOIR within the next twenty years will at last provide the capability to detect and characterise a potentially habitable Earth-like planet around a Sun-like star. Since the beginning, the Exoplanet field has been driven by the search for life. Within our lifetime these missions will allow us to determine how common planets like our own are, to probe the very nature of habitability and to find our place in the universe.

APPENDIX A

SSO QUALITY CHECKS

In this appendix I present the values of the readout noise (Figure A.1) and dark current (Figure A.2), as two main quality checks. Their values are monitored over the timespan of the SSO data sample, 1st June 2018 to 23rd March 2020 (as defined in Section 3.1.1), to flag any immediate issues with the telescopes.

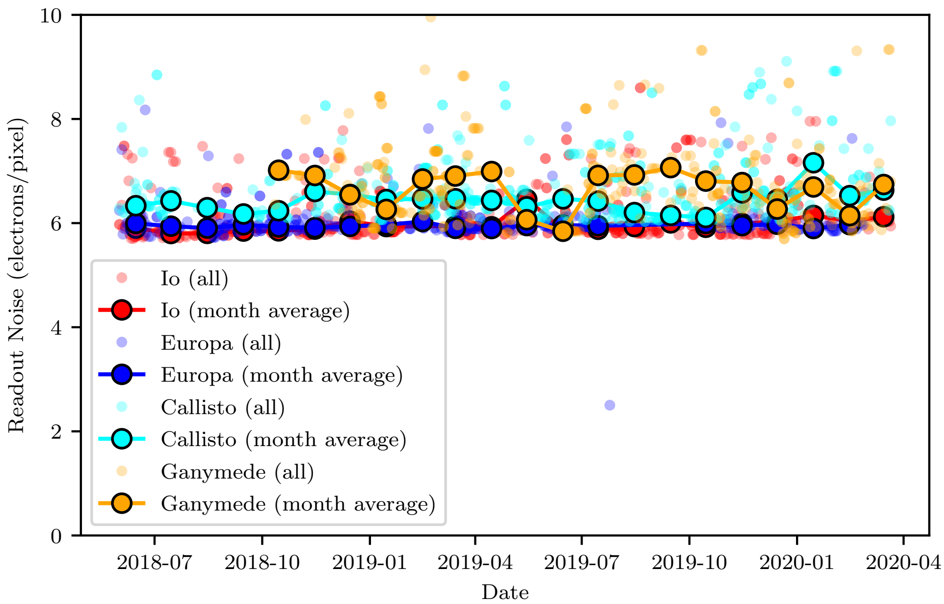

Figure A.1: Readout noise (in e^-) of the four telescopes throughout the duration of the SSO data sample. Io is shown in red, Europa in blue, Callisto in cyan and Ganymede in orange. Both the night-by-night values and monthly averages are shown. Ganymede demonstrates the most variable readout noise, however, this is still within the required telescope specifications (priv. com. D. Sebastian)

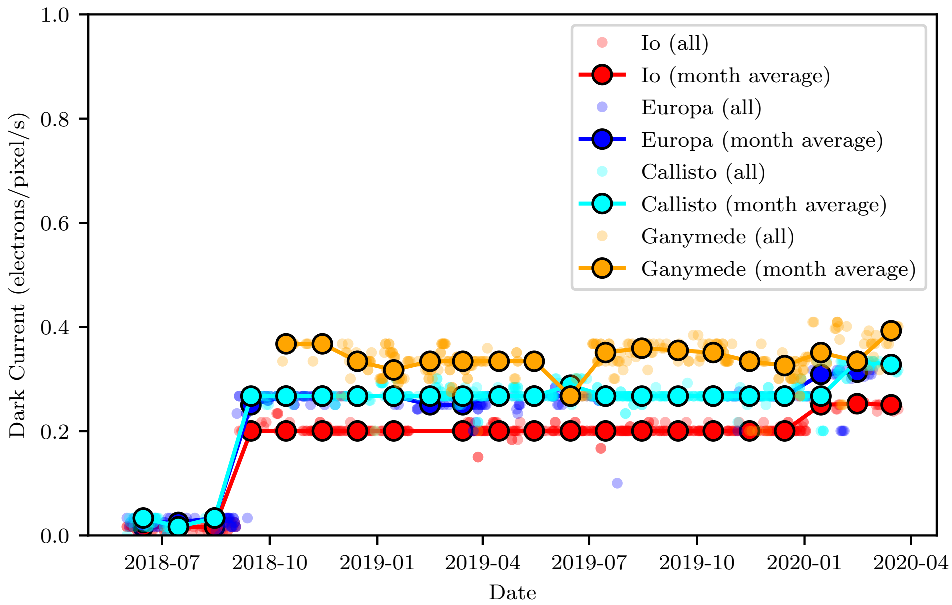

Figure A.2: Average dark current (in e^{-}/s/pixel) of the four telescopes throughout the duration of the SSO data sample. Io is shown in red, Europa in blue, Callisto in cyan and Ganymede in orange. Both the night-by-night values and monthly averages are shown. The jump in dark current at the end of September 2018 is due to the 10°C increase in CCD operating temperature. There is a smaller jump in dark current on the 8$^{\text{th}}$ January 2020. Before this date we had only taken 15, 30 and 60 s exposure dark images every day and after this date we included also 120 s dark images. Longer exposures increases the amount of thermally excited electrons, and subsequently, very slightly increases the average dark current value.

Appendix B

Blended Targets

In this appendix I demonstrate the situation where a target object has a very nearby contaminating star (Figure B.1). I show that in this case there is no optimal aperture, as the photometry in all apertures is adversely affected by the 'blending' of the target (Figure B.2).

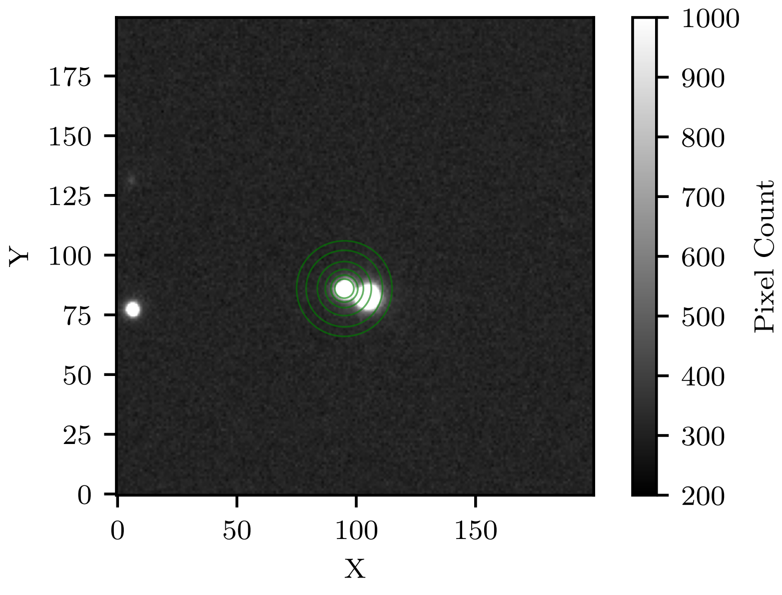

Figure B.1: An example of a field with a blended target. This science image was taken on 21st March 2020. The target is circled in green, with apertures increasing in size from aperture 3 (with a radius of 4 pixels) to aperture 8 (20 pixels). At the largest aperture the fluxes from both stars are contained.

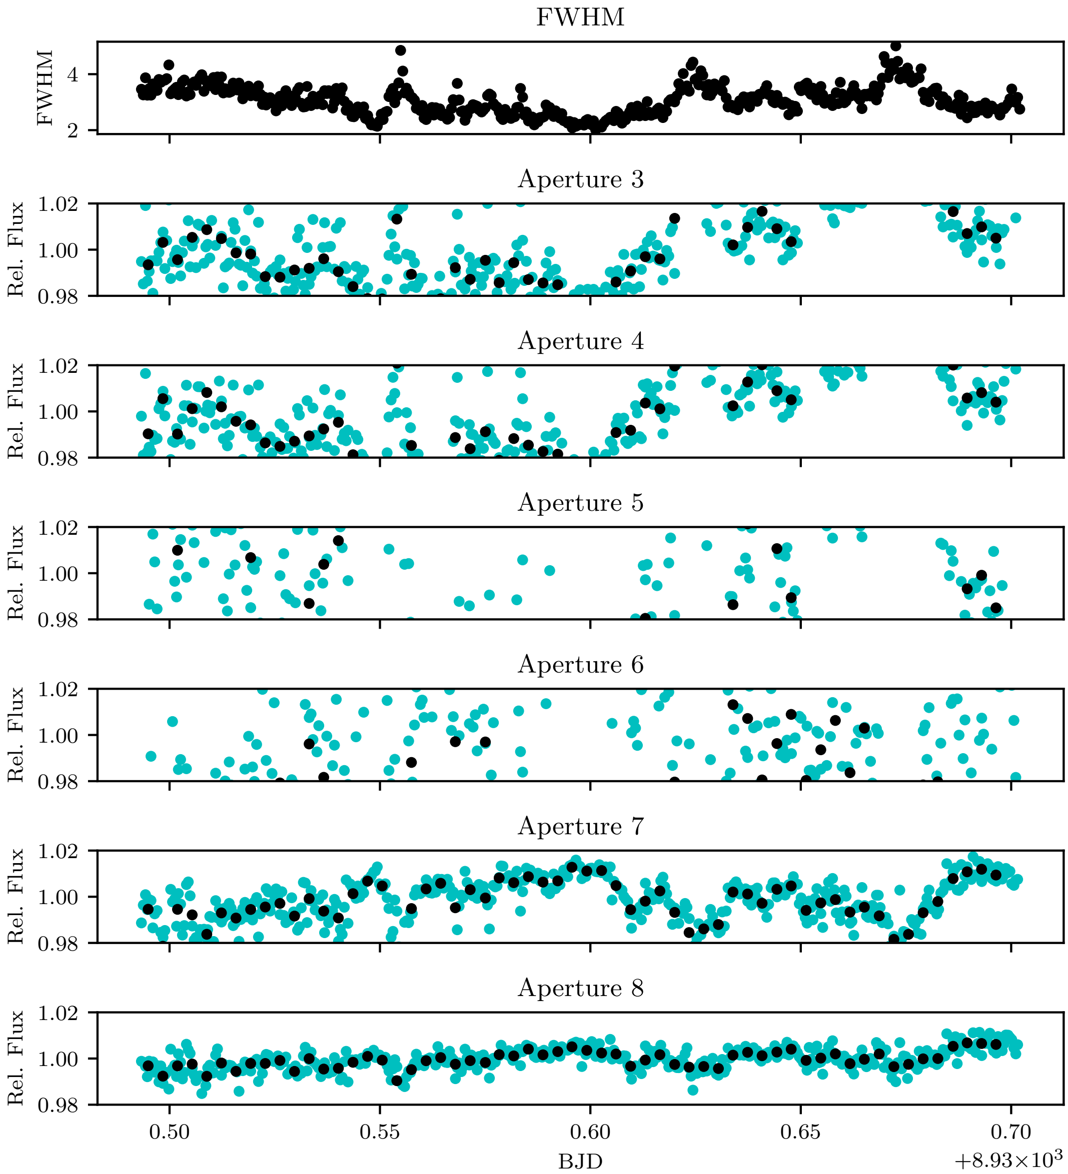

Figure B.2: The top plot is the FWHM for the night of 21st March 2020 (same target as Figure B.1). The six plots below are this target's lightcurves for 6 different apertures (from 3–8). The cyan points are the unbinned data and the black points are binned every 5 minutes. The lightcurve for aperture 8 appears to be the best quality (and would be chosen as the optimal aperture), however, it contains the flux from 2 stars, and therefore we cannot know whether any structure comes from the target or the nearby star.

Appendix C

SSO Data Sample

In this appendix I present a table of the 154 objects contained within the SSO Data Sample (see Section 3.1) that have been observed for longer than 20 hours. Each object is identified using the corresponding ID in the Gaia DR2 catalogue, and its spectral type classification, effective temperature, radius, mass and J-mag are as calculated by Sebastian et al. (2021). I indicate whether any flares have been identified for each target, using the flare detection method in Section 4.2, and the rotation periods and classes are as assigned in Section 4.3.

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
1227705135863076864	M6	2817	0.12	0.10	11.01	–	N	Y
2328674716056981888	L0	2177	0.10	0.08	13.58	7.22569	L	N
2331849006126794880	M8	2573	0.11	0.09	11.65	–	N	Y
2332937316480303616	M6	2826	0.13	0.10	12.38	–	N	Y
2340736324254735488	M9	2329	0.11	0.08	12.40	6.65037	L	N
2349207644734247808	M7	2629	0.11	0.09	13.02	0.686611	B	Y
2393563872239260928	M6	2778	0.12	0.10	11.75	9.52537	L	N
2531195858721613056	M9	2388	0.10	0.08	11.69	8.10832	U	N
2615804824667108736	M4	3110	0.17	0.15	11.52	3.95463	B	N
2631857350835259392	M5	2940	0.14	0.11	11.11	0.126905	A	Y
2886691573123995520	M6	2817	0.13	0.10	11.52	4.6768	U	N
2915187925218983296	M8	2605	0.11	0.09	12.53	2.21553	U	N
2945455521833759744	M8	2539	0.11	0.09	13.01	12.0269	U	N
2985035874544160384	L2	2015	0.09	0.07	13.08	15.0355	L	N
3005440443830195968	M5	2924	0.14	0.11	10.98	3.25168	U	Y
3013672487388307456	M6	2821	0.12	0.10	11.82	6.06948	L	Y
3015582682682095232	M5	3059	0.16	0.13	12.02	–	N	N
3015584160150844416	M6	2784	0.12	0.10	12.30	–	N	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
3136232952593850496	M5	2952	0.14	0.12	11.17	22.8084	L	Y
3175523485214138624	M6	2837	0.13	0.10	12.51	–	N	Y
3192982561632526208	M7	2625	0.11	0.09	12.18	3.54788	U	N
3200303384927512960	M7	2714	0.11	0.10	10.66	0.333683	A	Y
3257243312560240000	M6	2819	0.12	0.10	11.30	2.87669	U	Y
3280675417175220224	M7	2711	0.12	0.09	11.62	7.27517	U	Y
3421840993510952192	M6	2767	0.11	0.10	10.70	4.2967	B	Y
3460806448649173504	L2	2021	0.10	0.07	12.81	0.331179	A	N
3473814014803013248	M6	2826	0.12	0.10	12.87	9.80213	B	N
3474993275382942208	M6	2814	0.12	0.10	11.98	–	N	N
3479243059623564160	M7	2701	0.12	0.09	12.89	3.18133	B	N
3486653767994941952	M6	2860	0.13	0.10	12.12	7.53557	L	N
3493736924979792768	M7	2687	0.12	0.09	10.93	0.142204	A	Y
3504014060164255104	M7	2647	0.11	0.09	11.79	16.3255	L	Y
3540491079971287552	M8	2476	0.11	0.08	12.63	–	N	N
3542490095189650560	M6	2852	0.13	0.10	11.73	–	N	Y
3853943806185777664	M8	2502	0.11	0.09	12.30	–	N	N
3899128064731508736	M6	2810	0.12	0.10	11.26	2.47154	A	N
3987041475434695936	M6	2912	0.13	0.11	11.30	–	N	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
4053559111471124608	M5	2939	0.13	0.11	11.43	0.0904842	A	Y
4073379904492310656	M6	2847	0.13	0.10	12.61	–	N	Y
4127182375667696128	M6	2868	0.13	0.11	11.30	0.568644	U	Y
4349305645979265920	M6	2871	0.13	0.11	11.35	3.31232	U	Y
4364185165038935040	M9	2344	0.10	0.08	12.05	3.75716	U	N
4404521333221783680	M8	2618	0.11	0.09	11.90	–	N	Y
4405150047715392128	M6	2783	0.12	0.10	11.30	4.26036	B	N
4450376396936878336	M4	3082	0.16	0.14	11.61	–	N	Y
4489306942580947584	M7	2720	0.12	0.09	11.91	–	N	Y
4620448786799883136	M6	2911	0.13	0.11	11.34	–	N	Y
4654435618927743872	M5	2923	0.13	0.11	11.19	1.23056	U	Y
4656782698309778432	M5	2923	0.14	0.11	10.35	0.165194	A	Y
4680959790759112320	M6	2806	0.12	0.10	13.20	8.93941	U	N
4709450095539630208	M4	3152	0.19	0.16	11.47	0.5688	A	Y
4767876769050721920	M6	2798	0.12	0.10	12.52	6.06472	L	Y
4769527685759537280	M5	2953	0.14	0.12	12.34	2.65202	U	N
4774420478143249536	M6	2780	0.12	0.10	11.69	–	N	Y
4825880783419986432	M7	2656	0.11	0.09	12.00	65.4684	L	Y
4833780362148081536	M5	2932	0.13	0.11	12.66	1.1473	U	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
4854708878788267264	M5	3040	0.16	0.13	12.16	3.08549	U	N
4860376345833699840	L0	2313	0.10	0.08	10.72	–	N	Y
4871414343064541824	M5	2925	0.13	0.11	11.61	0.466827	A	Y
4923179694098110848	M6	2758	0.11	0.10	12.46	–	N	N
4928644747924606848	M8	2507	0.11	0.09	12.23	11.8185	L	Y
4928939932436444544	M7	2627	0.11	0.09	13.55	–	N	N
4929042942932181888	M6	2864	0.13	0.11	11.54	5.49037	U	Y
4967628688601251200	M7	2711	0.12	0.10	12.44	0.681987	B	N
4971892010576979840	M7	2642	0.11	0.09	11.62	–	N	Y
4983839338285335680	M8	2558	0.11	0.09	13.15	–	N	Y
5047423236725995136	M8	2539	0.11	0.09	11.69	–	N	Y
5055805741577757824	M7	2642	0.11	0.09	11.36	25.6616	L	Y
5062172669818053760	M6	2905	0.13	0.11	12.20	0.777832	B	N
5064526827293371008	M6	2855	0.13	0.10	12.77	3.33577	U	N
5072067381112863104	M6	2777	0.12	0.10	12.68	12.0423	L	Y
5076670520902556544	M6	2802	0.12	0.10	12.28	0.574373	A	Y
5116817882319879936	M8	2586	0.12	0.09	12.05	32.0324	U	N
5128049359237940224	M8	2512	0.11	0.09	13.08	0.142722	A	N
5156623295621846016	M7	2682	0.11	0.09	12.55	6.75219	B	Y

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
5207997014358306816	M6	2827	0.13	0.10	11.30	–	N	N
5235273458327588224	M7	2649	0.11	0.09	12.08	–	N	N
5321049521489667328	M6	2784	0.12	0.10	12.02	–	N	Y
5368479017141170304	M9	2355	0.11	0.08	11.43	10.0273	L	Y
5383082833650178816	M6	2822	0.13	0.10	11.82	5.04329	U	Y
5392287051645815168	M7	2653	0.11	0.09	12.89	0.182261	A	Y
5424690587034891264	M6	2861	0.12	0.11	11.54	0.114879	A	Y
5424690587034982144	M6	2880	0.13	0.11	11.48	0.133995	A	Y
5432670704985289088	M6	2802	0.13	0.10	10.90	–	N	Y
5449420287163389184	M7	2682	0.12	0.09	12.94	–	N	Y
5469802724480366848	M7	2626	0.11	0.09	12.79	–	N	Y
5541664742899490176	M6	2797	0.12	0.10	13.09	11.261	B	N
5565156633450986752	M6	2905	0.14	0.11	11.09	1.64451	B	Y
5595785866305529216	M9	2440	0.11	0.08	13.11	0.78274	B	N
5602408602804816768	L0	2194	0.10	0.08	13.16	–	N	N
5604989328392927232	M6	2789	0.12	0.10	12.40	0.230339	A	Y
56252256123908096	M7	2632	0.11	0.09	11.76	0.613862	A	Y
5637175400984142336	M7	2657	0.11	0.09	12.89	0.572865	U	N
5652718166073269888	M6	2802	0.13	0.10	12.39	–	N	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
5657734928392398976	M7	2693	0.12	0.09	13.00	–	N	Y
5676353096222033792	L2	2032	0.10	0.07	12.78	–	N	N
5682008056323530368	M6	2825	0.12	0.10	11.07	–	N	Y
5698599996038907520	M6	2823	0.12	0.10	12.26	–	N	N
5704261999871958656	M6	2912	0.13	0.11	12.34	28.1255	L	Y
5723739672264914176	L1	2088	0.09	0.07	12.80	0.145299	B	N
5865271012355608704	M7	2665	0.11	0.09	12.56	–	N	Y
5888509602928271488	M7	2624	0.11	0.09	12.66	0.392684	B	Y
5918136424731578240	M6	2822	0.13	0.10	12.02	0.196078	A	Y
5923656557253383936	M9	2453	0.10	0.08	13.63	14.8292	B	N
5970493789759784192	M5	2938	0.13	0.12	12.08	1.36163	B	Y
5972252905283041664	M6	2916	0.13	0.12	12.85	–	N	Y
5986422934440315648	M6	2775	0.12	0.10	12.76	–	N	N
6018209128388194688	M5	3050	0.15	0.13	10.37	0.286606	A	Y
6056881391901174528	M5	2953	0.14	0.12	11.15	–	N	Y
6060965630906108672	M6	2885	0.13	0.11	11.71	1.52668	U	Y
6067637295632046464	M9	2337	0.10	0.08	13.21	–	N	N
6119528334597735296	M6	2812	0.12	0.10	11.84	0.232153	A	Y
6127211412606464640	M6	2763	0.12	0.10	12.61	–	N	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
6166055474788409088	M6	2815	0.12	0.10	12.55	–	N	N
6166161199701848832	M4	3165	0.24	0.20	11.13	–	N	N
6166184667404110464	M5	2959	0.14	0.12	11.96	2.41985	B	N
6219675839377104128	M7	2711	0.12	0.09	12.29	24.2047	L	Y
6225165597853415552	L0	2298	0.11	0.08	13.25	13.0064	L	N
6234648404610590208	M8	2613	0.12	0.09	11.65	–	N	N
6265453524968112640	M8	2502	0.11	0.09	11.38	1.74805	U	N
6340981796172195584	M4	3160	0.20	0.18	11.59	0.223211	A	Y
6353933390312920832	M6	2881	0.13	0.11	12.67	–	N	N
6357834388848708224	M6	2793	0.12	0.10	12.35	2.03399	U	Y
6368902966247688320	M6	2904	0.13	0.11	12.32	0.405271	B	Y
6401461910728172672	M6	2883	0.13	0.11	12.37	–	N	N
6405457982659103872	M6	2889	0.13	0.11	10.92	0.143118	A	Y
6421389047155380352	M6	2900	0.13	0.11	10.47	–	N	Y
6442331445127207808	M9	2430	0.12	0.08	13.20	11.4956	L	N
6448329418495400960	M7	2741	0.12	0.10	12.24	16.0032	B	N
6473160651658069120	M6	2829	0.13	0.10	11.62	–	N	Y
6474508889136390528	M6	2808	0.13	0.10	12.95	–	N	N
6481588305207220864	M8	2545	0.11	0.09	12.94	–	N	N

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
6494861747014476288	M6	2906	0.13	0.11	12.39	2.02021	U	N
6502618938988110848	L0	2258	0.10	0.08	13.39	0.115832	A	N
6504700451938373760	M6	2810	0.13	0.10	12.62	13.5774	L	N
6505691764750006656	M6	2897	0.13	0.11	13.12	25.6717	B	Y
6506200976072829184	M7	2712	0.12	0.10	12.25	0.579832	U	Y
6508226723167166336	M7	2709	0.12	0.09	12.57	8.94936	B	N
6515783804023662336	M7	2691	0.11	0.09	13.05	5.76835	L	N
6572811211549713792	M6	2916	0.13	0.11	11.45	1.49344	U	Y
6576611428676004608	M8	2474	0.11	0.08	13.43	2.08607	U	N
6594527803147581568	M6	2908	0.13	0.11	12.13	1.7671	U	Y
659464504288593536	M6	2825	0.12	0.10	11.05	15.2381	L	Y
6633500847494138496	M8	2512	0.11	0.08	13.23	–	N	N
6634463852178515456	M6	2752	0.12	0.10	12.82	–	N	N
6637017330494491648	M7	2736	0.12	0.10	11.81	12.7356	L	N
6637022759333088768	M5	2958	0.14	0.12	12.03	–	N	Y
6733860940302404864	M7	2666	0.12	0.09	12.14	0.269839	A	Y
6759481141756109056	M7	2744	0.12	0.10	11.45	–	N	Y
6766848728654380160	M8	2605	0.12	0.09	12.35	–	N	N
6784262587654030336	M9	2370	0.11	0.08	13.47	12.2504	L	Y

Continued on next page

Gaia DR2 ID	Spectral Type	Teff (K)	Radius (R)	Mass (M)	J-mag	Period (d)	Rotation Class	Flaring?
6798222846275675520	M5	2928	0.13	0.11	11.36	–	N	N
6858958940879829248	M6	2791	0.12	0.10	11.81	0.422189	A	Y
6914281796143286784	M6	2829	0.13	0.10	11.45	1.27649	B	Y

Table C.1: The SSO Data Sample

Appendix D

Correcting Flare Occurrence Rates for Completeness

In this appendix, I present the smoothed recovery fractions as a function of flaring energy (calculated using the procedure in Section 4.5.6) for all stars, and the average recovery fraction, separated into spectral type bins. For each spectral type bin, I also present the MCMC fits to the measured flare occurrence rates, demonstrating the fit to Equation 4.16 and corresponding 'intrinsic' flaring rates. The spectral type bins are M4–5 (Figures D.1 and D.2), M6 (Figures D.3 and D.4), M7 (Figures D.5 and D.6), M8 (Figures D.7 and D.8) and M9-L0 (Figures D.9 and D.10).

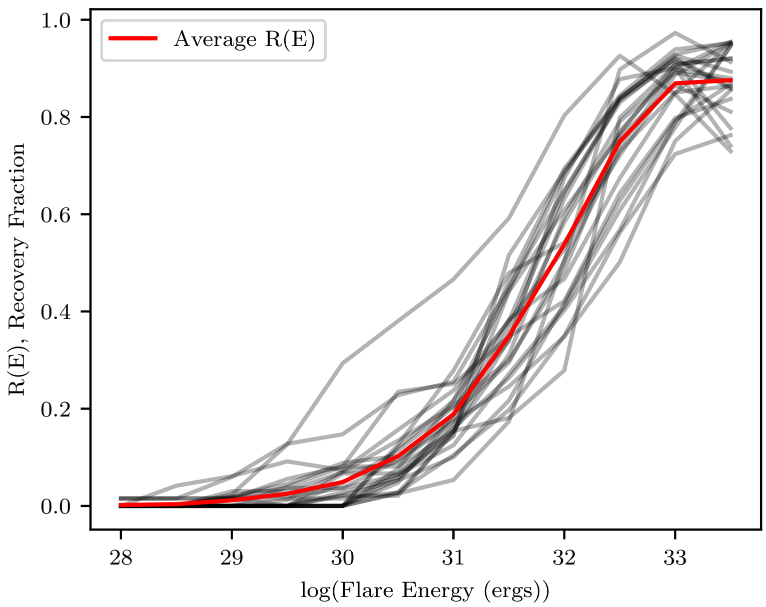

Figure D.1: The smoothed recovery fractions for all objects in the SSO data sample with spectral types M4–M5 are shown in black. The red line is the average recovery fraction across all stars.

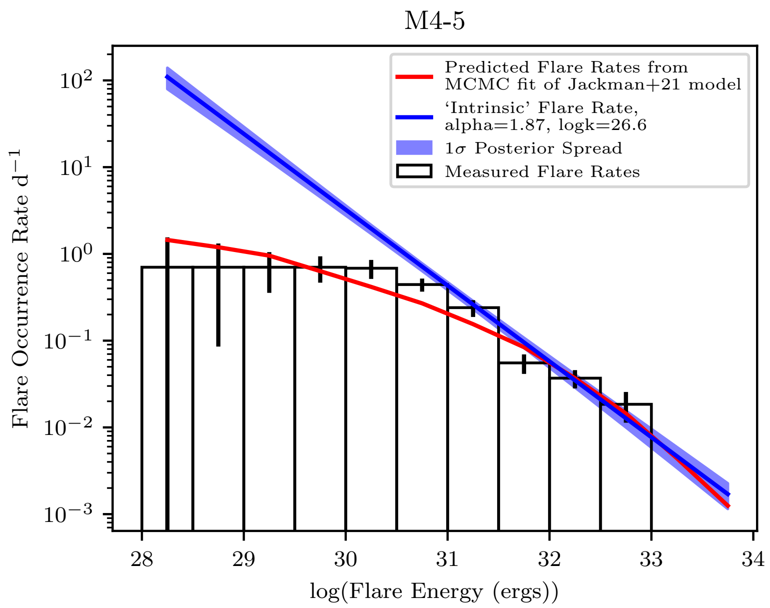

Figure D.2: The best MCMC fit of Equation 4.16 to the observed flare occurrence rates for the spectral type bin M4–M5. The measured flare rates from the SSO data sample, and corresponding errors, are shown in black. The best fit of Equation 4.16, which accounts for the incompleteness of the flare recovery process, to the observed occurrence rates is shown in red. The 'intrinsic' flare rate (and the 1σ posterior spread) is shown in blue.

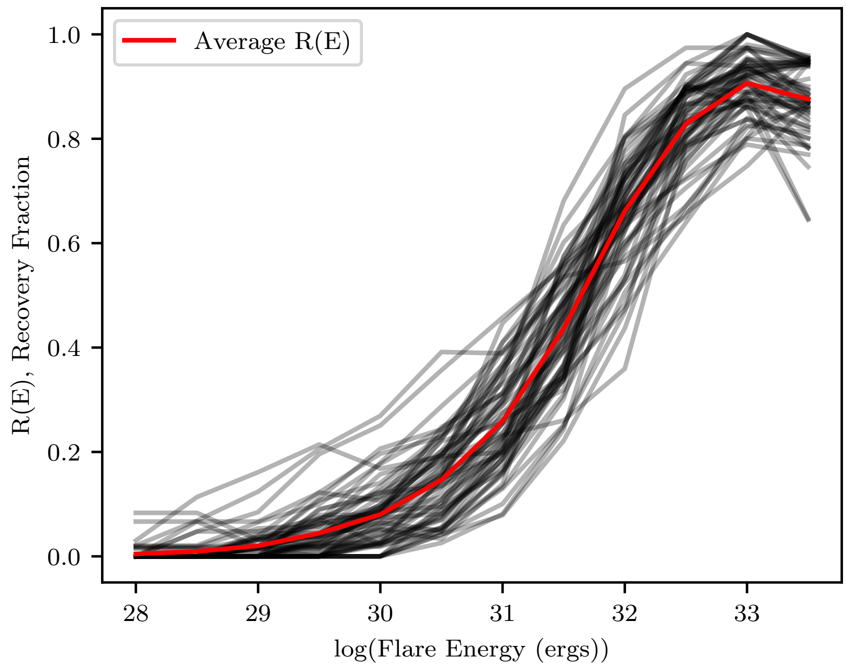

Figure D.3: The same plot as Figure D.1, however for spectral type M6.

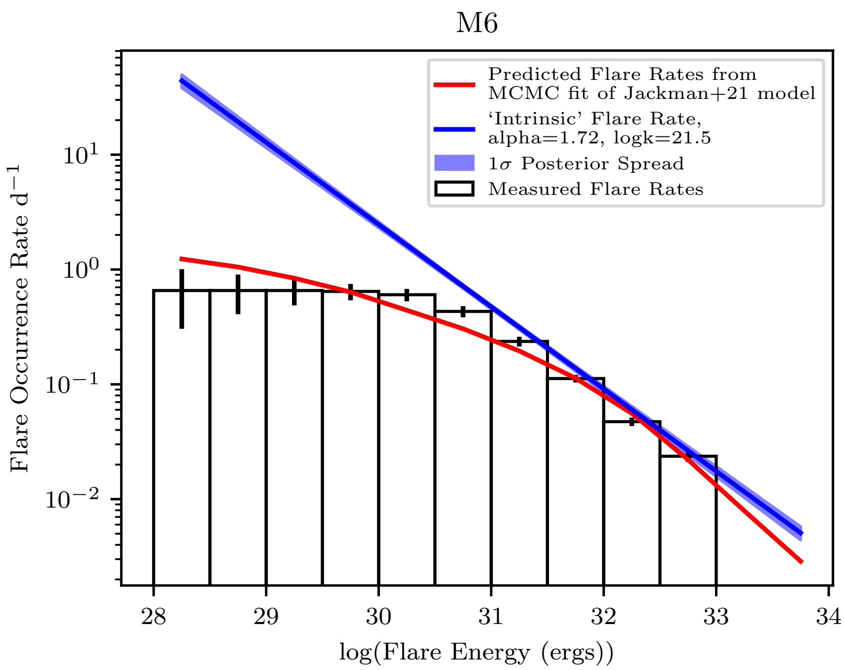

Figure D.4: The same plot as Figure D.2, however for spectral type M6.

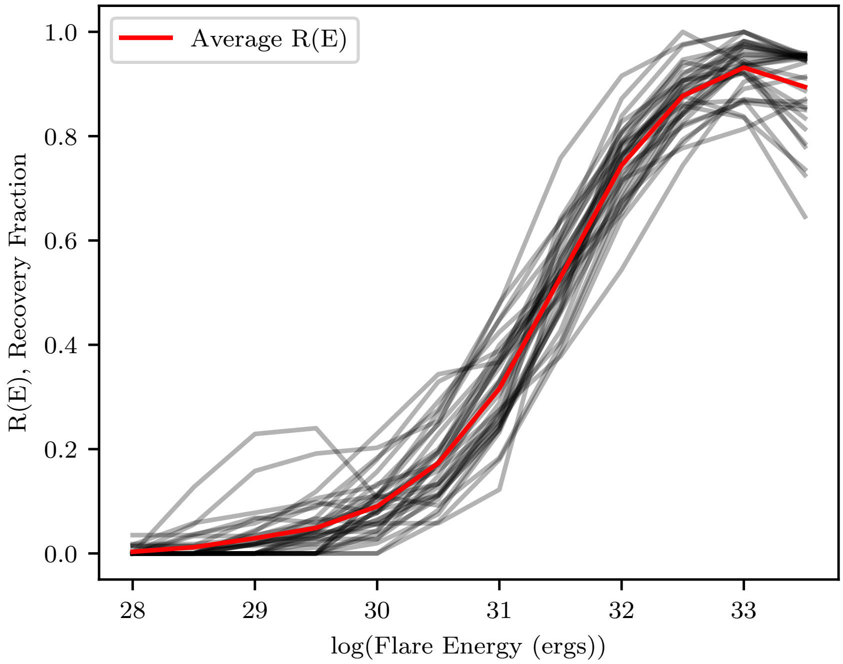

Figure D.5: The same plot as Figure D.1, however for spectral type M7.

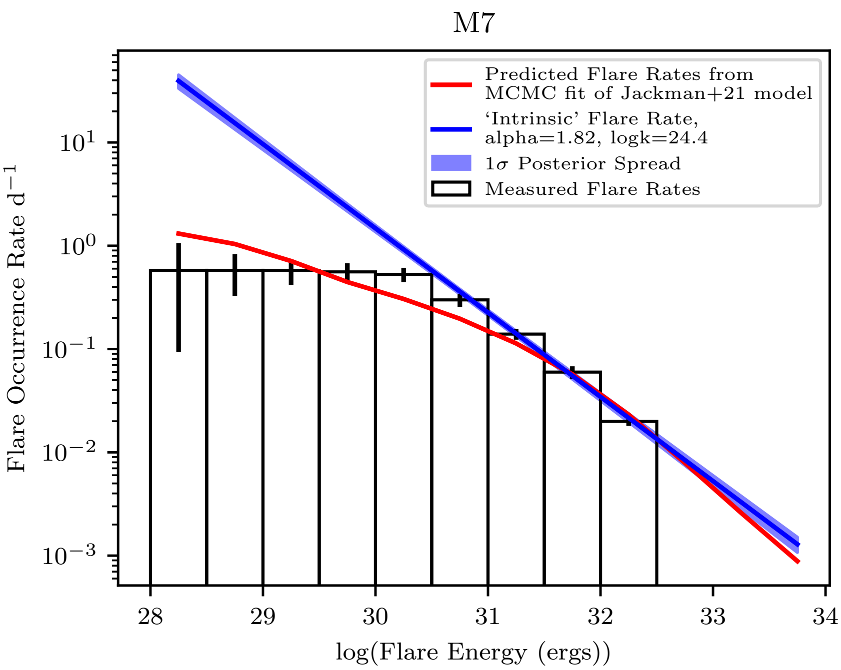

Figure D.6: The same plot as Figure D.2, however for spectral type M7.

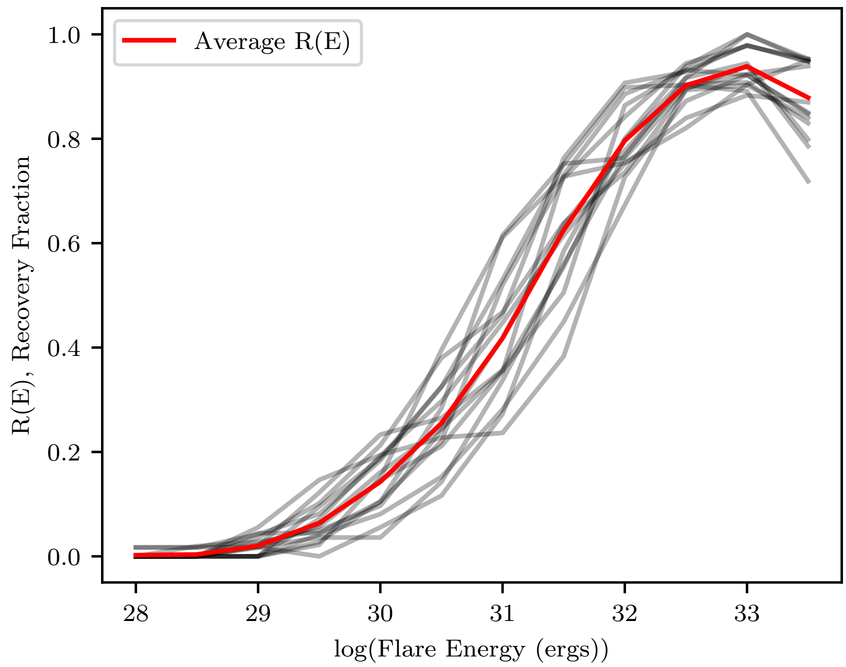

Figure D.7: The same plot as Figure D.1, however for spectral type M8.

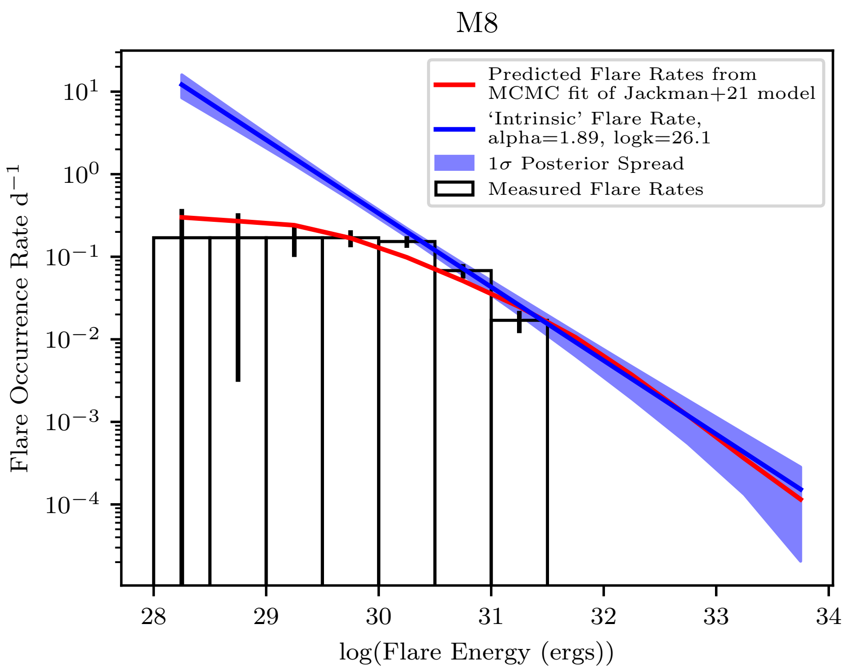

Figure D.8: The same plot as Figure D.2, however for spectral type M8.

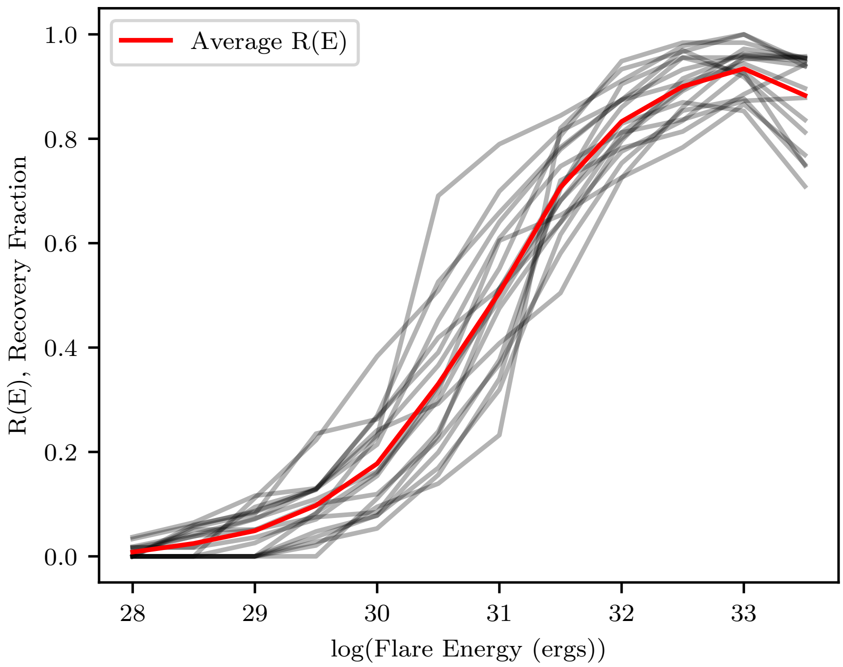

Figure D.9: The same plot as Figure D.1, however for spectral type M9–L0.

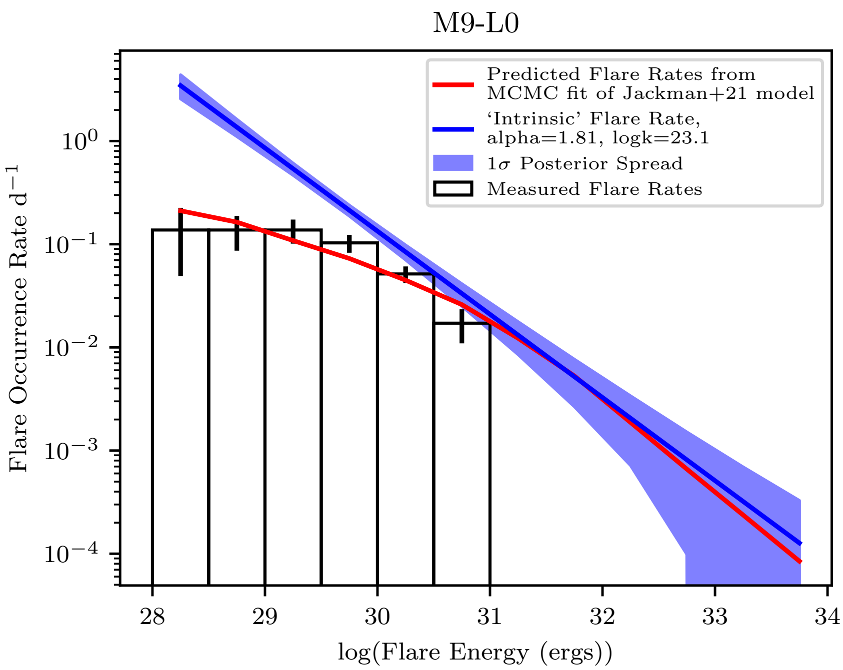

Figure D.10: The same plot as Figure D.2, however for spectral type M9–L0.

Funding and Data Acknowledgements

This research has been made possible through STFC studentship n° 302132223.

The SPECULOOS Project, whose data is central to these results, has received funding from the European Research Council (ERC) under the FP/2007–2013 ERC grant agreement n° 336480, and under the European Union's Horizon 2020 research and innovation programme (grants agreements n° 679030 & 803193/BEBOP); from an Actions de Recherche Concertée (ARC) grant, financed by the Wallonia–Brussels Federation, from the Balzan Prize Foundation, from the BELSPO/BRAIN2.0 research program (PORTAL project) and from F.R.S-FNRS (Research Project ID T010920F). This work was also partially supported by a grant from the Simons Foundation (PI Queloz, grant number 327127), as well as by the MERAC foundation (PI Triaud). PI Gillon is F.R.S.-FNRS Senior Research Associate. PI Demory acknowledges support from the Swiss National Science Foundation (PP00P2-163967 and PP00P2-190080). This research is also supported by funding from STFC grants n° ST/S00193X/1 and ST/S00305/1.

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.

This thesis makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by the National Aeronautics and Space Administration and the National Science Foundation.

Many PYTHON (Van Rossum & Drake 1995, 2009) packages have been used in the course of this work: MATPLOTLIB (Hunter 2007), NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), ASTROPY (The Astropy Collaboration et al. 2013, 2018), EMCEE (Foreman-Mackey et al. 2013), GEORGE (Ambikasaran et al. 2016). I have also used the CASUTOOLS (Irwin et al. 2004), FITSIO (`https://github.com/esheldon/fitsio`) - a PYTHON wrapper of CFITSIO (Pence 1999), and ASTROMETRY.NET (Lang et al. 2010) programs extensively in the SSO Pipeline.

Bibliography

Adams F. C., Bodenheimer P., Laughlin G., 2005, Astronomische Nachrichten, 326, 913

Agol E., Fabrycky D. C., 2018, Handbook of Exoplanets, pp 797–816

Agol E., Steffen J., Sari R., Clarkson W., 2005, Monthly Notices of the Royal Astronomical Society, 359, 567

Aigrain S., Pont F., Zucker S., 2012, Monthly Notices of the Royal Astronomical Society, 419, 3147

Aigrain S., Parviainen H., Pope B., 2016, Monthly Notices of the Royal Astronomical Society, 459, 2408

Allred J. C., Kowalski A. F., Carlsson M., 2015, Astrophysical Journal, 809

Alonso R., 2018, Handbook of Exoplanets, pp 1441–1467

Alonso R., et al., 2004, The Astrophysical Journal, 613, L153

Alsubai K. A., et al., 2014, Acta Astronomica, 63, 465

Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2016, IEEE Transactions on Pattern Analysis and Machine Intelligence, 38, 252

Andrae R., et al., 2018, Astronomy & Astrophysics, 616, A8

Anfinogentov S., Nakariakov V. M., Mathioudakis M., Van Doorsselaere T., Kowalski A. F., 2013, The Astrophysical Journal, 773, 156

Anglada-Escudé G., et al., 2016, Nature, 536, 437

Armstrong D. J., de Mooij E., Barstow J., Osborn H. P., Blake J., Saniee N. F., 2017, Nature Astronomy, 1, 0004

Artigau , et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. SPIE, p. 914715, doi:10.1117/12.2055663

Auvergne M., et al., 2009, Astronomy & Astrophysics, 506, 411

Bailer-Jones C., Lamm M., 2003, Monthly Notices of the Royal Astronomical Society, 339, 477

Baker A. D., Blake C. H., Sliski D. H., 2017, Publications of the Astronomical Society of the Pacific, 129, 085002

Bakos G., Noyes R. W., Kovács G., Stanek K. Z., Sasselov D. D., Domsa I., 2004, Publications of the Astronomical Society of the Pacific, 116, 266

Bakos G. , et al., 2013, Publications of the Astronomical Society of the Pacific, 125, 154

Ballard S., 2019, The Astronomical Journal, 157, 113

Ballard S., Johnson J. A., 2016, The Astrophysical Journal, 816, 66

Barclay T., Pepper J., Quintana E. V., 2018, The Astrophysical Journal Supplement Series, 239, 2

Bardalez Gagliuffi D. C., et al., 2019, The Astrophysical Journal, 883, 205

Barkaoui K., Gillon M., Benkhaldoun Z., Emmanuel J., Elhalkouj T., Daassou A., Burdanov A., Delrez L., 2017, Journal of Physics: Conference Series, 869, 012073

Barkaoui K., et al., 2019, The Astronomical Journal, 157, 43

Barnes R., Raymond S. N., Jackson B., Greenberg R., 2008, Astrobiology, 8, 557

Barros S. C. C., Demangeon O., Díaz R. F., Cabrera J., Santos N. C., Faria J. P., Pereira F., 2020, Astronomy & Astrophysics, 634, A75

Bashi D., Helled R., Zucker S., Mordasini C., 2017, Astronomy and Astrophysics, 604, A83

Batalha N. M., et al., 2013, The Astrophysical Journal Supplement Series, 204, 24

Beatty J. T., et al., 2005, Proceedings of the National Academy of Science, 102, 9306

Beaulieu J.-P., et al., 2006, Nature, 439, 437

Bennett D. P., et al., 2008, The Astrophysical Journal, 684, 663

Benz A. O., Güdel M., 2010, Annual Review of Astronomy and Astrophysics, 48, 241

Benz W., et al., 2020, Experimental Astronomy, 51, 109

Berta-Thompson Z. K., et al., 2015, Nature, 527, 204

Berta Z. K., Irwin J., Charbonneau D., Burke C. J., Falco E. E., 2012a, The Astronomical Journal, 144, 145

Berta Z. K., et al., 2012b, The Astrophysical Journal, 747, 35

Berta Z. K., Irwin J., Charbonneau D., Berta Z. K., Irwin J., Charbonneau D., 2013, The Astrophysical Journal, 775, 91

Bevis M., Businger S., Herring T. A., Rocken C., Anthes R. A., Ware R. H., 1992, Journal of Geophysical Research: Atmospheres, 97, 15787

Blake C. H., Shaw M. M., 2011, Publications of the Astronomical Society of the Pacific, 123, 1302

Blake C. H., Bloom J. S., Latham D. W., Szentgyorgyi A. H., Skrutskie M. F., Falco E. E., Starr D. S., 2008, Publications of the Astronomical Society of the Pacific, 120, 860

Bond I. A., et al., 2017, Monthly Notices of the Royal Astronomical Society, 469, 2434

Bonfanti A., et al., 2021, Astronomy and Astrophysics, 646

Bonfils X., et al., 2013, Astronomy & Astrophysics, 549, A109

Bonfils X., et al., 2018, Astronomy and Astrophysics, 618

Borsato L., et al., 2021, Monthly Notices of the Royal Astronomical Society, 506, 3810

Borucki W. J., et al., 2010, Science, 327, 977

Borucki W. J., et al., 2011, The Astrophysical Journal, 736, 19

Bouchy F., et al., 2005, Astronomy & Astrophysics, 444, L15

Bouchy F., et al., 2017, The Messenger, 169, 21

Bowler B. P., 2016, Publications of the Astronomical Society of the Pacific, 128, 102001

Broeg C., Fernández M., Neuhåuser R., 2005, Astronomische Nachrichten, 326, 134

Broeg C., et al., 2013, EPJ Web of Conferences, 47, 03005

Bryant E. M., et al., 2021, Monthly Notices of the Royal Astronomical Society: Letters, 504, L45

Buccino A. P., Lemarchand G. A., Mauas P. J., 2007, Icarus, 192, 582

Burdanov A., Delrez L., Gillon M., Jehin E., 2018, in Deeg H. J., Belmonte J. A., eds, , Handbook of Exoplanets. p. 130, doi:10.1007/978-3-319-55333-7_130

Byrd R. H., Lu P., Nocedal J., Zhu C., 1995, SIAM Journal on Scientific Computing, 16, 1190

Carrington R. C., 1859, Monthly Notices of the Royal Astronomical Society, 20, 13

Carroll B. W., Ostlie D. A., 1996, An Introduction to Modern Astrophysics. Addison-Wesley Pub

Carter J. A., et al., 2012, Science, 337, 556

Castro-Almazán J. A., Muñoz-Tuñón C., García-Lorenzo B., Pérez-Jordán G., Varela A. M., Romero I., 2016, Observatory Operations: Strategies, Processes, and Systems VI, 9910, 99100P

Chabrier G., 2003, Publications of the Astronomical Society of the Pacific, 115, 763

Chabrier G., Baraffe I., 1997, Astronomy and Astrophysics, 327, 1039

Charbonneau D., Brown T. M., Latham D. W., Mayor M., 2000, The Astrophysical Journal, 529, L45

Charbonneau D., et al., 2009, Nature, 462, 891

Chauvin G., 2018, in Schmidt D., Schreiber L., Close L. M., eds, Vol. 10703, Adaptive Optics Systems VI. SPIE, p. 5, doi:10.1117/12.2314273

Chauvin G., et al., 2004, Astronomy & Astrophysics, 425, L29

Chauvin G., et al., 2017, Astronomy & Astrophysics, 605, L9

Chen J., Kipping D., 2016, The Astrophysical Journal, 834, 17

Chen X.-H., Dempster A. P., Liu J. S., 1994, Biometrika, 81, 457

Christiansen J. L., et al., 2012, Publications of the Astronomical Society of the Pacific, 124, 1279

Claudi R., 2016, in , Vol. 428, Methods of Detecting Exoplanets. Springer, Cham, pp 183–252, doi:10.1007/978-3-319-27458-4_4

Claudi R., et al., 2021, Life, 11, 1

Cloutier R., Menou K., 2020, The Astronomical Journal, 159, 211

Collier Cameron A., et al., 2006, Monthly Notices of the Royal Astronomical Society, 373, 799

Cooke B. F., et al., 2020, Monthly Notices of the Royal Astronomical Society, 500, 5088

Covone G., Ienco R. M., Cacciapuoti L., Inno L., 2021, Monthly Notices of the Royal Astronomical Society, 505, 3329

Cowan N. B., Fujii Y., 2018, Handbook of Exoplanets, pp 1469–1484

Crossfield I. J. M., et al., 2019, The Astrophysical Journal Letters, 883, L16

Curiel S., Ortiz-León G. N., Mioduszewski A. J., Torres R. M., 2020, The Astronomical Journal, 160, 97

Davenport J. R. A., 2016, The Astrophysical Journal, 829

Davenport J. R., et al., 2014, Astrophysical Journal, 797, 122

Davenport J. R. A., Covey K. R., Clarke R. W., Boeck A. C., Cornet J., Hawley S. L., 2019, The Astrophysical Journal, 871, 241

Deleuil M., et al., 2018, Astronomy & Astrophysics, 619, A97

Delfosse X., et al., 2013, Astronomy & Astrophysics, 553, A8

Deline A., et al., 2020, Astronomy & Astrophysics, 635, A22

Delrez L., et al., 2018a, Monthly Notices of the Royal Astronomical Society, 475

Delrez L., et al., 2018b, Ground-based and Airborne Telescopes VII, 10700, 49

Deming D., 2008, Nature, 456, 714

Deming L. D., Seager S., 2017, Journal of Geophysical Research: Planets, 122, 53

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

Demory B.-O., Gillon M., Madhusudhan N., Queloz D., 2016a, Monthly Notices of the Royal Astronomical Society, 455, 2018

Demory B.-O., Queloz D., Alibert Y., Gillen E., Gillon M., 2016b, The Astrophysical Journal Letters, 825, L25

Demory B. O., et al., 2020, Astronomy and Astrophysics, 642, A49

Dieterich S. B., Henry T. J., Jao W.-C., Winters J. G., Hosey A. D., Riedel A. R., Subasavage J. P., 2014, The Astronomical Journal, 147, 94

Dittmann J. A., et al., 2017, Nature, 544, 333

Donati J.-F., et al., 2018, in , Handbook of Exoplanets. Springer, Cham, pp 903–929, doi:10.1007/978-3-319-55333-7_107

Dong S., Zhu Z., 2013, The Astrophysical Journal, 778, 53

Dressing C. D., Charbonneau D., 2013, The Astrophysical Journal, 767, 95

Dressing C. D., Charbonneau D., 2015, The Astrophysical Journal, 807, 45

Dressing C. D., et al., 2019, The Astronomical Journal, 158, 87

Duan J., et al., 1996, Journal of Applied Meteorology, 35, 830

Ducrot E., et al., 2018, The Astronomical Journal, 156, 218

Ducrot E., et al., 2020, Astronomy and Astrophysics, 640

Dunham E. W., Mandushev G. I., Taylor B. W., Oetiker B., 2004, Publications of the Astronomical Society of the Pacific, 116, 1072

Edgeworth K. E., 1946, Nature, 157, 481

Edwards B., et al., 2020, The Astronomical Journal, 161, 44

Faria J. P., et al., 2016, Astronomy & Astrophysics, 588, A31

Feliz D. L., Plavchan P., Bianco S. N., Jimenez M., Collins K. I., Alvarado B. V., Stassun K. G., 2021, The Astronomical Journal, 161, 247

Filippazzo J. C., Rice E. L., Faherty J., Cruz K. L., Gordon M. M. V., Looper D. L., 2015, The Astrophysical Journal, 810, 158

Fletcher R., 1987, Practical Methods of Optimization. John Wiley & Sons, Ltd, doi:10.1002/9781118723203

Fontanive C., et al., 2020, The Astrophysical Journal Letters, 905, L14

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306

Foreman-Mackey D., Hogg D. W., Morton T. D., 2014, Astrophysical Journal, 795, 64

Foreman-Mackey D., Agol E., Ambikasaran S., Angus R., Foreman-Mackey D., Agol E., Ambikasaran S., Angus R., 2017, The Astronomical Journal, 154, 220

Fortier A., Beck T., Benz W., Broeg C., Cessa V., Ehrenreich D., Thomas N., 2014, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, 9143, 91432J

Fressin F., et al., 2013, The Astrophysical Journal, 766, 81

Fridlund M., Baglin A., Lochard J., Conroy L., 2006, Technical report, The CoRoT Mission Pre-Launch Status - Stellar Seismology and Planet Finding

Furukawa Y., Nakazawa H., Sekine T., Kobayashi T., Kakegawa T., 2015, Earth and Planetary Science Letters, 429, 216

Gaia Collaboration et al., 2016a, Astronomy & Astrophysics, 595, A1

Gaia Collaboration et al., 2016b, Astronomy & Astrophysics, 595, A2

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

Gaia Collaboration et al., 2021, Astronomy & Astrophysics, 649, A1

Gaidos E., et al., 2014, Monthly Notices of the Royal Astronomical Society, 443, 2561

Gaidos E., Mann A. W., Kraus A. L., Ireland M., 2016, Monthly Notices of the Royal Astronomical Society, 457, 2877

García-Mejía J., Charbonneau D., Fabricant D., Irwin J. M., Fata R., Zajac J. M., Doherty P. E., 2020, in Marshall H. K., Spyromilio J., Usuda T., eds, Vol. 11445, Ground-based and

Airborne Telescopes VIII. SPIE, p. 168

Gardner J. P., et al., 2006, Space Science Reviews, 123, 485

Gaudi B. S., 2010, in Seager S. ed., , Exoplanets. University of Arizona Press, Tucson, AZ, pp 79–110

Gauza B., Béjar V. J. S., Pérez-Garrido A., Osorio M. R. Z., Lodieu N., Rebolo R., Pallé E., Nowak G., 2015, The Astrophysical Journal, 804, 96

Gershberg R. E., 1972, Astrophysics and Space Science, 19, 75

Giacobbe P., et al., 2012, Monthly Notices of the Royal Astronomical Society, 424, 3101

Gibbs A., et al., 2020, The Astronomical Journal, 159, 169

Gibson N. P., Aigrain S., Roberts S., Evans T. M., Osborne M., Pont F., 2012, Monthly Notices of the Royal Astronomical Society, 419, 2683

Gilbert E. A., et al., 2020, The Astronomical Journal, 160, 116

Gill S., et al., 2020a, Monthly Notices of the Royal Astronomical Society, 491, 1548

Gill S., et al., 2020b, The Astrophysical Journal Letters, 898, L11

Gilliland R. L., et al., 2011, The Astrophysical Journal Supplement Series, 197, 6

Gillon M., Jehin E., Magain P., Chantry V., Hutsemékers D., Manfroid J., Queloz D., Udry S., 2011, EPJ Web of Conferences, 11, 06002

Gillon M., Jehin E., Fumel A., Magain P., Queloz D., 2013, EPJ Web of Conferences, 47, 03001

Gillon M., et al., 2016, Nature, 533, 221

Gillon M., et al., 2017, Nature, 542, 456

Gillon M., et al., 2018, in Gilmozzi R., Marshall H. K., Spyromilio J., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 10700, Ground-based and Airborne Telescopes VII. SPIE, p. 49, doi:10.1117/12.2312475

Gizis J. E., Burgasser A. J., Berger E., Williams P. K. G., Vrba F. J., Cruz K. L., Metchev S., 2013, Astrophysical Journal, 779

Gizis J. E., Paudel R. R., Schmidt S. J., Williams P. K. G., Burgasser A. J., 2017, The Astrophysical Journal, 838, 22

Glazier A. L., Howard W. S., Corbett H., Law N. M., Ratzloff J. K., Fors O., Ser D. d., 2020, The Astrophysical Journal, 900, 27

Gould A., Loeb A., 1992, The Astrophysical Journal, 396, 104

Gould A., et al., 2010, The Astrophysical Journal, 710, 1800

Greene T. P., Line M. R., Montero C., Fortney J. J., Lustig-Yaeger J., Luther K., 2016, The Astrophysical Journal, 817, 17

Grunblatt S. K., Howard A. W., Haywood R. D., Grunblatt S. K., Howard A. W., Haywood R. D., 2015, The Astrophysical Journal, 808, 127

Guillot T., 2005, Annual Review of Earth and Planetary Sciences, 33, 493

Günther M. N., et al., 2018, Monthly Notices of the Royal Astronomical Society, 478, 4720

Günther M. N., et al., 2020a, Submitted to AAS journals

Günther M. N., et al., 2020b, The Astronomical Journal, 159, 60

Hadden S., Lithwick Y., 2014, The Astrophysical Journal, 787, 80

Han C., et al., 2013, The Astrophysical Journal, 778, 38

Han C., et al., 2016, The Astronomical Journal, 152, 95

Han C., et al., 2020, Astronomy & Astrophysics, 641, A105

Hardegree-Ullman K. K., Cushing M. C., Muirhead P. S., Christiansen J. L., 2019, The Astronomical Journal, 158, 75

Harris C. R., et al., 2020, Nature, 585, 357

Hawley S. L., Davenport J. R., Kowalski A. F., Wisniewski J. P., Hebb L., Deitrick R., Hilton E. J., 2014, Astrophysical Journal, 797, 121

Haywood R. D., et al., 2014, Monthly Notices of the Royal Astronomical Society, 443, 2517

He M. Y., Triaud A. H. M. J., Gillon M., 2017, Monthly Notices of the Royal Astronomical Society, 464, 2687

Henry T. J., Jao W.-C., Subasavage J. P., Beaulieu T. D., Ianna P. A., Costa E., Méndez R. A., 2006, The Astronomical Journal, 132, 2360

Henry T. J., et al., 2018, The Astronomical Journal, 155, 265

Heyl J. S., Gladman B. J., 2007, Monthly Notices of the Royal Astronomical Society, 377, 1511

Hilton E. J., West A. A., Hawley S. L., Kowalski A. F., 2010, The Astronomical Journal, 140, 1402

Hippke M., Heller R., 2019, Astronomy & Astrophysics, 623

Hodgson R., 1859, Monthly Notices of the Royal Astronomical Society, 20, 15

Holman M. J., Murray N. W., 2005, Science, 307, 1288

Howard A. W., et al., 2012, The Astrophysical Journal Supplement Series, 201, 15

Howard W. S., Corbett H., Law N. M., Ratzloff J. K., Glazier A., Fors O., Ser D. d., Haislip J., 2019, The Astrophysical Journal, 881, 9

Howard W. S., et al., 2020a, The Astrophysical Journal, 895, 140

Howard W. S., et al., 2020b, The Astrophysical Journal, 902, 115

Howell S. B., et al., 2014, Publications of the Astronomical Society of the Pacific, 126, 398

Hsu D. C., Ford E. B., Ragozzine D., Ashby K., 2019, The Astronomical Journal, 158, 109

Hsu D. C., Ford E. B., Terrien R., 2020, Monthly Notices of the Royal Astronomical Society, 498, 2249

Hunter J. D., 2007, Computing in Science and Engineering, 9, 90

Husser T.-O., Berg S. W.-v., Dreizler S., Homeier D., Reiners A., Barman T., Hauschildt P. H., 2013, Astronomy & Astrophysics, 553, A6

Hwang K.-H., et al., 2018, The Astronomical Journal, 157, 23

Ilin E., Schmidt S. J., Davenport J. R. A., Strassmeier K. G., 2019, Astronomy & Astrophysics, 622, 133

Ilin E., et al., 2021, Monthly Notices of the Royal Astronomical Society, 507, 1723

Irwin M. J., et al., 2004, Optimizing Scientific Return for Astronomy through Information Technologies, 5493, 411

Irwin J., Charbonneau D., Nutzman P., Falco E., 2008, Proceedings of the International Astronomical Union, 4, 37

Irwin J., Berta Z. K., Burke C. J., Charbonneau D., Nutzman P., West A. A., Falco E. E., 2011, The Astrophysical Journal, 727, 56

Jackman J. A., et al., 2019, Monthly Notices of the Royal Astronomical Society: Letters, 485, L136

Jackman J. A. G., et al., 2021, Monthly Notices of the Royal Astronomical Society, 000, 1

Jeffers S. V., et al., 2018, Astronomy & Astrophysics, 614, A76

Jehin E., et al., 2011, The Messenger, 145, 2

Jehin E., et al., 2018, The Messenger, 174, 2

Jones A., Noll S., Kausch W., Szyszka C., Kimeswenger S., 2013, Astronomy & Astrophysics, 560, A91

Jontof-Hutter D., et al., 2016, The Astrophysical Journal, 820, 39

Jung Y. K., et al., 2018a, The Astronomical Journal, 155, 219

Jung Y. K., et al., 2018b, The Astronomical Journal, 156, 208

Kaltenegger L., Traub W. A., 2009, Astrophysical Journal, 698, 519

Kasting J. F., Whitmire D. P., Reynolds R. T., 1993, Icarus, 101, 108

Kerber F., et al., 2010, in Stepp L. M., Gilmozzi R., Hall H. J., eds, Vol. 7733, Ground-based and Airborne Telescopes III. International Society for Optics and Photonics, p. 77331M, doi:10.1117/12.856390

Kerber F., et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. p. 84463N, doi:10.1117/12.924340

Kerber F., Querel R. R., Hanuschik R., 2014, Observatory Operations: Strategies, Processes, and Systems V, 9149, 91490M

King I. R., 1971, Publications of the Astronomical Society of the Pacific, 83, 199

Kipping D. M., 2008, Monthly Notices of the Royal Astronomical Society, 389, 1383

Kirkpatrick J. D., 2005, Annual Review of Astronomy & Astrophysics, 43, 195

Kirkpatrick J. D., Henry T. J., Simons D. A., 1995, The Astronomical Journal, 109, 797

Kirkpatrick J. D., et al., 1999, The Astrophysical Journal, 519, 802

Kirkpatrick J. D., et al., 2012, The Astrophysical Journal, 753, 156

Kondo I., et al., 2019, The Astronomical Journal, 158, 224

Kopparapu R. K., et al., 2013, The Astrophysical Journal, 765, 131

Kotani T., et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. SPIE, p. 914714, doi:10.1117/12.2055075

Kovács G., Zucker S., Mazeh T., 2002, Astronomy & Astrophysics, 391, 369

Kowalski A. F., Hawley S. L., Wisniewski J. P., Osten R. A., Hilton E. J., Holtzman J. A., Schmidt S. J., Davenport J. R., 2013, Astrophysical Journal, Supplement Series, 207, 15

Kreidberg L., 2018, Handbook of Exoplanets, pp 2083–2105

Kreidberg L., et al., 2014, Nature, 505, 69

Krijt S., Bowling T. J., Lyons R. J., Ciesla F. J., 2017, The Astrophysical Journal, 839, L21

Kubas D., et al., 2012, Astronomy & Astrophysics, 540, A78

Kunimoto M., Bryson S., 2020, Research Notes of the AAS, 4, 83

Kuzuhara M., et al., 2013, The Astrophysical Journal, 774, 11

Lacy C. H., Moffett T. J., Evans D. S., 1976, Astrophysical Journal Supplement Series, 30, 85

Lammer H., et al., 2007, Astrobiology, 7, 185

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

Law N. M., et al., 2015, Publications of the Astronomical Society of the Pacific, 127, 234

Lehmer O. R., Catling D. C., Parenteau M. N., Hoehler T. M., 2018, The Astrophysical Journal, 859, 171

Leleu A., et al., 2021, Astronomy and Astrophysics, 649, 14

Lendl M., et al., 2020, Monthly Notices of the Royal Astronomical Society, 492, 1761

Li T., et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. International Society for Optics and Photonics, p. 84462L, doi:10.1117/12.924792

Li T., DePoy D. L., Marshall J. L., Nagasawa D. Q., Carona D. W., Boada S., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. International Society for Optics and Photonics, p. 91476Z, doi:10.1117/12.2055167

Li D., Blake C. H., Nidever D., Halverson S. P., 2017, Publications of the Astronomical Society of the Pacific, 130, 014501

Liebes S., 1964, Physical Review, 133, B835

Lienhard F., et al., 2020, Monthly Notices of the Royal Astronomical Society, 497, 3790

Lillo-Box J., et al., 2020, Astronomy and Astrophysics, 642

Lin C.-L., et al., 2021, The Astronomical Journal, 162, 11

Lingam M., Loeb A., 2019a, Monthly Notices of the Royal Astronomical Society, 485, 5924

Lingam M., Loeb A., 2019b, The Astrophysical Journal, 883, 143

Lissauer J. J., et al., 2011, The Astrophysical Journal Supplement Series, 197, 8

Liu M. C., et al., 2011, Astrophysical Journal, 740

Lodato G., Delgado-Donate E., Clarke C. J., 2005, Monthly Notices of the Royal Astronomical Society: Letters, 364, L91

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

Louie D. R., Deming D., Albert L., Bouma L. G., Bean J., Lopez-Morales M., 2018, Publications of the Astronomical Society of the Pacific, 130, 044401

Lovis C., Fischer D. A., 2010, in Seager S., ed., , Exoplanets. University of Arizona Press, Tucson, pp 27–53

Luger R., et al., 2017, Nature Astronomy, 1, 0129

Luger R., Foreman-Mackey D., Hedges C., Luger R., Foreman-Mackey D., Hedges C., 2021, The Astronomical Journal, 162, 124

Luque R., et al., 2018, Astronomy & Astrophysics, 620, A171

Lurie J. C., Davenport J. R., Hawley S. L., Wilkinson T. D., Wisniewski J. P., Kowalski A. F., Hebb L., 2015, Astrophysical Journal, 800, 95

Macintosh B., et al., 2015, Science, 350, 64

Mahadevan S., et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. SPIE, p. 91471G, doi:10.1117/12.2056417

Mandel K., Agol E., 2002, The Astrophysical Journal, 580, L171

Mann A. W., et al., 2019, The Astrophysical Journal, 871, 63

Mao S., Paczynski B., Mao S., Paczynski B., 1991, The Astrophysical Journal Letters, 374, L37

Martínez R. R., Lopez L. A., Shappee B. J., Schmidt S. J., Jayasinghe T., Kochanek C. S., Auchettl K., Holoien T. W. S., 2019, The Astrophysical Journal, 892, 144

Mayor M., Queloz D., 1995, Nature, 378, 355

Mayor M., et al., 2003, The Messenger, 114, 20

McCormac J., Pollacco D., Skillen I., Faedi F., Todd I., Watson C. A., 2013, Publications of the Astronomical Society of the Pacific, 125, 548

McCullough P. R., Stys J. E., Valenti J. A., Fleming S. W., Janes K. A., Heasley J. N., 2005, Publications of the Astronomical Society of the Pacific, 117, 783

Medina A. A., Winters J. G., Irwin J. M., Charbonneau D., 2020, The Astrophysical Journal, 905, 107

Ment K., et al., 2020, The Astronomical Journal, 161, 23

Merline W. J., Howell S. B., 1995, Experimental Astronomy, 6, 163

Mighell K., 1999, in ASP Conference Series. pp 50–55

Miguel Y., Cridland A., Ormel C. W., Fortney J. J., Ida S., 2019, Monthly Notices of the Royal Astronomical Society, 491, 1998

Miyazaki S., et al., 2018, The Astronomical Journal, 156, 136

Moffett T. J., 1974, Technical Report 273, UV CETI FLARE STARS: OBSERVATIONAL DATA

Mondrik N., Newton E., Charbonneau D., Irwin J., 2018, The Astrophysical Journal, 870, 10

Morales J. C., et al., 2019, Science, 365, 1441

Morgan W. W., Keenan P. C., 1973, Annual Review of Astronomy and Astrophysics, vol. 11, p.29, 11, 29

Morton T. D., Swift J., 2014, The Astrophysical Journal, 791, 10

Muirhead P. S., et al., 2015, The Astrophysical Journal, 801, 18

Mulders G. D., 2018, Handbook of Exoplanets, pp 1–26

Mulders G. D., Pascucci I., Apai D., 2015, The Astrophysical Journal, 798, 112

Mullan D. J., Bais H. P., 2018, The Astrophysical Journal, 865, 101

Murray C. D., Correia A. C. M., 2010, in Seager S. ed., , Vol. 526, Exoplanets. University of Arizona Press, Tucson, pp 15–23

Murray C. A., et al., 2020, Monthly Notices of the Royal Astronomical Society, 495, 2446

Murray C. A., et al., 2022, Monthly Notices of the Royal Astronomical Society, 513, 2615

Newton E. R., Irwin J., Charbonneau D., Berta-Thomspon Z. K., West A. A., 2016, in Proceedings of the International Astronomical Union. Cambridge University Press, pp 124–125, doi:10.1017/S174392131500602X

Newton E. R., Irwin J., Charbonneau D., Berlind P., Calkins M. L., Mink J., 2017, The Astrophysical Journal, 834, 85

Newton E. R., Mondrik N., Irwin J., Winters J. G., Charbonneau D., 2018, The Astronomical Journal, 156, 217

Nielsen L. D., et al., 2019, Monthly Notices of the Royal Astronomical Society, 489, 2478

Niraula P., et al., 2020, The Astronomical Journal, 160, 172

Noll S., Kausch W., Barden M., Jones A. M., Szyszka C., Kimeswenger S., Vinther J., 2012, Astronomy & Astrophysics, 543, A92

Nowak G., et al., 2020, Astronomy and Astrophysics, 642

Noyes R. W., Hartmann L. W., Baliunas S. L., Duncan D. K., Vaughan A. H., 1984, Astrophysical Journal, 279, 763

Nutzman P., Charbonneau D., 2008, Publications of the Astronomical Society of the Pacific, 120, 317

O'brien S. M., et al., 2021, Monthly Notices of the Royal Astronomical Society, 509, 6111

Osborn J., Föhring D., Dhillon V. S., Wilson R. W., 2015, Monthly Notices of the Royal Astronomical Society, 452, 1707

Paczynski B., 1986, The Astrophysical Journal, 304, 1

Paczyński B., 1996, Annual Review of Astronomy and Astrophysics, 34, 419

Parmentier V., Crossfield I. J. M., 2018, Handbook of Exoplanets, pp 1419–1440

Parviainen H., 2015, Monthly Notices of the Royal Astronomical Society, 450, 3233

Patel B. H., Percivalle C., Ritson D. J., Duffy C. D., Sutherland J. D., 2015, Nature Chemistry, 7, 301

Paudel R. R., Gizis J. E., Mullan D. J., Schmidt S. J., Burgasser A. J., Williams P. K. G., Berger E., 2018, The Astrophysical Journal, 858, 55

Paudel R. R., Gizis J. E., Mullan D. J., Schmidt S. J., Burgasser A. J., Williams P. K., Youngblood A., Stassun K. G., 2019, Monthly Notices of the Royal Astronomical Society, 486, 1438

Pecaut M. J., Mamajek E. E., 2013, The Astrophysical Journal Supplement Series, 208, 9

Pence W., 1999, in Mehringer D. M., Plante R. L., Roberts D. A., eds, ASP Conference Series. p. 487

Pepe F., et al., 2000, Optical and IR Telescope Instrumentation and Detectors, 4008, 582

Pepe F., et al., 2002, The Messenger, 110, 9

Pepper J., et al., 2007, Publications of the Astronomical Society of the Pacific, 119, 923

Perryman M., 2018, The Exoplanet Handbook, second edition edn

Petigura E. A., Howard A. W., Marcy G. W., 2013, Proceedings of the National Academy of Sciences, 110, 19273

Pollacco D. L., et al., 2006, Publications of the Astronomical Society of the Pacific, 118, 1407

Pont F., Zucker S., Queloz D., 2006, Monthly Notices of the Royal Astronomical Society, 373, 231

Querel R. R., Kerber F., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. International Society for Optics and Photonics, p. 914792, doi:10.1117/12.2055500

Quinn S. N., et al., 2012, The Astrophysical Journal Letters, 756, L33

Quirrenbach A., 2010, in Seager S., ed., , Exoplanets. University of Arizona Press, Tuscon, pp 157–174

Quirrenbach A., Consortium C., 2020, in Evans C. J., Bryant J. J., Motohara K., eds, Vol. 11447, Ground-based and Airborne Instrumentation for Astronomy VIII. SPIE, p. 262, doi:10.1117/12.2561380

Quirrenbach A., et al., 2010, in McLean I. S., Ramsay S. K., Takami H., eds, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III. International Society for

Optics and Photonics, p. 773513, doi:10.1117/12.857777

Quirrenbach A., et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. International Society for Optics and Photonics, p. 91471F, doi:10.1117/12.2056453

Quirrenbach A., et al., 2016, in Evans C. J., Simard L., Takami H., eds, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI. SPIE, p. 990812, doi:10.1117/12.2231880

Rackham B. V., Apai D., Giampapa M. S., 2018, The Astrophysical Journal, 853, 122

Rackham B. V., Apai D., Giampapa M. S., 2019, The Astronomical Journal, 157, 96

Raetz S., Stelzer B., Damasso M., Scholz A., 2020, Astronomy & Astrophysics, 637, A22

Rajpaul V., Aigrain S., Osborne M. A., Reece S., Roberts S., 2015, Monthly Notices of the Royal Astronomical Society, 452, 2269

Ranjan S., Wordsworth R., Sasselov D. D., 2017, The Astrophysical Journal, 843, 110

Rasmussen C. E., Williams C. K. I., 2006, MIT Press, pp 63–71

Rauer H., et al., 2013, Experimental Astronomy, 38, 249

Rebolo R., Osorio M. R. Z., Martín E. L., 1995, Nature, 377, 129

Refsdal S., Bondi H., 1964, Monthly Notices of the Royal Astronomical Society, 128, 295

Reid I. N., Cruz K. L., 2002, The Astronomical Journal, 123, 2806

Reid N., Hawley S., 2006, in , New Light on Dark Stars. Springer Berlin Heidelberg, pp 101–146, doi:10.1007/3-540-27610-6_3

Reiners A., et al., 2018, Astronomy & Astrophysics, 612, A49

Reylé C., Jardine K., Fouqué P., Caballero J. A., Smart R. L., Sozzetti A., 2021, Astronomy & Astrophysics, 650, A201

Ribas I., et al., 2018, Nature, 563, 365

Ricker G. R., et al., 2015, Journal of Astronomical Telescopes, Instruments, and Systems, 1, 14003

Rimmer P. B., Shorttle O., 2019, Life, 9, 12

Rimmer P. B., Xu J., Thompson S. J., Gillen E., Sutherland J. D., Queloz D., 2018, Science Advances, 4, eaar3302

Rodriguez J. E., et al., 2020, The Astronomical Journal, 160, 117

Ryu Y.-H., et al., 2019, The Astronomical Journal, 158, 151

Sabotta S., et al., 2021, Astronomy & Astrophysics

Sackett P. D., 1998, Planets Outside the Solar System: Theory and Observations, pp 189–228

Sagear S. A., Skinner J. N., Muirhead P. S., 2020, The Astronomical Journal, 160, 19

Sahlmann J., Lazorenko P. F., Ségransan D., Martín E. L., Queloz D., Mayor M., Udry S., 2013, Astronomy & Astrophysics, 556, A133

Salpeter E. E., Salpeter E. E., 1955, The Astrophysical Journal, 121, 161

Sarkar S., Argyriou I., Vandenbussche B., Papageorgiou A., Pascale E., 2018, Monthly Notices of the Royal Astronomical Society, 481, 2871

Scalo J., et al., 2007, Astrobiology, 7, 85

Scargle J. D., 1982, The Astrophysical Journal, 263, 835

Schmidt S. J., et al., 2019, The Astrophysical Journal, 876, 115

Seager S., 2010, Exoplanets. University of Arizona Press, Tucson, `https://uapress.arizona.edu/book/exoplanets`

Seager S., Mallen-Ornelas G., 2003, The Astrophysical Journal, 585, 1038

Seager S., Kuchner M., Hier-Majumder C. A., Militzer B., 2007, The Astrophysical Journal, 669, 1279

Seager S., Deming D., Valenti J. A., 2009, Astrophysics and Space Science Proceedings, 10, 123

Seager S., Bains W., Petkowski J., 2016, Astrobiology, 16, 465

Sebastian D., et al., 2020, in Marshall H. K., Spyromilio J., Usuda T., eds, Vol. 11445, Ground-based and Airborne Telescopes VIII. SPIE, p. 276, doi:10.1117/12.2563563

Sebastian D., et al., 2021, Astronomy & Astrophysics, 645, A100

Segura A., Walkowicz L. M., Meadows V., Kasting J., Hawley S., 2010, Astrobiology, 10, 751

Seli B., Vida K., Moór A., Pál A., Oláh K., 2021, Astronomy & Astrophysics, 650, A138

Serrano L. M., et al., 2018, Astronomy & Astrophysics, 611, A8

Sestovic M., Demory B.-O., 2020, Astronomy & Astrophysics, 641, A170

Shah B. K., 1973, The American Statistician, 27, 123

Shappee B. J., et al., 2014, Astrophysical Journal, 788, 48

Shibayama T., et al., 2013, Astrophysical Journal, Supplement Series, 209, 5

Shvartzvald Y., et al., 2017, The Astrophysical Journal Letters, 840, L3

Silburt A., Gaidos E., Wu Y., 2015, The Astrophysical Journal, 799, 180

Silverberg S. M., Kowalski A. F., Davenport J. R. A., Wisniewski J. P., Hawley S. L., Hilton E. J., 2016, The Astrophysical Journal, 829, 129

Skrutskie M. F., et al., 2006, The Astronomical Journal, 131, 1163

Skumanich A., 1972, The Astrophysical Journal, 171, 565

Skumanich A., 1986, The Astrophysical Journal, 309, 858

Smith G. D., et al., 2021a, Monthly Notices of the Royal Astronomical Society, 507, 5991

Smith A. M. S., et al., 2021b, Astronomy & Astrophysics, 646, A183

Sozzetti A., 2005, Publications of the Astronomical Society of the Pacific, 117, 1021

Sozzetti A., et al., 2013, EPJ Web of Conferences, 47

Stelzer B., Damasso M., Scholz A., Matt S. P., 2016, Monthly Notices of the Royal Astronomical Society, 463, 1844

Stevenson K. B., et al., 2014, Science, 346, 838

Sullivan P. W., et al., 2015, The Astrophysical Journal, 809, 77

Sumi T., et al., 2016, The Astrophysical Journal, 825, 112

Sutherland J. D., 2017, Nature Reviews Chemistry, 1, 0012

Swain M. R., Vasisht G., Tinetti G., Bouwman J., Chen P., Yung Y., Deming D., Deroo P., 2008, The Astrophysical Journal, 690, L114

Swain M. R., et al., 2009, The Astrophysical Journal, 704, 1616

Takizawa K., Minagawa J., Tamura M., Kusakabe N., Narita N., 2017, Scientific Reports, 7, 1

Talens G. J. J., Spronck J. F. P., Lesage A.-L., Otten G. P. P. L., Stuik R., Pollacco D., Snellen I. A. G., 2017, Astronomy & Astrophysics, 601, A11

Tamburo P., Muirhead P. S., 2019, Publications of the Astronomical Society of the Pacific, 131, 114401

The Astropy Collaboration et al., 2013, Astronomy & Astrophysics, 558, A33

The Astropy Collaboration et al., 2018, The Astronomical Journal, 156, 123

The LUVOIR Team 2019, Technical report, The LUVOIR Mission Concept Study Final Report

Tilley M. A., Segura A., Meadows V., Hawley S., Davenport J., 2019, Astrobiology, 19, 64

Tinetti G., et al., 2007, Nature, 448, 169

Tinetti G., et al., 2018, Experimental Astronomy 2018 46:1, 46, 135

Todorov K., Luhman K. L., McLeod K. K., 2010, The Astrophysical Journal Letters, 714, L84

Tozzi A., et al., 2016, in Evans C. J., Simard L., Takami H., eds, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI. SPIE, p. 99086C, doi:10.1117/12.2231898

Traub W., Oppenheimer B., 2010, in Seager S., ed., , Exoplanets. University of Arizona Press, Tucson, pp 111–156

Triaud A. H., et al., 2020, Nature Astronomy, 4, 650

Trifonov T., Rybizki J., Kürster M., 2019, Astronomy & Astrophysics, 622, L7

Trifonov T., et al., 2021, Science, 371, 1038

Udalski A., Zebruń K., Szymański M., Kubiak M., Soszyński I., Szewczyk O., Wyrzykowski Pietrzyński G., 2002, Acta Astronomica, 52, 115

Van Rossum G., Drake F. L., 1995, Python reference manual

Van Rossum G., Drake F. L., 2009, Python 3 Reference Manual. CreateSpace, `https://www.python.org/`

Vanderspek R., et al., 2019, The Astrophysical Journal Letters, 871, L24

Vida K., Kővári Z., Pál A., Oláh K., Kriskovics L., 2017, The Astrophysical Journal, 841, 124

Virtanen P., et al., 2020, Nature Methods, 17, 261

Wang Y. H., 1993, Statistica Sinica, 3, 295

Weiss L. M., et al., 2013, The Astrophysical Journal, 768

Wells R. D., et al., 2021, Astronomy & Astrophysics, 653

West A. A., et al., 2004, The Astronomical Journal, 128, 426

West A. A., Hawley S. L., Bochanski J. J., Covey K. R., Reid I. N., Dhital S., Hilton E. J., Masuda M., 2008, The Astronomical Journal, 135, 785

West A. A., Weisenburger K. L., Irwin J., Berta-Thompson Z. K., Charbonneau D., Dittmann J., Pineda J. S., 2015, Astrophysical Journal, 812

West R. G., et al., 2019, Monthly Notices of the Royal Astronomical Society, 486, 5094

Wheatley P. J., et al., 2018, Monthly Notices of the Royal Astronomical Society, 475, 4476

Williams P. K., Berger E., Irwin J., Berta-Thompson Z. K., Charbonneau D., 2015, Astrophysical Journal, 799, 192

Winn J. N., 2010, in , Exoplanets

Wolszczan A., Frail D. A., 1992, Nature, 355, 145

Wright N. J., Drake J. J., 2016, Nature, 535, 526

Wright N. J., Newton E. R., Williams P. K., Drake J. J., Yadav R. K., 2018, Monthly Notices of the Royal Astronomical Society, 479, 2351

Xu J., Ritson D. J., Ranjan S., Todd Z. R., Sasselov D. D., Sutherland J. D., 2018, Chemical Communications, 54, 5566

Yang H., Liu J., 2019, The Astrophysical Journal Supplement Series, 241, 29

Yang J., Liu Y., Hu Y., Abbot D. S., 2014, The Astrophysical Journal, 796

Yang H., et al., 2017, The Astrophysical Journal, 849, 36

Yang J.-Y., Xie J.-W., Zhou J.-L., 2020, The Astronomical Journal, 159, 164

Youdin A. N., 2011, The Astrophysical Journal, 742, 38

Young A. T., 1967, The Astronomical Journal, 72, 747

Zapolsky H. S., Salpeter E. E., 1969, The Astrophysical Journal, 158, 809

Zechmeister M., et al., 2019, Astronomy & Astrophysics, 627, A49

Zhang Z., Zhou Y., Rackham B. V., Apai D., 2018, The Astronomical Journal, 156, 178

Zhang X., et al., 2020, The Astronomical Journal, 159, 116

Zhu C., Byrd R. H., Lu P., Nocedal J., 1997, ACM Transactions on Mathematical Software, 23, 550

de Wit J., Seager S., 2013, Science, 342, 1473