




**Development and Application of a Small-Aperture Observatory for
Time-Series Photometry and Performance Simulations for the
PLATO Mission**

by

Morgan Adam Mitchell

A thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy in Physics

University of Warwick, Department of Physics

June 2026

**UNIVERSITY
OF WARWICK**

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	viii
Declarations	ix
Abstract	x
Abbreviations and quantities	xi
Chapter 1 Introduction	1
1.1 Astronomical Image Detectors	2
1.1.1 The Human Eye	2
1.1.2 Photographic Film and Plates	3
1.1.3 Charge-Coupled Devices	3
1.1.4 Complementary Metal-Oxide Semiconductor Active Pixel Sensors	6
1.2 Star Clusters and Variable Stars	8
1.2.1 Open Clusters	8
1.2.2 Globular Clusters	10
1.2.3 Variability in Globular Clusters	13
1.3 Exoplanet Evolution and Discovery	21
1.3.1 Definition of a Planet	21
1.3.2 Stellar and Planetary Formation	21
1.3.3 Planet Migration	22
1.3.4 Late-Stage Evolution of Planetary Systems	23
1.3.5 Detection of Exoplanets and Exo-Minor Bodies	23
1.3.6 Exoplanet Discovery Missions	28
1.4 Thesis Outline	32
Chapter 2 Methodologies	33
2.1 Astronomical Image Reduction	33
2.1.1 Bias Correction	33
2.1.2 Dark Current Correction	34

2.1.3	Flat-Field Correction	35
2.1.4	Solving Astrometry	36
2.1.5	Background Subtraction and Photometry Extraction	37
Chapter 3 The TWenty Inch Survey Telescope (TWIST)		41
3.1	TWIST	41
3.1.1	Telescope	41
3.1.2	Camera	43
3.1.3	Software	44
3.1.4	Visibility	44
3.1.5	Dark Current	44
3.1.6	Read noise	46
3.1.7	Photon Transfer Curve and Linearity	46
3.1.8	Spectral Response	49
3.1.9	Photometric precision	50
3.2	Transit of TOI-1899 b	54
3.2.1	Observation Planning	55
3.2.2	Observations	56
3.2.3	Image Reduction	56
3.2.4	Selecting Reference Stars	57
3.2.5	Photometry Extraction	57
3.2.6	Check for systematics	57
3.2.7	Light curve normalization	62
3.2.8	Time conversion and fitting	62
3.3	Pulsations of the ZZ Ceti G29-38	63
3.4	Planetary debris transiting the white dwarf ZTF J0328–1219	65
Chapter 4 Photometric Study of Messier 3		67
4.1	Introduction	67
4.2	Observations and Data Reduction	69
4.2.1	Observations	69
4.2.2	Data Reduction	72
4.3	Periodic Variable Analysis	76
4.3.1	Periodicity Search	76
4.3.2	SX Phoenicis Variables	79
4.3.3	RR Lyrae Variables	87
4.3.4	Unclassified Variables	91
4.4	Non-Periodic Variability	95
4.4.1	Microlensing	95
4.5	Conclusion	99

Chapter 5	<i>PLATO</i> Timing Predictions for the <i>Kepler</i> Field	100
5.1	Introduction	100
5.2	Simulations	102
5.2.1	Target selection	102
5.2.2	Stellar variability	103
5.2.3	Transit injection	104
5.2.4	Payload simulation	105
5.2.5	Photometry extraction and simulation runs	105
5.2.6	Detrending	107
5.3	Photometric performance	107
5.3.1	Comparison of PSF and aperture photometry	107
5.3.2	Comparison with modelled and <i>Kepler</i> noise	109
5.3.3	Correlated noise	112
5.4	Transit Fitting	112
5.4.1	Light curve preparation	113
5.4.2	Transit model and likelihood	113
5.4.3	Parameter estimation and posterior sampling	114
5.4.4	Classification of timing constraints	114
5.5	Transit timing performance and classification across SNR	115
5.5.1	Transit SNRs of the sample planets	116
5.5.2	Timing performance across SNR	116
5.5.3	Validation of timing measurements	118
5.5.4	Performance variation with sampling cadence	118
5.5.5	Transit timing yield for known <i>Kepler</i> systems	118
5.5.6	Individual systems	119
5.6	Conclusions	120
Chapter 6	Conclusions	123
6.1	Thesis Summary	124
6.2	Future Work	125

List of Tables

3.1	Details and specifications of TWIST	43
3.2	TWIST noise model quantities	51
3.3	TOI-1899 candidate reference stars	57
4.1	Photometric and astrometric properties of new SXP stars	78
4.2	Fourier properties of newly discovered SXP stars	85
4.3	Periods and amplitudes of RRL variables	88
4.4	Blazhko fit parameters	93
4.5	Unclassified variable properties	95
5.1	SNRs for Diverging <i>Kepler</i> TTV systems with <i>PLATO</i>	121

List of Figures

1.1	Hubble image of M45	9
1.2	Centaurus constellation from the Uranometria	11
1.3	Messier 22 in the infrared	12
1.4	HR diagram showing the locations of pulsating star classes	15
1.5	Example Blazhko RR Lyrae light curves	18
1.6	RR Lyrae Period–Luminosity relations per bandpass	19
1.7	Cumulative exoplanet discoveries by method	24
1.8	Exoplanet discovery space and representative survey sensitivities	25
1.9	Example microlensing light curve	27
1.10	Geometry of the <i>PLATO</i> field	30
1.11	Photometric precision comparison of exoplanet transit missions	31
3.1	The TWenty-Inch Survey Telescope in its enclosure.	42
3.2	Observing limits of TWIST	45
3.3	TWIST dark current	46
3.4	TWIST read noise	47
3.5	TWIST photon transfer curve	48
3.6	TWIST linearity curve	48
3.7	TWIST spectral response	49
3.8	TWIST zeropoint magnitude vs colour	50
3.9	TWIST noise vs magnitude dim-star focussed	52
3.10	TWIST noise vs magnitude bright-star focussed	53
3.11	Temporal scaling of photometric noise with TWIST	54
3.12	Predicted Transit Interval of TOI-1899 b	55
3.13	Visibility of TOI-1899 from TWIST	56
3.14	Normalized flux ratios of the reference stars	58
3.15	Raw fluxes of TOI-1899 and reference stars	59
3.16	Median half-flux diameters on 24 June 2024	60
3.17	Sky background flux on 24 June 2024	60
3.18	Apertures of TOI-1899 and reference stars	61
3.19	Photometry of TOI-1899 from TWIST on 24 June 2024	62
3.20	Photometry of G29-38 from TWIST in Nov 2025	64
3.21	Photometry of ZTF J0328–1219 from TWIST in Jan–Feb 2026	66

3.22	White light <i>JWST</i> NIRSpec light curve of ZTF J0328–1219 from 2026 Feb 7–8 . . .	66
4.1	Overview of the M3 survey observations	69
4.2	The M3 survey field	71
4.3	Overview of the flat-field correcting process	73
4.4	Finding charts for the new SXP variables	78
4.5	M3 colour–magnitude diagram	81
4.6	Extracted frequencies for SXP variables V301 and V304	83
4.7	Extracted frequencies for SXP variables V302 and V303	84
4.8	Example phased RRL light curves	89
4.9	Catalogued vs recovered RRL amplitudes	89
4.10	Blazhko phase-folded light curves	92
4.11	Unclassified variables V286 and V287 light curves	94
4.12	M dwarf flare candidates from the TWIST M3 survey	98
5.1	Period-radius diagram for the <i>Kepler</i> sample planets	102
5.2	<i>Kepler</i> targets within the <i>PLATO</i> LOPN1 field	104
5.3	Noise ratio, number of contaminants, contamination ratio, noise redness, and aperture size vs host brightness	108
5.4	Noise-magnitude plot for the <i>Kepler</i> hosts	109
5.5	Noise and noise redness vs contamination ratio	110
5.6	Distribution of <i>PLATO</i> vs <i>Kepler</i> noise	111
5.7	Transit timing performance statistics	117

Acknowledgments

First of all, I would like to thank my supervisor, Don Pollacco, for putting his faith in me, and for giving me the freedom to forge my own path and direct my own journey; I wouldn't have had it any other way.

They say it takes a village to raise a [PhD student], and what a truly exemplary village the Department of Astronomy at the University of Warwick has been. I am boundlessly grateful for all the conversations and advice from just about every member of staff in the group. In particular, I must thank James McCormac, without whom much of this work would just not have been possible. I must also thank Paul Chote, for trusting me with his telescopes and teaching me the ways of the observational astronomer.

I truly doubt I could have made it this far if not for my peers in (and out of) the group, so many of whom I hope to count as lifelong friends. I regret that there is no space here for an exhaustive list, but I would like to recognize in particular Ginger, Isobel, James Blake, Anna-Maria, Alastair, and Laura. Thanks for putting up with me, and for all of the advice I wish I would have taken sooner. Finally, I would like to thank the mini-boss of NGTS, Ioannis, for being by my side through thick and thin from the beginning.

I have no doubt that I will look back upon these years as some of the best of my life, thanks to you all.

THIS THESIS IS DEDICATED TO MY BROTHERS AND SISTERS.

Declarations

I declare that this thesis is my own work except where explicitly stated otherwise. This thesis has not been submitted to this or any other academic institution for any other degree or qualification.

Parts of this work have been published or submitted for publication in MNRAS. Chapters 3 and 4 are based on parts of Mitchell et al. (2026a). The astrometric solver described in Section 2.1.4 was developed by Paul Chote, and is now described in Chote et al. (2025).

Chapter 5 is based on Mitchell et al. (2026b). Fig 5.2 was created by Yoshi Eschen.

Abstract

Time-series photometry is ubiquitous in modern astronomy, enabling the study of stellar variability, transient events, and planetary systems through precise measurements of brightness changes. Unified by the theme of time-series photometry in astronomy, this thesis presents the characterization and scientific application of a new small-aperture facility, and performance simulations for the upcoming ESA M3 *PLATO* mission.

Firstly, this work describes the commissioning and characterization of the TWenty Inch Survey Telescope (TWIST), a robotic 0.5 m observatory developed for high-cadence photometric observations with a scientific CMOS camera. Measurements of instrumental noise sources and photometric performance demonstrate that TWIST is capable of delivering precision photometry suitable for a wide range of variability studies. Observations obtained with the facility are used to validate its scientific performance across timescales – from minutes-long white-dwarf debris transits to months-long variable-star pulsation modulation.

I then report on a 67-night survey of the globular cluster Messier 3 with TWIST. Analysis of the resulting photometry reveals four previously unknown SX Phoenicis variables and provides updated characterization of the cluster’s RR Lyrae population, including more than 50 Blazhko modulators. New periods, amplitudes, and classifications are obtained for several known variables, demonstrating the continued discovery potential of long-baseline photometric monitoring, even in one of the most extensively observed globular clusters in the Galaxy.

Finally, end-to-end simulations are used to assess *PLATO*’s ability to measure transit timings for known exoplanet systems in the *Kepler* field. By combining realistic detector modelling, stellar variability, contamination effects, and transit fitting, the expected timing performance of the mission is quantified and related to transit signal-to-noise ratio. Although *PLATO* is predicted to deliver lower photometric precision than Kepler, the mission’s long temporal baseline would provide valuable constraints on the long-term evolution of selected planetary systems and their transit timing variations.

Abbreviations and quantities

Abbreviations

BJD Barycentric Julian Date

CCD Charge-Coupled Device

CMOS Complementary Metal-Oxide Semiconductor

EB Eclipsing Binary

ESA European Space Agency

ESPRESSO Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations

FAP False Alarm Probability

FWHM Full-Width at Half Maximum

HFD Half-Flux Diameter

JWST James Webb Space Telescope

NASA National Aeronautics and Space Administration

NGTS Next Generation Transit Survey

PDC(SAP) Pre-search Data Conditioning (Simple Aperture Photometry)

PLATO PLAnetary Transits and Oscillations of stars

PSF Point-Spread Function

RV Radial Velocity

RRL RR Lyrae-type variable star

SAP Simple Aperture Photometry

SXP SX Phoenicis-type variable star

TCE Threshold Crossing Event

TESS Transiting Exoplanet Survey Satellite

TWIST TWenty Inch Survey Telescope

TTV Transit Timing Variation

WASP Wide-Angle Search for Planets

Quantities

R_⊕ Earth radius: 6371 km

R_☉ Solar radius: 695,000 km

M_⊕ Earth mass: 5.972×10^{24} kg

M_☉ Solar mass: 1.988×10^{30} kg

R_J Jupiter radius: 71,492 km

au Astronomical unit: 1.496×10^8 km

M_J Jupiter mass: 1.898×10^{27} kg

pc Parsec: 3.086×10^{13} km

Chapter 1

Introduction



*“How far can you bear me?” I said to Gwaihir.
“Many leagues” said he, “but not to the ends of the earth. I was sent to
bear tidings, not burdens.”*

‘The Lord of the Rings: The Fellowship of the Ring’, JRR Tolkien, 1954

1.1 Astronomical Image Detectors

Building up to the introduction of the new observatory described in Chapter 3, this section provides a high-level historical and technical overview of optical astronomical image detectors, beginning with biological detection in the human visual system and culminating with contemporary charge-coupled device (CCD) and complementary metal-oxide-semiconductor (CMOS) sensors.

1.1.1 The Human Eye

The human eye was the first astronomical detector and remained the sole means of observation from the beginning of our history until the nineteenth century. Long before the development of optical instrumentation, the night sky was the object of human fascination. Astronomical traditions developed independently around the world, often alongside astrological interpretations. Ancient Greek and Babylonian astronomers even compiled detailed planetary ephemerides and stellar catalogues (Evans, 1998).

The peak of pre-telescopic observational astronomy was reached in the late sixteenth century by Tycho Brahe, whose naked-eye astrometry achieved typical positional uncertainties of around 1 arcmin (Maeyama, 2002). These measurements were precise enough to allow Johannes Kepler to infer his three laws of planetary motion decades before Newton’s formulation of universal gravitation.

With the invention of the telescope in the early seventeenth century, the human eye remained the recording element, now combined with optical magnification. Visual observations led William Herschel to discover Uranus in 1781. By the 1840s, further visual observations of Uranus over about three-quarters of its orbital period led Le Verrier to conclude that its motion could only be explained under Newtonian gravity by the presence of an additional planet, whose position in space he predicted in 1846. Johann Galle used this prediction later that year at the Berlin Observatory to detect Neptune with nothing but a telescope, his notes, and his eyes. Thus, the eye successfully served as the pre-eminent image sensor, even as mathematical understanding of celestial motion emerged.

Physiologically, the eye functions by focusing light in stages, starting with the cornea, which accounts for most of the optical power. Light then travels through the eye’s anterior chamber before reaching the pupil, which is adjustable in size by the surrounding iris. After the pupil is the crystalline lens, which can be adjusted in shape to achieve the final focused image at the retina. Photoreceptor cells on the retina detect light signals, which are then sent to the brain via the optic nerve (Valberg, 2005).

Unaided, the human eye is capable of a (wavelength- and pupil size-dependent) angular resolution of around 1 arcmin (Yanoff & Duker, 2009), a limiting visual magnitude of around 6–7, a spectral sensitivity from around 400–700 nm according to CIE standards 1931 and 1951. In low-light conditions with its rod cells, the eye can integrate for up to approximately 300 ms, with a quantum efficiency of around 25 per cent (Rieke & Baylor, 1998). Across all conditions, the eye is capable of a dynamic range of around 10^{10} , and an instantaneous dynamic range in a single fixation commonly quoted to be around 10–14 stops ($\sim 10^3$ – 10^4 : 1).

1.1.2 Photographic Film and Plates

Photosensitive chemical mixtures were used to capture images as far back as the early eighteenth century. From this point, two main hurdles had to be overcome in order to enable practical photography. Firstly, the problem of ‘fixing’ an image – making the exposed pattern permanent by removing further light sensitivity. Secondly, increasing sensitivity such that images could be captured in seconds or minutes rather than days. The first of these hurdles was overcome sometime in the 1820s by Nicéphore Niépce using a process that coated metal or glass with a solution of bitumen in lavender oil. Upon completion of the exposure, the unhardened bitumen was dissolved away, and chemical etching made the image permanent (Willfried Baatz, 1997). Although revolutionary, the requirement for days of exposure time rendered this approach impractical.

The second hurdle was overcome in 1839 with the invention of the daguerrotype by Louis Daguerre. The breakthrough of his approach was to use a mercury fume treatment to amplify the contrast of a much more briefly exposed image. With this approach, exposure times of less than a minute were possible in direct sunlight (Willfried Baatz, 1997). Using this method, in 1840, John Draper became the first person to successfully capture a photograph of the moon (Trombino, 1980), marking the beginning of astrophotography. At that time, the low sensitivity of the daguerrotype and the evaporation-limited exposure times of the wet-chemical processes that followed limited the sensitivity of photography to less than that of the human eye. It was not until the 1880s and the popularization of dry-plate photography that astronomical photography became a useful research tool (Sebastian, 2001).

Around this time, a paradigm shift in astronomy emerged. For the first time, an objective, permanent record of night-sky observations could be made. Integrated exposures lasting hours were now possible, facilitating the recovery of detailed nebulosity and faint stars that are otherwise impossible to resolve with the human eye. The subsequent history of film/plate astrophotography is largely the history of observational astronomy itself over the next century, and thus I make no attempt to summarize it here.

Aside from their ability to integrate long exposures and exist permanently, the properties of photographic film are, in several ways, similar to or *worse* than those of the human eye. Typical quantum efficiencies of film are of the order of one per cent, maximum dynamic ranges within an image are around 100:1, and a high degree of non-linearity is present.

1.1.3 Charge-Coupled Devices

Around a decade after the invention of the integrated circuit (IC), William Boyle and George E. Smith at Bell Labs were experimenting with metal–oxide–semiconductor (MOS) structures with an open mind as to their possible applications. They realized that many MOS capacitors could be fabricated in a single IC, with each capacitor storing charge in a manner analogous to magnetic bubble storage. By applying clocked voltages to the electrodes, charge packets could be transferred sequentially from one capacitor to the next. In their introductory paper (Boyle & Smith, 1970), they suggested that such a device could be used as a shift register, an imaging device, a display device, or even to perform logic; the application of this technology to image sensing proved particularly fruitful.

The operation of CCD image sensors relies on the photoelectric effect in semiconductor materials.

Silicon, the material used in most CCDs, has a band gap of approximately 1.1 eV at room temperature. When a photon with energy exceeding the band gap is absorbed, it can excite an electron from the valence band into the conduction band, producing an electron–hole pair. In a CCD pixel, applying voltages to the MOS gate electrodes creates a potential well in the underlying silicon, forming a depletion region in which mobile carriers are swept toward the well by the electric field. Photoelectrons generated within or near this region are therefore collected and stored in the potential well during the exposure. The accumulated charge packet is then transferred through the array during readout by sequentially modulating the gate voltages, which shifts the potential wells and moves the stored electrons toward the output amplifier (Janesick, 2001). At the output, the charge packet is converted to a voltage by dumping the electrons onto a small capacitance, producing a voltage signal proportional to the number of collected electrons. This signal is buffered and amplified by an on-chip amplifier. The final stage is analogue-to-digital conversion, in which the voltage is sampled and converted into a discrete digital value (usually expressed in analogue-to-digital units, ADU), completing the transformation from incident photons to a measurable digital signal.

Early CCD image sensors employed a full-frame architecture, in which each pixel serves simultaneously as the photodetector, the charge storage well, and a stage in the shift register used for readout (Fossum & Hondongwa, 2014). During exposure, the clocking signals are held fixed so that photoelectrons accumulate in the pixel wells; during readout the charge packets are transferred sequentially through the array to a serial register and output amplifier. Because the same structure performs both photodetection and charge transfer, full-frame CCDs generally required a mechanical shutter to prevent image smear during readout.

Subsequent developments sought to separate photodetection from charge transfer in order to reduce smear and enable faster operation. One important step was the introduction of the interline transfer (ILT) CCD, in which each column of pixels is paired with a vertical CCD shift register. Photoelectrons generated in a photodiode are rapidly transferred sideways into the adjacent light-shielded register at the end of the exposure, after which the registers shift the charge to the readout node. This architecture allows charge transfer to occur while the photodiodes begin integrating the next exposure, greatly reducing image smear and eliminating the need for a mechanical shutter in many applications.

Early ILT devices, however, suffered from incomplete transfer of charge from the photodiode to the CCD register, leading to image lag. This problem was largely solved with the invention of the pinned photodiode (PPD) by Teranishi et al. (1982). The PPD uses a buried p^+np structure that allows the photodiode to be fully depleted under the transfer gate, ensuring that essentially all stored charge is transferred into the CCD register. The buried structure also suppresses surface dark current, improving noise performance. The introduction of the pinned photodiode therefore eliminated image lag and significantly improved detector performance, and the same basic structure remains fundamental to most modern imaging sensors.

Another major advance, particularly for astronomical applications, was the development of backside-illuminated CCDs. In front-illuminated devices, incoming photons must pass through the pixel circuitry before reaching the active silicon, reducing quantum efficiency. By thinning the substrate and illuminating the device from the reverse side, photons enter the silicon directly, greatly increasing quantum efficiency, particularly at short wavelengths (Janesick, 2001). Modern backside-

illuminated CCDs achieve peak quantum efficiencies up to around 95 per cent – around two orders of magnitude greater than film.

The wavelength range over which a CCD is sensitive is set primarily by the band gap of silicon and the wavelength-dependent absorption depth of photons in the material. The band gap of silicon (~ 1.1 eV) corresponds to a long-wavelength cutoff near $1.1\ \mu\text{m}$, beyond which photons do not have sufficient energy to generate electron–hole pairs. At shorter wavelengths, photons are readily absorbed, but the depth at which absorption occurs increases strongly with wavelength: blue photons are absorbed within $\lesssim 1\ \mu\text{m}$ of the surface, while red and near-infrared photons can penetrate tens to hundreds of microns before being absorbed (see Fig 3.3(b) of Janesick 2001). In conventional CCDs with relatively shallow depletion regions, long-wavelength photons may be absorbed outside the region of efficient charge collection, reducing quantum efficiency. This limitation is mitigated in deep-depleted CCDs, which use high-resistivity silicon and higher substrate bias voltages to extend the depletion region to depths of $\sim 100\ \mu\text{m}$ or more. As a result, a much larger fraction of red and near-infrared photons are absorbed within the depletion region and efficiently collected, significantly enhancing sensitivity at long wavelengths. Overall, CCDs may be optimized for sensitivity over a wavelength range spanning around four orders of magnitude, from around $0.1\ \text{nm}$ – $1\ \mu\text{m}$ – the hard X-ray to the near infrared (Janesick, 2001).

CCDs are subject to many sources of noise (see Chapter 7 of Janesick 2001). However, assuming an appropriate design and integration, the noise is dominated by three components: dark current noise, read noise, and fixed pattern noise.

Dark current arises from the thermal excitation of electrons from the valence band to the conduction band, which are then swept into the potential well. This dark current limits the maximum possible exposure time of an image. Additionally, the number of dark current electrons generated in a pixel during a given exposure follows a Poisson distribution, with the image-to-image variance equal to the mean integrated dark current; a larger dark current leads to greater dark current noise. In astronomical applications, this process is mitigated by cooling, often using Peltier-effect devices or cryogenic liquids, which can reduce dark current to a few electrons per pixel per hour. Thus, dark current noise primarily impacts in very low background environments where long exposures are used, for example, in spectroscopy, narrowband photometry, or space-based imaging.

Read noise in CCDs arises during the conversion of the stored charge packet into a measurable voltage and its subsequent amplification. Although the amplifier noise sets a fundamental limit, read noise refers to all noise sources that are independent of the signal level. Since CCDs typically have one amplifier (or at most a few) to service the whole chip, the read noise is strongly dependent on the readout speed – faster readout requires higher bandwidth in the output electronics, which increases the integrated noise power, leading to higher read noise. Conversely, slower readout allows for narrower bandwidth and longer integration times in the measurement, reducing the effective noise. As a result, CCD systems often trade readout speed against noise performance, with low-noise applications such as astronomical imaging typically operating at relatively slow pixel readout rates. For applications requiring the lowest read noise, readout times of minutes per image are commonplace.

This slow readout has negative consequences for the performance of CCDs in high-cadence applications. Increasing the cadence naturally increases the number of readouts but also typically demands the use of faster readouts, which introduce even more noise. A lower duty cycle (the fraction

of total time spent actively exposing) is another natural consequence of high cadence observing, which reduces the accumulated signal and thus the signal-to-noise ratio. This impact can be partly mitigated by the use of frame-transfer CCDs. In these devices, only half of the sensor is exposed to light, while the other half serves as a storage area. Upon completion of the exposure, the charge is rapidly transferred from the light-sensitive side to the opaque side and is read out while the light-sensitive side starts the next exposure. If the readout time is less than the exposure time, these devices can achieve a duty cycle of almost 100 per cent while allowing noise-efficient readout. Frame transfer CCDs are used by the *TESS* (Ricker et al., 2015) and upcoming *PLATO* (Rauer et al., 2025) space telescopes. The dynamic range of a sensor is given by the ratio of the full-well capacity (the maximum charge capacity of one pixel) to the read noise, and in CCDs, it can reach hundreds of thousands in high-capacity modes.

Unlike read noise and dark current noise, fixed pattern noise is spatial rather than temporal and consists of two components: photo response non-uniformity (PRNU) and dark signal non-uniformity (DSNU). Both components arise from imperfect manufacturing. Under identical illumination, two pixels will not necessarily produce the same average response (PRNU). Under dark, zero-illumination conditions, two pixels will not generate the same average dark signal either (DSNU). Since neither the PRNU nor the DSNU evolves rapidly with time, they are largely correctable with calibration frames, discussed in Section 2.1.

CCDs have been the premier astronomical detectors across the near-ultraviolet to near-infrared wavelengths since the 1980s (Lesser, 2015). Aside from their use in current observatories, they continue to be selected for upcoming ground-based and space-based observatories, including the *PLATO* space telescope and the HARMONI instrument on the Extremely Large Telescope (ELT; Ramsay et al. 2021). However, their inherently slow serial readout, together with their relatively high manufacturing complexity and cost, has led to CCDs becoming increasingly specialized devices, reserved for a limited range of scientific applications.

Since the mid-2000s, CCDs have been largely supplanted in consumer imaging by active pixel sensors (APSs), in which charge transfer and amplification are performed locally within each pixel rather than through serial readout to a small number of output amplifiers. Modern APS devices are predominantly implemented using CMOS technology – the standard semiconductor fabrication platform underpinning most modern integrated circuits, including CPUs and memory – and therefore benefit from decades of industrial development. Furthermore, the consumer imaging market is orders of magnitude larger than the scientific market, and this scale has driven rapid advances in APS performance over the past two decades, to the point that they now rival, and in some cases exceed, the performance of CCDs in many applications (e.g. Apergis et al. 2025, 2026).

1.1.4 Complementary Metal-Oxide Semiconductor Active Pixel Sensors

As a matter of terminology, Complementary Metal-Oxide-Semiconductor (CMOS) refers to an integrated circuit fabrication technology in which complementary n-type and p-type MOS transistors (nMOS and pMOS) are employed to create low-power, integrated electronic circuits. In contrast to CCDs, which fundamentally are specialized charge-transfer devices, CMOS image sensors are constructed as transistor-based integrated circuits in which charge is converted to a voltage and processed locally within each pixel. The avoidance of serial charge transfer is the central advantage of active

pixel sensors, enabling highly parallel readout architectures that substantially reduce readout time and power consumption relative to CCDs. Aside from the active pixel electronics, CMOS image sensors rely on the same underlying physical principles as CCD sensors. In both cases, photons are detected via the photoelectric effect in silicon, generating electron-hole pairs within a depleted region of the semiconductor. The resulting charge is collected and stored in potential wells defined by applied electric fields, and subsequently measured through electronic readout circuitry. Both technologies are fabricated using closely related silicon processing techniques and rely on similar material properties.

Given the same MOS foundation, early CMOS APS devices were explored soon after the invention of the CCD. However, these devices suffered a host of disadvantages compared to CCDs, arising both from the state of fabrication technology at the time and from their underlying architecture. Firstly, each pixel needed three transistors which, along with the circuitry connecting them, were located on the front side of the sensor. As a result, the optical fill factor – the fraction of the pixel area used for detection – was low, typically around 20–30 per cent (Fossum, 1997). Limitations in the contemporary minimum lithographic feature size constrained how small these electronics could be made, which in turn additionally limited the minimum achievable pixel size.

More significantly, these devices suffered from severe fixed pattern noise. In addition to the sources of FPN shared with CCDs, such as non-uniform quantum efficiency and pixel geometry, the lack of a shared output chain introduced much greater pixel-to-pixel variation. In CMOS APS devices, charge-to-voltage conversion occurs within each pixel, so variations in the local sense node capacitance lead directly to differences in gain – two pixels that have collected identical charge can produce different output voltages. In CCDs, by contrast, charge is transferred to a common output node for conversion, largely removing sensitivity to such pixel-level variations. Furthermore, each CMOS pixel contains its own amplifier, with pixel-to-pixel differences in gain and offset that can vary with temperature, further increasing the level of fixed pattern noise.

As well as fixed-pattern noise, early CMOS APS devices also exhibited significantly higher read noise than CCDs. This arose primarily from the fact that charge-to-voltage conversion and amplification were performed within each pixel, rather than at a small number of highly optimized output amplifiers. The in-pixel amplifiers were necessarily small and subject to significant pixel-to-pixel variation. Furthermore, each pixel was independently reset prior to readout, introducing reset (kTC) noise that was not effectively suppressed in early implementations. These effects resulted in substantially higher read noise than in CCDs – around 50 electrons for the CMOS compared to 3–5 electrons for the CCD (Cunningham et al., 2003).

A major advance in CMOS APS development was the realization of effective correlated double sampling (CDS) within the pixel architecture. In CDS, the pixel signal is measured relative to a reference level obtained immediately after resetting the sense node, such that subtracting the two measurements removes reset noise (kTC noise) and pixel-level offsets, significantly reducing both read noise and fixed pattern noise. Although the effectiveness of CDS had long been recognized, early CMOS pixel designs did not permit its efficient implementation, as charge collection and readout occurred at the same node and the two required measurements could not be obtained under sufficiently similar conditions for accurate subtraction.

This limitation was overcome through the introduction of the pinned photodiode (PPD) together with a transfer gate (the addition of a fourth transistor to each pixel). In this architecture, photo-

generated charge is stored within the photodiode and then transferred in a single, controlled step to a dedicated sense node, allowing both the reset and signal levels to be measured at the same node and enabling true correlated double sampling. The PPD, originally developed for CCDs, also ensures complete charge transfer, eliminating image lag, while its buried structure suppresses surface dark current, improving both noise performance and uniformity.

Like CCDs, early CMOS APSs used a front-side-illuminated architecture before the eventual development of backside-illuminated designs. The first mass-produced BSI CMOS APSs appeared in 2009 from Sony, developed to counter the wiring-induced low fill factors encountered with shrinking pixel sizes. Around this time, the prominent CCD camera manufacturer Andor Technology introduced its Neo CMOS camera, and with it the term ‘scientific CMOS’ (sCMOS). As a result of the consumer imaging-driven development of these sensors, CMOS APS technology was finally able to compete directly with CCDs in astronomy.

In the present day, scientific CMOS (sCMOS) cameras achieve read-noise performance comparable to, and in some regimes surpassing, that of CCDs. Recent ‘quantitative CMOS’ (qCMOS) architectures have demonstrated sub-electron read noise, enabling single-photon sensitivity at frame rates exceeding 100 Hz (e.g. Diouf et al. 2022). Additionally, back-illuminated CMOS sensors reach peak quantum efficiencies of $\gtrsim 95$ per cent in optimized designs (e.g. Apergis et al. 2025). In addition, sCMOS systems are typically available at lower cost than comparable CCDs, which have become increasingly specialized instruments.

However, CMOS detectors retain several limitations arising from their pixel-level readout architecture. In particular, pixel-to-pixel variations in gain and read noise introduce residual fixed-pattern noise even after calibration, while stochastic effects such as random telegraph signal (RTS) noise can produce non-Gaussian noise tails that are difficult to model and mitigate (Alarcon et al., 2023). Furthermore, the use of per-pixel amplifiers inevitably leads to greater spatial and temporal non-uniformity compared to CCDs. As a result, CCDs may still offer superior performance in applications requiring extreme photometric precision, long-exposure stability, or highly uniform noise properties, such as high-resolution spectroscopy or space-based high-precision photometry.

1.2 Star Clusters and Variable Stars

Switching from instrumentation to astronomy, the following section introduces the two broad classes of visible star clusters: open clusters and globular clusters. As a prelude to the globular cluster variability survey of Chapter 4, I introduce open clusters before moving on to an introduction and discussion of globular clusters and the variable stars found within them.

1.2.1 Open Clusters

Open clusters are young (generally $\lesssim 300$ Myr) loosely gravitationally bound stellar systems found typically in the Galactic disc (Zwart et al., 2010). An example cluster – the bright, nearby Pleiades – is shown in Fig. 1.1.



Figure 1.1: Colour-composite image of the Pleiades open cluster (M45) from the *Hubble Space Telescope*. Image credit: ESA/Hubble.

They form in giant molecular clouds, where supersonic turbulence produces a fractal-like hierarchical density field (McKee & Ostriker, 2007). Filament networks arise within these structures, and the densest filaments fragment gravitationally into many prestellar cores, which leads to clustered star formation (André et al., 2010).

However, most of these clustered star-forming structures do not go on to form open clusters; early cluster survival is largely dictated by how quickly stellar feedback, in the form of stellar winds, removes the remaining gas. The birth rate of embedded clusters – those contained within and optically obscured by their progenitor gas cloud – exceeds that of visible open clusters by an order of magnitude, and only $\sim 4\text{--}7$ per cent survive to ages comparable to the Pleiades (Lada & Lada, 2003). Survival is linked to the local gas density. In the densest regions of the interstellar medium, the gas free-fall time is short, so a larger fraction of the gas is converted into stars before feedback clears it out. This leads to higher local star-formation efficiencies, making the resulting systems more likely to remain gravitationally bound. In contrast, clusters forming in lower-density environments convert less gas into stars and are therefore more easily disrupted once the gas is expelled, leading to high so-called cluster infant mortality (Kruijssen, 2012). The consequence of this is that, although almost all stars are born in cluster environments, less than 10 per cent of the Galactic stellar disk population was ever a member of an open cluster (Piskunov et al., 2006).

Typical properties of open clusters are total masses $\lesssim 10^3 M_\odot$, virial radii¹ ~ 1 pc, and half-mass relaxation times $\lesssim 100$ Myr (Zwart et al., 2010). A census by Piskunov et al. (2006) complete up to ~ 850 pc indicates a formation rate of $0.23 \pm 0.03 \text{ kpc}^{-2} \text{ Myr}^{-1}$ (area projected into the Galactic plane), a characteristic lifetime of 322 ± 31 Myr, and a total Galactic population of around 10^5 such clusters. Dissolution in the disc is driven by Galactic tidal forces and encounters with giant molecular clouds, which cause tidal shocks and inject energy into the clusters (Gieles et al., 2006).

Compared to the general population of known variable stars, a larger proportion of the variable

¹Defined as $r_{\text{vir}} = \frac{GM^2}{2|U|}$ for cluster mass M and potential energy U .

stars in open clusters are of eruptive and rotating types, and a smaller proportion are pulsating types (Zejda et al., 2012).

1.2.2 Globular Clusters

At the opposite extreme of the star cluster landscape from open clusters are globular clusters (GCs): long-lived, tightly bound, metal-poor stellar systems whose present-day populations are typically on the order of ~ 10 Gyr old with characteristic masses $M \sim 10^5 M_{\odot}$ (Kruijssen, 2015). Typical virial radii are around 10 pc, and half-mass relaxation times around 1 Gyr – both around an order of magnitude greater than those of open clusters (Zwart et al., 2010). Unlike open clusters, globular clusters are distributed approximately spherically symmetrically throughout the Galactic halo, although some anisotropies are observed (e.g. Arakelyan et al. 2018). Under the assumption of isotropy with respect to the Galactic centre, Shapley (1918a) used the apparent anisotropy of the distribution (as seen from Earth) to reason that the Sun was not close to the Galactic centre, and to estimate the centre’s location, arriving within a few degrees of the modern value². Globular clusters are found in almost all galaxies across a wide range of masses and morphologies, and their demographics have long been used to constrain galaxy formation histories (e.g. Brodie & Strader 2006).

The widely used Milky Way compilation by Harris (2010) lists 157 objects classified as GCs and provides homogenized parameters, including positions, distances, velocities, metallicities, and photometric and structural properties. Dynamical modelling that fits kinematic and surface-density profiles with large suites of N -body models yields present-day masses and structural parameters for large GC samples (e.g. 112 Milky Way GCs; Baumgardt & Hilker 2018). In the *Gaia* era, the accessible dynamical constraints have expanded sharply, with mean parallaxes and proper motions measured for 170 Galactic clusters in a homogeneous *Gaia* EDR3 analysis (Vasiliev & Baumgardt, 2021).

A large-sample age determination by VandenBerg et al. (2013) indicates that many Milky Way GCs formed early, with absolute ages for 55 clusters spanning ~ 11 – 12.5 Gyr, showing increasing age with decreasing metallicity. Homogeneous *HST*-based relative ages by Marín-Franch et al. (2009) for 64 Galactic GCs reveal two distinct populations. The first, an old population with a ~ 400 Myr age dispersion and no age-metallicity correlation, is consistent with a rapid assembly of the Galactic halo. They identify a second group that follows an age-metallicity relation, which may tentatively indicate an origin arising from the accretion of satellite galaxies. Indeed, the largest globular clusters of both the Milky Way and Andromeda galaxies, ω Cen and Mayall II are both thought to be formed of the cores of disrupted dwarf satellite galaxies (Meylan et al., 2001).

Discovery

In principle, globular clusters are visible to the naked eye, reaching apparent magnitudes as bright as $V = 3.9$ mag in the case of ω Cen. To unaided observers, however, they appear only as faint, slightly diffuse patches of light, lacking any resolvable structure. As a result, these objects remained effectively hidden in plain sight for much of human history. Indeed, ω Cen was already recorded by Claudius Ptolemy in his 2nd-century *Almagest*, where it was catalogued simply as a star on the

²His distance estimate, around 20 kpc, which was highly uncertain per his own account, proved to be less accurate.

centaur's back. It was later assigned the Bayer designation ω within the constellation Centaurus (shown in Fig. 1.2), as though it were an individual star, by Johann Bayer in his 1603 *Uranometria*, whence its modern name is derived.



Figure 1.2: The constellation Centaurus, reproduced from the 1661 edition of Bayer's *Uranometria*. ω Cen is found at the arch of the Centaur's back, denoted by a star symbol.

Although somewhat dimmer at $V = 5.1$ mag, the globular cluster M22 (shown in Fig. 1.3) was likely the first to be identified as non-stellar in nature. Located just 44 arcmin from the ecliptic plane, Messier (1781) notes that M22 was discovered by German amateur astronomer Abraham Ihle while observing Saturn in 1665. Though distinguishable from a point source, the telescopes of the day lacked the power to resolve the individual stars therein. Messier himself referred to M22 as a nebula in his entry to the catalogue from which it derives its modern name. Messier was, however, able to resolve individual stars in the nearby GC M4 in 1764, becoming the first to do so in any globular cluster. Starting in 1782, William Herschel systematically observed Messier's catalogued objects with his large telescopes, resolving the clusters into individual stars. Herschel also discovered many new clusters while carrying out his large surveys and, in 1789, coined the term *globular cluster* (Herschel, 1789).



Figure 1.3: The Galactic globular cluster Messier 22. Stacked false-colour infrared image from the VISTA telescope at the European Southern Observatory. Image credit: ESO/VVVX survey.

Formation and dynamical evolution

In contrast to the well-understood formation of open clusters, the origin of globular clusters (GCs) is significantly more complex. Despite numerous proposed enrichment and formation scenarios, no single model currently reproduces the full set of observed constraints, particularly those related to multiple stellar populations. Empirically, characteristic chemical abundance patterns are observed in nearly all ancient GCs, yet are absent in massive clusters younger than ~ 2 Gyr, suggesting that their present-day characteristics are linked to the conditions at the time of cluster formation, which no longer exist, or subsequent long-term evolution (Bastian & Lardo, 2018).

The present-day globular cluster mass function is peaked, with a near-universal characteristic mass scale of $M_{\text{peak}} \sim 2 \times 10^5 M_{\odot}$ (e.g. Jordán et al. 2007), with only a weak dependence on host galaxy mass. This contrasts with the approximately power-law initial cluster mass functions observed in young cluster populations, and is therefore widely interpreted as the result of long-term dynamical evolution and preferential destruction of low-mass clusters (Forbes et al., 2018). However, evaporation alone cannot account for the near-universality of the peak, as disruption rates depend on both the host galaxy environment and cluster orbits. A formation framework proposed by Kruijssen (2015) addresses this by invoking a two-phase evolutionary history: clusters initially undergo rapid disruption in dense, gas-rich high-redshift discs, before migrating into galactic haloes where subsequent evolution is dominated by slower tidal evaporation. Consistent with this picture, GCs are widely thought to have formed predominantly at high redshift, in environments with higher gas pressures and surface densities than those typical of the local Universe. In the E-MOSAICS model (Pfeffer et al., 2018), the present-day ($z = 0$) GC population arises naturally as the subset of massive clusters formed at early epochs ($z \gtrsim 1-2$) that survive early disruption and subsequently evolve under the combined effects of stellar evolution, two-body relaxation, and tidal shocks.

Dynamically, GCs evolve through two-body relaxation, in which repeated weak gravitational encounters between stars redistribute energy and angular momentum over time (Heggie & Hut, 2003).

This drives key secular processes including mass segregation, core contraction (and in some cases core collapse), and the gradual escape of stars across the tidal boundary. External perturbations, such as disc and bulge shocks, can inject additional energy into clusters, accelerating mass loss and modifying their internal structure. As a result, present-day GC properties reflect both their initial conditions and their subsequent dynamical histories. A quantitative model incorporating relaxation and tidal shocking by Gnedin & Ostriker (1997) predicts substantial long-term attrition, with more than half of the present Galactic GC population expected to be destroyed over the next Hubble time, implying that the surviving clusters represent only a small fraction of the original population.

Within the Milky Way, the GC age–metallicity relation exhibits a bifurcation into two distinct branches: an old population with nearly constant age, and a younger branch associated primarily with accreted systems (Forbes & Bridges, 2010). Analyses combining *Gaia*-based kinematics with cluster ages and orbital properties suggest that ~ 40 per cent of Galactic GCs formed in situ, while ~ 35 per cent are associated with specific accretion events, including Gaia–Enceladus, Sagittarius, the Helmi stream progenitor, and Sequoia (Massari et al., 2019). The oldest clusters, with ages $\gtrsim 12.5$ Gyr, formed near the epoch of reionization and thus serve as tracers of the earliest phases of galaxy assembly (Forbes et al., 2018).

Multiple populations

GCs were long treated as simple stellar populations, motivating their widespread use as benchmarks for stellar evolution of metal-poor stars. However, it is now well established that most Galactic GCs host multiple stellar populations that are distinguishable in both photometry and spectroscopy (Gratton et al., 2012; Massari et al., 2016; Bastian & Lardo, 2018; Lee & Snedden, 2021).

A defining observational signature separating globular clusters from open clusters is star-to-star variation in light elements (e.g. He, C, N, O, Na, Al) exhibiting the chemical imprint of high-temperature hydrogen burning in extinct, earlier generations of stars (Bastian & Lardo, 2018). The anticorrelation between [Na/Fe] and [O/Fe] abundances shown in red giants (e.g. by Carretta et al. 2009), which appears ubiquitous among the GCs studied at high spectral resolution, provides a particularly clear view of this phenomenon, separating halo-like compositions from second-generation populations.

Unlike the light elements, significant star-to-star variations in [Fe/H] abundances are not common within globular clusters. This feature distinguishes typical multiple population-hosting GCs from dwarf galaxies, where broad variation is commonplace. ω Cen is an outlier, showing a large [Fe/H] range of around 1.5 dex, which supports its proposed dwarf galaxy core origin (Bastian & Lardo, 2018).

1.2.3 Variability in Globular Clusters

Globular clusters are uniquely powerful laboratories for stellar variability because their stars share approximately common distances and ages, while spanning an array of different evolutionary phases (horizontal branch, giant branches, blue stragglers) where many types of variability emerge. Their common distances are particularly useful for constraining distance scales as calibrators of period–luminosity relations.

Earliest discovery and early history

The study of variable stars in globular clusters began in the photographic plate era, where repeated exposures of clusters enabled the recognition of brightness changes by direct plate comparisons. The first periodic variable star to be discovered in a globular cluster, a Type II Cepheid (W Virginis) variable, was identified by Pickering (1890) with this method. Through this approach, it was quickly discovered that some clusters are rich in short-period variables (Pickering & Bailey, 1895).

The use of cluster variable pulsation properties to measure distances was motivated by the discovery of the classical Cepheid period–luminosity relation by Leavitt & Pickering (1912), based on the approximately equally distant members of the Magellanic clouds. This method provided a prototype for using pulsating variables as standard candles, which was used by (Shapley, 1918b) to estimate distances to several Galactic globular clusters. These calculations formed the basis of Shapley’s estimate of the location of the Galactic centre, discussed above, although they relied on an uncertain (and substantially underestimated) distance estimate to the Magellanic clouds and were affected by misidentifications of (dimmer) Type II Cepheids as classical Cepheids.

The field progressed from variable star discovery toward systematic classification when RR Lyrae populations in five well-studied Milky Way GCs (M3, M5, M15, M53, and ω Cen) were found to separate into two groups in mean period and the relative proportion of RR ab and RR c types (Oosterhoff, 1939). This discovery introduced the now-called Oosterhoff dichotomy (OoI versus OoII; discussed below) as an empirical reality requiring a physical explanation, although it is not yet fully understood today. Follow-up work in the mid-20th century refined the periods of variables in specific clusters (e.g., in M92; Oosterhoff 1944), allowing classification within this new dichotomy.

By the 1970s, catalogue compilation became a central enabling technology, with large literature compilation (notably from the David Dunlap Observatory programme) making cross-cluster comparisons and statistical analysis routine and establishing stable identification for thousands of cluster variables (Sawyer Hogg, 1973). A widely used modern collection summarizing the literature to 2001 reported on the order of ~ 3000 known variables in Galactic GCs and provided summary statistics (including mean periods by RR Lyrae subtype; Clement et al. 2001). This catalogue is still regularly updated as of 2026³.

Overview of variable-star types in globular clusters

The diversity of variable stars in GCs arises from the intersection of old, low-mass, metal-poor stellar evolution and high stellar encounter rates in dense stellar environments, which enhance interacting binaries and compact objects and drive blue straggler formation (Catelan, 2009). Most photometric variables observed in GCs are radial or non-radial pulsators located within the instability strip on the HR diagram. Pulsations of these stars are driven by the κ (opacity) mechanism, where increasing temperature causes increased opacity within the outer layers of the stars due to the partial second ionization of helium. This process drives a cycle of expansion and cooling, followed by contraction and heating, resulting in periodic photometric variability (Lamers & Levesque, 2017). This mechanism drives the variability of three of the most commonly identified variable star types in globular clusters: Cepheids, SX Phoenicis, and RR Lyrae variables. The *Gaia* DR2 colour–absolute

³The latest updates to the catalogue include results stemming from the present work.

magnitude diagram positions of stars assigned to these variable classes and the other pulsating variable star classes discussed in this section by the classifier of Rimoldini et al. (2019) are shown in Fig 1.4.




Figure 1.4: *Gaia* DR2 colour–absolute magnitude diagram (as absolute median *G*-band magnitude vs. median G_{BP} – median G_{RP}) of a selection of classification candidates in the *Gaia* DR2 variable star classifier pipeline: fundamental mode (RRab, blue dots), first-overtone (RRc, orange ‘x’ marks), and double-mode (RRd, cyan ‘+’ marks) RR Lyrae types, classical (CEP, green rhombi) and type-II (T2CEP, red squares) Cepheids, δ Scuti/SX Phoenicis types (DSCT_SXPHE, black triangles), and Mira/semiregular types (MIRA_SR, magenta circles). Faint outliers typically denote contaminating objects, which in the case of long-period variables are represented mostly by young stellar objects. All of these sources satisfy the conditions listed in Appendix B of Eyer et al. (2019), among which a relative parallax precision better than 20 per cent, with no correction for interstellar or circumstellar extinction or reddening. The background points in grey act as a reference of objects within a radius of 1 kpc from the Sun. Figure reproduced from Rimoldini et al. (2019).

Cepheids are low-surface-gravity, radially pulsating stars in the upper instability strip, more luminous than horizontal-branch variables and spanning multiple post-main-sequence evolutionary states. Classical (δ) Cepheids are high-mass ($\sim 3\text{--}13 M_{\odot}$; Rathour et al. 2025), shell hydrogen burning or core helium burning Population I stars, which are essentially never found in globular clusters. Where they are observed coincident with globular clusters, they may be assumed to belong to foreground or background populations (e.g. Kuehn et al. 2012).

Conversely, type II Cepheids are mainly old variables and are routinely found in GCs. They are divided into three subclasses based on their pulsation periods: BL Herculis (shorter periods) through to W Virginis and RV Tauri (longer periods), which roughly correspond to evolutionary states from the double shell burning post-early asymptotic giant branch (AGB) to the post-AGB and thermally pulsing AGB states with increasing luminosity and decreasing temperature (Bono et al., 2020). They are brighter than RR Lyrae and therefore valuable for more distant distance measurements, albeit at much lower number counts per cluster (Wallerstein, 2002). Typical masses are in the interval $\sim 0.52\text{--}0.59 M_{\odot}$ (Bono et al., 1997).

Anomalous Cepheids are a metal-poor (Ripepi et al., 2024) subtype with masses $\sim 1\text{--}2 M_{\odot}$ and

periods $\sim 0.5\text{--}2.5$ d (Ngeow et al., 2022) which have been discovered at lower rates than their type II counterparts in globular clusters (around 1:20; Wallerstein 2002). They are often anomalously massive given the ages of the clusters they inhabit. Their light curves are visually very similar to those of RR Lyrae where the periods are less than one day, and to those of classical Cepheids otherwise. Compared to the expectation for a given period, they are more luminous than RR Lyrae and less luminous than classical Cepheids (Soszyński et al., 2015).

SX Phoenicis stars are short-period pulsators associated with the blue-straggler star (BSS) population in clusters and are the Population II analogues of the δ Scuti class (Balona & Nemec, 2012). BSSs were originally discovered in the globular cluster M3 by Sandage (1953). They lie on or slightly above the zero-age main sequence, but brighter and bluer than the main sequence turnoff for a given cluster, implying ages younger than the clusters they inhabit. This inconsistency is generally understood to be caused by the merging of, or otherwise by mass transfer between, old, lower-mass stars. A fraction of these BSSs fall within the low-luminosity tail of the instability strip and pulsate – often in multiple radial and non-radial modes – due to the κ mechanism; these stars are SX Phoenicis variables. Typical masses are $\sim 1\text{--}1.2 M_{\odot}$ (Nemec & Mateo, 1990). Catalogue-level work by Rodríguez & López-González (2000) emphasizes that their observed cluster incidence is limited by both intrinsic rarity and the survey cadence and precision needed to identify low-amplitude, high-frequency, multi-periodic behaviour.

Another class of pulsating stars commonly found in GCs is the cool, luminous semi-regular variables (SRVs), a subtype of long-period variables (LPVs). LPVs include Mira variables, semi-regular variables, slow irregular variables, and occasionally the RV Tauri type mentioned above, among others. SRVs may themselves be further divided into four or five additional subtypes (Kudashkina, 2019). Virtually all stars towards the tips of the giant and asymptotic branches show semi-regular/quasiperiodic photometric variability, and like the variable types mentioned above, LPVs (including SRVs) demonstrate their own period–luminosity relations (Trabucchi & Mowlavi, 2022). Within globular clusters, SRVs are low-mass bright giant or supergiant stars with low-density extended envelopes and internally driven pulsation (Feast, 1996; Trabucchi et al., 2021).

Eclipsing binaries in GCs are rare, *extrinsically* variable and astrophysically valuable systems because their visible properties encode substantial useful information, not only about themselves but also about their host clusters. Combining spectroscopic and photometric data can yield precise masses, radii, and luminosities of the individual components. Knowledge of the apparent brightnesses and intrinsic luminosities allows direct, independent cluster distance measurements. Additionally, mass-radius and mass-luminosity diagrams of the component stars can be used to estimate the ages of the clusters in which they reside (Thompson et al., 2020; Rożyczka et al., 2022). Based on the statistics of Clement et al. (2001) at their time of writing, there were around two known GC eclipsing binary systems for every three globular clusters.

RR Lyrae stars, the Blazhko Effect, and Period–Luminosity Relations

Unlike Population I stars, which evolve off the giant branch into a red clump on the colour–magnitude diagram, Population II stars form a horizontal branch, with approximately constant visual luminosities but spanning a wide range of temperatures; the existence of a horizontal branch is a defining feature of globular clusters. Where the instability strip crosses the horizontal branch exists the so-called RR

Lyrae gap. Historically, colour–magnitude diagrams omitted variable stars with poorly determined mean colours and brightnesses, leading to an obvious gap in the horizontal branch. The word ‘gap’ is perhaps a misnomer, however, because stars *are* there, and these are the RR Lyrae (RRL) variables. Like the other GC horizontal branch stars, RRLs are low-mass, core-helium-burning giants that have ascended the RGB but have not yet entered the AGB (Catelan, 2009). RRLs in particular have masses $\sim 0.7 M_{\odot}$, temperatures 6100–7400 K, and $(B - V)_0$ colours 0.18–0.4 mag (Smith, 2004).

Like Cepheids, RRLs by definition pulsate radially in one or more modes. Fundamental mode pulsators are known as RR0⁴ or RR*ab* – the combination of Solon Bailey’s early, and now typically merged, *a* and *b* types. These stars generally fall towards the cooler end of the instability strip and have periods of ~ 0.4 – 0.8 d and peak-to-peak amplitudes of ~ 0.5 – 1.3 mag. They are distinguished by the characteristic sawtooth shape of their light curves, showing a rapid brightening, followed by a slow decay, and then a small but rapid dimming before brightening again. First-overtone pulsators are known as RR1 or RR*c* types, which are generally hotter, and have shorter periods (~ 0.2 – 0.5 d) and smaller amplitudes (~ 0.3 – 0.6 mag) than their fundamental mode counterparts, as well as sinusoidal light curves. RRL stars pulsating in both the fundamental and first overtone modes are designated RR01 or RR*d*. RR*ab* stars are the most common among GCs, followed by RR*c* and RR*d* types.

Discovered at the beginning of the last century (Blažko, 1907), the Blazhko effect (illustrated in Fig. 1.5) is the quasi-periodic amplitude and/or period modulation seen in RRLs and occasionally Cepheid variables (Szabó, 2013). This modulation is present in around half of RR*ab* stars (Jurcsik et al., 2009) and around 6 per cent of RR*c* variables (Netzel et al., 2018) and was identified in around 9 per cent of the Galactic bulge RR*d* stars identified by OGLE (Soszynski et al., 2014; Smolec et al., 2015). These modulations can be singly or multiply periodic, with periods ranging from a couple of days to several thousand days, and amplitude modulations typically vary from a few per cent to over 100 per cent between minimum and maximum. The effect appears to be related to period doubling, detected by *Kepler* (Szabó et al., 2010), which refers to the alternation between smaller and larger amplitudes between successive pulsations. Szabó et al. (2010) and Buchler & Kolláth (2011) argue that a 9:2 resonance between the fundamental and ninth overtone modes can explain both the period doubling and the quasi-periodic modulation, although it cannot account for a multiperiodic Blazhko effect (e.g. Sódor et al. 2011), and thus the full mechanistic picture has yet to be uncovered.

⁴In this system, introduced by Alcock et al. (2000), the designation RR*lmn* indicates the presence of radial pulsations in the *l*th, *m*th, and *n*th overtone modes, etc, with 0 corresponding to the fundamental mode. This system allows greater specificity with respect to the modes present (especially rare second overtone pulsations), but is less commonly used in the literature and is not used in the remainder of this thesis.




Figure 1.5: Phase-folded *CoRoT* Blazhko light curves. Left: folded by the Blazhko period, right: folded by the pulsational period. Representative averaged light curves are shown in two extreme modulation phases in each case in the right panels. Figure reproduced from Szabó et al. (2014).

As with the other kinds of pulsating variables, RRLs exhibit a period–luminosity relation. However, in contrast to the other types, the apparent *visual* luminosities of RRLs are not strongly correlated with their periods. This is because absolute visual brightnesses of these stars are not strongly correlated with their temperatures – the horizontal branch is horizontal when V is expressed on the vertical axis of the CMD. However, Longmore et al. (1986) showed that the $2.2\ \mu\text{m}$ K band, which is much more sensitive to temperature along the horizontal branch, shows a strong correlation with fundamental-mode pulsation period. Catelan et al. (2004) showed that the B brightness and all of the Johnson-Cousins-Glass bandpass brightnesses red of V show a stronger correlation with period than V , with the K band showing the strongest correlation for a given cluster (see Fig. 1.6). The longest wavelengths show the least scatter because greater brightness in redder bandpasses is associated with both higher luminosity and lower temperature along the HB, both of which are correlated with longer periods. In contrast, greater brightness in blue bandpasses is associated with both increasing temperature and increasing luminosity, which have opposite effects on the expected period. That is, a small change in K , whether due to changes in temperature or total luminosity, has

approximately the same effect on the period (Catelan et al., 2004).




Figure 1.6: RR Lyrae Period–Luminosity relations in several different passbands. Upper panels: U (left), B, V, R (right). Lower panels: I (left), J, H, K (right). The correlation coefficient is shown in all panels. All plots refer to an HB simulation with $Z = 0.001$ and an intermediate HB type. Figure reproduced from Catelan et al. (2004).

Despite the small scatter in these infrared PL relations within a given cluster, the slope and intercept of these relations vary from cluster to cluster, with both metallicity and horizontal branch morphology. Fortunately, the dependence on horizontal branch morphology, often parameterized by the ratio of the difference between the number of blue and red HB stars on either side of the RRL gap to the total number of HB stars, is weak in the K band and only the metallicity is significant. Thus, period–infrared luminosity–metallicity relations can be highly effective (Catelan et al., 2004). In addition to PL relations, metallicity–luminosity relations have often been used (e.g. $M_V = 0.20[\text{Fe}/\text{H}] + 0.81$; Sarajedini 2025). These relations depend on the ability to measure the reddening-corrected apparent V magnitudes of many RRL stars in a cluster, as well as metallicity to calculate distances and are less applicable to individual stars.

The development of high-precision space-based astrometry from *Hipparcos* and later *Gaia* introduced an important external calibration for RR Lyrae luminosities through the direct measurement of trigonometric parallaxes (Vasiliev & Baumgardt, 2021). This step reduces the number of rungs of the cosmic distance ladder needed to reach RRLs in GCs, reducing the systematic uncertainty in the overall distance estimates. Although individual stars may be too dim or too distant to allow precise individual parallaxes, the large survey scale of *Gaia* allows precise overall measurements when data from many GC stars are combined. These data enable direct tests of PL relations across multiple passbands (e.g. Clementini et al. 2017).

The Oosterhoff Dichotomy

The structure of RR Lyrae populations in Galactic globular clusters is shaped by the Oosterhoff dichotomy, mentioned above. Specifically, the dichotomy is that distributions of several different properties – the ratio of RR_{ab} to RR_c pulsators, the mean period of the RR_{ab} stars, and the cluster metallicity – are all bimodal across the Galactic GCs (Garro et al., 2024), and the modes of the different properties are correlated with each other. Shorter pulsation periods, a lower proportion of first-overtone pulsators, and higher metallicities are all associated with one another. The higher-metallicity group is the Oosterhoff I (OoI) class, and clusters with the opposite properties form the Oosterhoff II (OoII) class. The classes of clusters are also separated from each other in Bailey (period vs amplitude) diagrams. For a given period, OoII cluster RRLs have a larger amplitude on average than those of the OoI class (the period shift effect; Jurcsik et al. 2021). Additionally, the OoI clusters, with their higher metallicities, tend to have redder HB morphologies (more HB stars red of the RRL gap compared to blue of the gap), which follows the trend of Population I stars forming a single red clump at the equivalent evolutionary stage.

Several physical explanations have been proposed to account for this dichotomous behaviour. To explain the ratio of RR_{ab} to RR_c pulsators, van Albada & Baker (1973) invokes a hysteresis mechanism, whereby the pulsation mode depends on the evolutionary direction of a star traversing the instability strip in its journey from the blue to red ends of the HB and back, linking the observed Oosterhoff groups to the underlying evolutionary tracks. In OoI clusters, the RRL stars may be predominantly entering the instability strip from the red end of the HB, while in OoII clusters, the RRL stars may be entering from the blue end. Lee et al. (1990) suggests that the period-shift effect can also be explained by evolutionary status: if RR Lyrae stars in OoII clusters are systematically more evolved and therefore more luminous than those in OoI clusters, their longer pulsation periods follow naturally.

As discussed in Section 1.2.3, metallicity almost uniquely defines the slope and intercept of RRL period–luminosity relations, so it naturally follows that the differences in mean periods across the Oosterhoff classes are related to metallicity. Moreover, given the evolutionary mechanisms driving mode selection and period shifting, it naturally follows that membership of the Oosterhoff classes is related to horizontal branch morphology – where the HB stars are clustered on the HR diagram. It has long been known that HB morphology is primarily driven by metallicity, and thus metallicity is considered the ‘first parameter’ that dictates HB morphology in what is known as the second- (and third-) parameter problem, and these parameters are generally thought to be helium abundance or age. However, neither of these parameters is capable of explaining the full picture of observed morphologies (Catelan, 2009).

It is important to mention, however, that the relatively clear separation of the Galactic globular clusters into two Oosterhoff classes is not universally realized. Between the OoI and OoII classes exists an underpopulated mean RR_{ab} period interval $0.58 \leq \langle P_{ab} \rangle \leq 0.62$ d (Kuehn et al., 2012) known as the Oosterhoff gap, and clusters falling into this gap are called Oosterhoff intermediate clusters. Aside from the intermediate clusters, some clusters do not lie in between but instead show strong, otherwise usually mutually exclusive features of both classes. For example, the very metal-rich (characteristic of OoI) Galactic clusters NGC 6388 and NGC 6441 exhibit anomalously long RR Lyrae periods, even compared to OoII types. For these clusters, the metallicity luminosity relation breaks

down, which may imply that no simple and universally applicable relation exists after all (Pritzl et al., 2000).

Finally, complicating the understanding of the Oosterhoff dichotomy is that it does not appear to exist beyond the Milky Way, at least in the nearby dwarf galaxies, with most clusters showing mixed or intermediate properties (Soszyński et al., 2016). NGC 1786 is an example of a poorly defined cluster in the LMC, showing intermediate RRab periods, OoI-like RRc periods, and a ratio of fundamental to first-overtone pulsators characteristic of OoII clusters (Kuehn et al., 2012).

1.3 Exoplanet Evolution and Discovery

As an introduction to the exoplanet transit timing study in Chapter 5, the following section provides an exposition on the formation, evolution, and discovery of planets – especially those beyond the Solar System.

1.3.1 Definition of a Planet

In 2006, at the 26th General Assembly of the International Astronomical Union (IAU), a resolution (B5) was passed defining the term ‘planet’ in the following manner:

A ‘planet’ is a celestial body that (a) is in orbit around the Sun, (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, and (c) has cleared the neighbourhood around its orbit.

This definition applies to objects within our Solar System, but the IAU developed a working definition of the term exoplanet for objects outside our Solar System in 2018 (Lecavelier des Etangs & Lissauer, 2022):

Objects with true masses below the limiting mass for thermonuclear fusion of deuterium (currently calculated to be 13 Jupiter masses for objects of solar metallicity) that orbit stars, brown dwarfs or stellar remnants and that have a mass ratio with the central object below the L_4 / L_5 instability ($M/M_{\text{central}} < 2/(25 + \sqrt{621}) \approx 1/25$) are ‘planets’, no matter how they formed. The minimum mass/size required for an extrasolar object to be considered a planet should be the same as that used in our Solar System, which is a mass sufficient both for self-gravity to overcome rigid body forces and for clearing the neighborhood around the object’s orbit.

Not included in this definition are rogue or free-floating planets, which are not associated with a stellar system, having either formed in isolation or later separated from their host stars (Miret-Roig, 2023). However, the term ‘planet’ is usually used to describe these objects.

1.3.2 Stellar and Planetary Formation

Stars form in cold molecular clouds that fragment hierarchically into dense structures. *Herschel* observations show that star-forming filaments have a characteristic inner width of about 0.1 pc, and prestellar cores are preferentially embedded within the densest filaments (Arzoumanian et al., 2011). Because collapsing gas retains angular momentum, infall is not purely radial on small scales and a circumstellar accretion disc forms around the young star (Wurster & Li, 2018). It is in these protoplanetary discs that planets form.

Planets are generally thought to form by one of two processes: core accretion, which forms terrestrial and giant planets, and gravitational instability, which forms only giant planets (Helled et al., 2014; Durisen et al., 2007). In the core-accretion picture, planets are built from the bottom up – dust grains coagulate into larger solids, growth passes through runaway phases, and planetary embryos approach an isolation mass once they have consumed most of the solids in their local feeding zones (Pollack et al., 1996). Beyond the major snow lines, the extra availability of condensable solids makes it easier to build cores massive enough to trigger runaway gas accretion before the gas disc disperses; inside those snow lines, lower densities of solids favour the production of terrestrial planets instead (Alexander et al., 2013).

Core accretion naturally links giant-planet occurrence to stellar metallicity and broadly reproduces the contrast between rocky inner planets and volatile-rich outer planets seen in our Solar System (Fischer & Valenti, 2005). One of its weaknesses is the radial-drift problem: pebble-sized objects are expected to rapidly spiral into the star on timescales shorter than those needed to coalesce to sizes more resilient to inspiralling (Birnstiel et al., 2009). There is an additional difficulty of building giant-planet cores quickly enough in planetesimal-driven models before typical gas discs disappear (Haisch et al., 2001).

Disc instability is an alternative route. In sufficiently massive and rapidly cooling discs, self-gravity can amplify spiral structure and fragment the disc on dynamical timescales into bound clumps (Durisen et al., 2007). This channel is mostly used to explain very massive planets at large orbital distances, where classical core accretion is slowest. The directly imaged HR 8799 system is a standard reference case (Marois et al., 2008, 2010).

1.3.3 Planet Migration

Newly formed planets are not expected to remain exactly where they assembled. Low-mass planets embedded in gas discs are predicted to experience rapid Type I inward migration through a combination of Lindblad and corotation torques, whereas sufficiently massive planets can open gaps in the disc and transition to the slower Type II regime, whose pace is tied to the disc’s viscous diffusion timescale (Kley & Nelson, 2012).

After a few Myr, the circumstellar disc disperses (Haisch et al., 2001), and migration is driven by gravitational interactions with other planets or additional stars. Planet–planet scattering can excite eccentricities, move planets onto widely separated orbits, and eject planets altogether (Ford et al., 2003). Kozai–Lidov cycles due to a third body can pump eccentricity and inclination; if the periastron then becomes small enough, tidal dissipation can shrink and circularize the orbit into a hot-Jupiter-like configuration (Fabrycky & Tremaine, 2007).

The dynamics of multiple planets are intrinsically complex – the famous three-body problem – but Hill stability is a useful first diagnostic. For adjacent planets, Hill stability forbids pairs of planets switching places (maintains the distance order of planets from the star), but does not guarantee full stability: instability can still emerge later through escape of the outer planet (Petit et al., 2018), and thus predicting long-term stability usually requires numerical modelling rather than a single analytic criterion.

If migration is slow enough, planets can lock into mean-motion resonances. The *Kepler* multi-planet sample shows an asymmetry around first-order resonances, with an excess of pairs exhibiting

period ratios just above 2:1 and 3:2 and a deficit just below those values (Fabrycky et al., 2014). Resonances can preserve orbital order, as in the Neptune–Pluto 3:2 resonance, or help carve out unstable zones, as in the Kirkwood gaps produced by resonant and secular chaos in the Solar System asteroid belt (Lecar et al., 2001).

1.3.4 Late-Stage Evolution of Planetary Systems

The long-term evolution of planetary systems is also inextricably linked to the evolution of their host stars, and stellar lifetimes depend steeply on stellar mass. A non-rotating $60 M_{\odot}$ star spends only about 3.5 Myr on the main sequence, comparable to a typical disc lifetime, whereas the least massive, fully convective red dwarfs are predicted to survive for up to around 10 Tyr (Groh et al., 2014; Adams & Laughlin, 1997). The evolutionary baseline therefore differs enormously from system to system.

During the main sequence itself, tides can already alter the most compact systems. For close-in planets, coupled evolution of semi-major axis and eccentricity can proceed over Gyr timescales for commonly adopted tidal quality factors, so some present-day hot planets need not have reached their current separations immediately after disc migration ended (Jackson et al., 2008).

For a solar-mass star, the main-sequence lifetime is on the order of 10 Gyr, after which giant-branch evolution begins. A solar model by Schröder & Connon Smith (2008) predicts a maximum radius of about 1.2 AU at the tip of the red giant branch and predicts that the Sun will end as a white dwarf of roughly $0.54 M_{\odot}$, having lost about half of its initial mass overall. Around an initially solar-mass giant-branch star, Jupiter-like planets can be engulfed from initial orbits of up to around 3 AU, while terrestrial planets are less affected by tides but are still vulnerable if they begin too close to the maximum stellar radius (Mustill & Villaver, 2012).

A natural consequence of the post-main-sequence mass loss is that any planets surviving the giant branches will evolve onto wider or rearranged orbits around the white-dwarf remnant. Additionally, stellar mass loss can destabilize previously resonant or marginally stable configurations, and scattering by surviving planets can drive smaller bodies or planets onto white-dwarf-grazing or escaping trajectories (Veras, 2016). Post-main-sequence planetary systems, therefore, continue being dynamically active long after their host reaches its final stage of evolution.

1.3.5 Detection of Exoplanets and Exo-Minor Bodies

Once humans understood that the planets orbit the Sun, and that the Sun is just one of a vast number of stars in the universe, it became obvious that planets would exist outside our Solar System. Inspired by the work of Copernicus, the Italian philosopher Giordano Bruno may have been the first to recognize this as early as 1584 (Valentinuzzi, 2019). Yet, in the following four hundred years, precisely zero of these exoplanets were discovered. Based on the current IAU working definition of an exoplanet (Section 1.3.1), it was not until 1992 that the first exoplanets were discovered. Using the 305 m Arecibo Observatory and the method of pulsar timing, Wolszczan & Frail (1992); Wolszczan (1994) discovered an exotic system of three planets orbiting the millisecond pulsar PSR B1257+12. On the question of how common planets are around stars like our own, the first significant advancement came from Mayor & Queloz (1995), who used the radial velocity method to become the first to discover a planet around a solar-like star. Fostering a revolution in exoplanetary science, this advancement

was deemed worthy of a share of the 2019 Nobel Prize in Physics. As of June 2026, more than 6000 exoplanets have been discovered. Their cumulative discovery by method and year is shown in Fig. 1.7, while Fig. 1.8 shows their distribution in the mass–semi-major axis plane, together with the Solar-system planets and representative sensitivity limits for the principal detection techniques. The remainder of this section describes five of the most successful methods used to detect exoplanets.




Figure 1.7: Cumulative exoplanet discoveries by method for 11 different discovery methods. This plot was generated by the NASA Exoplanet Archive (Christiansen et al., 2025) on 2026 July 16.




Figure 1.8: Planet mass or $M_p \sin i$ as a function of reported semi-major axis or projected separation for confirmed planets with a reported mass estimate in the NASA Exoplanet Archive, coloured by discovery method. The upper axis gives the equivalent period for a $1 M_\odot$ host and the right-hand axis gives mass in Jupiter masses. Transit shading marks the region in which the geometric transit probability exceeds 1 per cent for a Sun-like star, with the 10 per cent locus also indicated. Radial-velocity lines are constant-semi-amplitude loci for $K = 10$ and 40 cm s^{-1} – the short-term ($< 1 \text{ h}$) and long-term ($< 3.5 \text{ yr}$) RV precision values of ESPRESSO (Figueira et al., 2025) – assuming a circular, edge-on orbit around a $1 M_\odot$ star. The *Gaia* DR4 curves are the $\text{RUWE} = 1.25$ forecasts of Wallace et al. (2025) for distances of 20 and 100 pc (see their Fig. 13). The Roman line is an illustrative microlensing sensitivity envelope based on Penny et al. (2019) (see their Fig. 9), and the SPHERE/SHINE curve is the 50 per cent mean detection-probability contour from Vigan et al. (2021) (see their Fig. 2). The overlays illustrate discovery space and selection biases rather than directly comparable survey completeness. Solar-system planet thumbnails are shown for visual reference and are not to scale. The images are public-domain NASA imagery obtained via Wikimedia Commons. Image credits: Mercury – NASA/Johns Hopkins University Applied Physics Laboratory/Arizona State University/Carnegie Institution of Washington; Venus – NASA/JPL; Earth – NASA/Apollo 17 crew; Mars – NASA/USGS; Jupiter – NASA/JPL/USGS; Saturn – NASA/JPL/Space Science Institute; Uranus – NASA/JPL-Caltech; Neptune – NASA.

Radial Velocity Method

In a two-body system, both the star and the planet orbit their common barycentre. The stellar reflex motion projected along the observer’s line of sight produces periodic Doppler shifts in its stellar spectrum, from which a Keplerian orbit may be inferred (Lovis & Fischer, 2010). Modern

radial-velocity surveys employ highly stable, high-resolution spectrographs, and the Doppler shift is commonly measured by cross-correlating each spectrum with reference template (Baranne et al., 1996; Lovis & Fischer, 2010). When sufficiently well sampled, the radial-velocity curve constrains the orbital period, eccentricity, argument of periastron, and velocity semi-amplitude. Combined with an estimate of the stellar mass, these observables yield the minimum planet mass, $M_p \sin i$, where $i = 90^\circ$ corresponds to an edge-on orbit. Radial velocity data alone do not provide the planetary radius, and because of the $\sin i$ degeneracy, they do not usually yield the true mass. Initially, and for two decades, the radial velocity method was the most successful method of exoplanet detection (Christiansen et al., 2025), before being overtaken by the transit method. As of June 2026, 1,186 planets have been discovered using the radial velocity method. As illustrated in Fig. 1.8, the radial velocity method is inherently biased in favour of massive planets on short-period orbits, and the state-of-the-art ESPRESSO (Pepe et al., 2010) spectrograph achieves a minimum RV uncertainty over long (~ 3.5 yr; Figueira et al. 2025) and even short (< 1 hr) timescales larger than the expected RV semi-amplitude of an Earth–Sun system viewed edge-on, indicating that the detection of such a system by means of the radial velocity method remains challenging.

Transits

A transit occurs when a planet passes in front of its host star as seen by the observer, producing a periodic reduction in the observed stellar flux (Winn et al., 2007). Transit light curves are commonly modelled with analytic or numerical eclipse models that include stellar limb darkening (Mandel & Agol, 2002; Winn, 2010). In the high signal-to-noise limit, transit photometry constrains the planet-to-star radius ratio, the scaled semi-major axis, the orbital period, and the orbital inclination. Because some of these parameters are partially degenerate, especially when ingress and egress are poorly resolved, precise and stable photometry is essential. Transit and radial-velocity measurements together yield both planetary mass and radius, and hence the mean bulk density. The shaded transit region in Fig. 1.8 shows the geometric bias of the method towards short-period systems viewed close to edge-on, and the dearth of transiting planets detected at separations greater than 1 au.

Ground-based wide-field surveys such as WASP demonstrated the effectiveness of transit searches for bright stars (Pollacco et al., 2006), while the *Kepler* mission discovered a field-transforming quantity of thousands of exoplanets with its continuous, space-based photometry (Borucki, 2016). To this day, the transit method is the most successful means of exoplanet discovery, with 4,653 confirmed planets as of June 2026.

Microlensing

After the transit and radial-velocity methods, microlensing is the third most prolific exoplanet discovery technique, with 278 confirmed planets in the NASA Exoplanet Archive as of June 2026 (Christiansen et al., 2025).

Microlensing occurs when a foreground massive object passes close to the line of sight to a more distant source star, bending and magnifying the source light. Because the lensed images are unresolved, the event is observed as a transient brightening, and the resulting light curve is typically a smooth peak lasting from hours to weeks, with short-lived anomalies if the lens is binary or hosts a planet (Paczynski, 1986; Gaudi, 2012). An example simulated light curve is shown in Fig. 1.9.




Figure 1.9: Illustrative microlensing light curve for an Earth-mass object in M22 passing in front of a 19.6 mag background star. The simulation assumes a relative proper motion of 1 mas yr^{-1} , a minimum separation of half the planet’s Einstein radius, and the point-source approximation.

Unlike transits, microlensing can produce large temporary magnifications, so it remains effective even for faint and distant source stars (Gaudi, 2012). The characteristic angular scale is the Einstein radius, which depends on the lens mass and the source and lens distances; the observed light curve is determined primarily by the closest approach and the Einstein timescale (Paczynski, 1994; Gaudi, 2012). In a simple event, these observables usually do not yield a unique lens mass, so masses are often inferred statistically unless extra information is available, for example from parallax, finite-source effects, lens flux, or high-resolution imaging years after the event (e.g. Batista et al. 2015).

Globular clusters are especially useful in this context because their distances and bulk proper motions can often be constrained independently. If a cluster lens magnifies a bulge star, these external constraints can tighten the inferred lens properties; the confirmed M22 event provides a clear example (Pietrukowicz et al., 2011).

Microlensing also has clear limitations. Events are rare because they require close alignment, they do not repeat, and the light curve alone provides no information about a lens’ radius or composition. On the other hand, the method is sensitive to low-luminosity hosts, planets with large orbital distances, and free-floating planets, making it complementary to the transit and radial-velocity techniques (Gaudi, 2012).

The Nancy Grace Roman Space Telescope (formerly WFIRST) is scheduled to launch on 2026 August 30 and will carry out a dedicated microlensing survey sensitive to Earth-mass planets and, in favourable cases, even the large-moon regime (Penny et al., 2019). The Roman sensitivity envelope in Fig. 1.8 illustrates the ability of microlensing surveys to probe low-mass planets at orbital separations of order a few astronomical units, a region that is less accessible to the transit and radial-velocity methods.

Direct Imaging

Direct imaging seeks to spatially separate planetary light from the much brighter light of the host star (Bowler, 2016). In practice, this requires starlight suppression through coronagraphy and adaptive optics (Guyon, 2018). Because current instruments are limited by extreme star–planet contrast ratios, present-day detections are strongly biased towards young, massive, self-luminous planets on wide orbits. This bias is apparent in Fig. 1.8, which includes the 50 per cent mean detection-probability contour from the SPHERE Infrared survey for Exoplanets (SHINE; Desidera et al. 2021). SHINE used the high-contrast SPHERE instrument on ESO’s Very Large Telescope to search nearby young stars for planetary and substellar companions, and its sensitivity lies predominantly at large separations and high companion masses. Although the resulting sample is much smaller than those provided by radial-velocity and transit surveys, direct imaging offers access to planetary atmospheres, luminosities, and wide-orbit architectures, while future high-contrast facilities are expected to extend this reach towards older and lower-mass planets (Chauvin, 2024).

Astrometry

Astrometry is the measurement of the position and motion of a star in the plane of the sky. In the presence of an orbiting companion, the star executes an orbit about the system barycentre, in addition to its parallax and proper motion (Sozzetti, 2024). Following in the footsteps of *Hipparcos*, *Gaia* is the defining astrometric mission of the modern era, providing repeated high-precision positional measurements for more than one billion sources (Prusti et al., 2016). Once parallax and proper motion have been modelled and removed, residual orbital motion can reveal unseen companions and, unlike radial-velocity data alone, can constrain the orbital inclination and therefore the true companion mass (Sozzetti, 2024). *Gaia* Data Release 4 is expected to include astrometric orbital solutions for many thousands of giant planets around nearby stars. The representative *Gaia* DR4 sensitivity curves in Fig. 1.8 show the predicted companion masses and orbital separations that could produce a detectable astrometric signal for Solar-mass hosts at distances of 20 and 100 pc. They illustrate that astrometric sensitivity is greatest for massive companions on intermediate- to long-period orbits and decreases for more distant systems. Astrometry is therefore complementary to radial-velocity and transit surveys, particularly for long-period giant planets that are difficult to detect in short-baseline photometric campaigns.

1.3.6 Exoplanet Discovery Missions

The remainder of this section offers a brief introduction to the main space-based exoplanet discovery missions: the now complete *Kepler* mission, the currently active *TESS* mission, and the upcoming *PLATO* mission, as a prelude to the transit-timing study of Chapter 5.

The *Kepler* mission

NASA’s *Kepler* mission was designed as a deep, single-field transit survey with a 0.95 m effective aperture and a science-quality field of view of approximately 105 deg^2 , optimized to detect Earth-sized planets through long, nearly continuous photometric monitoring (Borucki et al., 2010; Koch et al., 2010). In addition to transforming exoplanet demographics, *Kepler* made a major contribution

to asteroseismology by providing the long, high-precision light curves required to infer stellar radii, masses, and ages from oscillation spectra (Gilliland et al., 2010; Chaplin & Miglio, 2013). These stellar constraints are especially valuable for planetary interpretation because transit observations measure planet properties relative to those of the host star.

A key limitation of *Kepler*, however, is that most of its target stars are comparatively faint, which makes high-precision radial-velocity follow-up difficult for a large fraction of small and long-period candidates (Torres et al., 2017). As a result, most of its planet candidates were validated statistically rather than confirmed through direct mass measurements, and the reliability of the lowest signal-to-noise, longest-period candidates is more limited than that of the short-period population (Thompson et al., 2018; Burke et al., 2019). This is particularly important for temperate terrestrial candidates, where only a small number of transits are observed, and the false-alarm rate becomes increasingly important.

Many of the most interesting small habitable-zone planets discovered by *Kepler* orbit low-mass stars. These systems are compelling because their small stellar radii produce relatively deep transits, but their habitability remains suspect: M dwarf hosts can exhibit sustained magnetic activity, and their planets may be subject to high levels of high-energy irradiation, which may complicate atmospheric retention of close-in planets (Shields et al., 2016). Ultimately, though widely successful, *Kepler* did not succeed in its mission to find an Earth twin. The main legacy of *Kepler* is the development of a statistically powerful sample of exoplanets that covers a large region of period-radius phase space, and that it provided the necessary data to advance the field of asteroseismology beyond the brightest stars (e.g. Christensen-Dalsgaard 2012). Given the use of statistical confirmation, according to the NASA Exoplanet Archive (Christiansen et al., 2025), 2,784 confirmed planets were discovered by *Kepler* as of June 2026.

The *TESS* mission

The *Transiting Exoplanet Survey Satellite* (*TESS*) is another NASA mission, designed to be complementary to *Kepler*, trading collecting area for sky coverage and bright host stars. Rather than staring at a single field for several years, *TESS* surveys the sky in sectors of approximately $24^\circ \times 96^\circ$ spread across four non-overlapping cameras, observed for about 27 days each, with much longer cumulative coverage near the ecliptic poles (Ricker et al., 2015). This design makes *TESS* effective at finding short-period planets around bright nearby stars. The effective aperture size of each *TESS* camera is 10 cm, and the combined field of view is 2304 deg^2 .

Because *TESS* targets stars that are typically much brighter than those observed by *Kepler*, its planets are substantially more amenable to radial-velocity confirmation and detailed follow-up. This advantage comes with a trade-off in intrinsic transit sensitivity. Yield studies (Sullivan et al., 2015; Barclay et al., 2018) show that *TESS* is most sensitive to large, short-period planets and any small planets it finds are likely to orbit small host stars. Thus, *TESS* is not an Earth-twin survey, but a mission for building a bright-host sample of planets that are amenable to further characterization.

A further limitation of *TESS* is the large pixel angular scale of its cameras (21 arcsec vs 4 arcsec for *Kepler*), which increases the incidence of source blending and diluted eclipsing binaries within a photometric aperture (Sullivan et al., 2015); thus, it is all the more important that follow-up (e.g. radial velocity) be pursued. Owing to this increased need for confirmation on an individual basis,

only 897 *TESS* planets have been confirmed, compared to 7,931 *TESS* candidates as of June 2026 according to the NASA Exoplanet Archive (Christiansen et al., 2025).

The *PLATO* mission

Like *TESS*, ESA's *PLATO* mission is designed to detect and characterize transiting planets around bright stars. However, the *PLATO* mission focuses on the detection of terrestrial planets within the habitable zones of solar-type host stars (Rauer et al., 2025). Whereas *Kepler* used a single large aperture with a small field of view and *TESS* adopted an all-sky survey strategy with small cameras and relatively short observational sectors, *PLATO* will implement a hybrid approach. Its total field of view of 2132 deg^2 is comparable to that of *TESS*, while its array of co-pointing cameras (whose geometry is shown in Fig. 1.10) provides a combined collecting area equivalent to that of a single 59 cm aperture. Therefore, its photometric capability is expected to be closer to that of *Kepler*, although its precision will vary by around a factor of two for a given target brightness depending on the number of co-pointing cameras; see Fig. 1.11 for a comparison of *PLATO*'s photometric precision with other missions. Furthermore, *PLATO* will conduct a multi-year continuous observation of at least one field, similar to the long-baseline strategy employed by *Kepler*.




Figure 1.10: Geometry of the *PLATO* field, centred at the origin (0,0) of a generic spherical reference frame (units are degrees; projection is orthographic). The different shades of blue correspond to a different number of co-pointing cameras. From light to dark, the regions are viewed simultaneously by 6, 12, 18, or 24 cameras. Figure taken from Nascimbeni et al. (2022).




Figure 1.11: Comparison of the noise levels expected for *PLATO* with other missions. The approximate parameter space for two of *PLATO*’s target samples (green: P1, red: P5) is indicated for illustrative purposes. For *PLATO*, 24 N-CAMs and beginning-of-life performance are assumed (see Rauer et al. 2025; Section 11). The values for *Kepler* are taken from Van Cleve & Caldwell (2016), the values for *TESS* from Sullivan et al. (2015), and the values for *CHEOPS* from Fortier et al. (2024). The lines shown should be understood as reasonable approximations of the performance of each instrument and not used as reliable estimators for an individual target. Figure taken from Rauer et al. (2025).

Relative to the approximately month-long dwell time of an individual *TESS* sector, the extended *PLATO* baseline will substantially improve sensitivity to planets with intermediate orbital periods, including potentially habitable-zone planets orbiting solar-type stars. Compared to *Kepler*, *PLATO*’s much brighter sample should render the most favourable systems accessible to ultra-stable radial-velocity spectrographs, including HARPS3 (Thompson et al., 2016) and ESPRESSO (Pepe et al., 2010). The larger and brighter subset of transiting planets identified by *PLATO* should also be more amenable to atmospheric follow-up than the typical *Kepler* target. The precise photometry and long dwell time of the *PLATO* mission additionally permit the combination of precise transit photometry with asteroseismic characterization of host stars across a wide field of view. Accurate determinations of stellar radii and ages translate directly into improved constraints on planetary radii and bulk densities.

An additional potential advantage is that the long-duration *PLATO* fields may overlap with those of both *Kepler* and *TESS* (Nascimbeni et al., 2022; Eschen et al., 2024). In particular, if a future northern long-duration field were to include a substantial fraction of the original *Kepler* field, the greatly extended temporal baseline could be valuable for transit-timing-variation studies and for the re-assessment of marginal long-period systems (Jontof-Hutter et al., 2021).

1.4 Thesis Outline

Having introduced the instrumentation and scientific background to the present work, the remainder of the thesis is outlined as follows: Chapter 2 introduces the method of photometry extraction; Chapter 3 introduces and describes a characterization of the new TWIST observatory; Chapter 4 describes a photometric survey of the globular cluster M3 with TWIST; Chapter 5 describes an investigation of a science case for following up the *Kepler* field with the upcoming *PLATO* mission; Chapter 6 summarizes the findings of the thesis and outlines possible future investigations.

Chapter 2

Methodologies

2.1 Astronomical Image Reduction

This section outlines the process of astronomical image reduction used extensively throughout this work. This process is common to most ground-based observing at optical wavelengths and answers the question: how can we use digital images to measure the brightness of objects in space through a wobbling atmosphere that varies in opacity and surface brightness? The process of measuring these brightnesses is known as *photometry*.

The reduction starts with correcting for the peculiarities of digital cameras and follows an identical process regardless of whether a CCD or CMOS semiconductor detector is used. In both cases, each physical pixel collects charge accumulated from photoelectrons released when photons strike the detector at a rate of one electron per photon at optical wavelengths, modulated by the probability of capture – the quantum efficiency – which also depends on wavelength. Both sensor types exhibit photo-response non-uniformity, accumulate unwanted dark current, use electronic offsets, and show defective hot and non-responsive pixels – all of which must be accounted for. The basic suite of astronomical calibration images contains the bias, dark, and flat-field frames, which I introduce in turn below.

2.1.1 Bias Correction

To prevent negative values from being read out, every sensor uses a non-zero electronic voltage offset – a bias – which manifests itself as a signal floor in every digital image. This bias level, typically a few hundred analogue-to-digital units (ADU) in 16-bit readout modes, is largely uniform over the sensor and can be directly seen by taking an image with zero exposure time in darkness. This single image, however, contains both the constant underlying spatial noise and the frame-to-frame electronic read noise, which affects every image regardless of exposure time. To ascertain the underlying structure of the pixel-to-pixel variation in bias level, the impact of the read noise must be damped. This is done by simply taking many individual bias images and calculating the average value for each pixel. Howell (2000) recommends averaging 10 images, but the ideal number depends on the relative scales of the read and spatial noises. The resulting averaged bias image is commonly called the *master bias* frame, and is an approximation to the true underlying bias offset of the camera. Typically, this bias level is stable over time for a well-made camera, but it can vary with sensor temperature, so it is important to

take these images (and all other calibration images) at the same temperature as that of the on-target science images.

2.1.2 Dark Current Correction

Although typically at different levels for a given temperature, dark current affects both CCD and CMOS sensors. It arises naturally when electrons are excited across the silicon band gap into the conduction band – and is thus strongly dependent on temperature – and is most associated with imperfections at the pixel surface interfaces (Janesick, 2001). Unlike the bias level, there are large pixel-to-pixel variations in dark current levels, and the most impacted pixels are typically called hot pixels. The level of variation may be quantified by the dark-signal non-uniformity (DSNU) parameter.

The effect of dark current on the signal can be corrected in a manner similar to bias correction, by taking a series of dark calibration frames to map the rate of dark charge accumulation across the sensor, averaging them together, and subtracting the master bias frame. The resulting master dark frame containing the expected level of dark charge accumulation is then subtracted from each bias-subtracted on-target (science) image. The dark frames are, unsurprisingly, taken in total darkness with the camera shutter closed or a blocking filter in front of the sensor if available.

Since dark current is generally assumed to be constant throughout an exposure, any non-zero exposure time can, in principle, be used by scaling the dark frame to the exposure time used for the science images. However, dark current is often small and cannot be accurately measured from a single, science-frame exposure-length image. The number of dark electrons accumulated in an exposure is quantized and governed by a Poisson distribution, so the combined master dark frame must be formed of enough total dark exposure time to accumulate at least tens of electrons per pixel to provide any meaningfully accurate measurement of the dark current. The hottest of the hot pixels set a ceiling on the maximum exposure time, as the rate of dark signal generation cannot be determined from a saturated pixel.

In the case of the CMOS camera used in this work in Chapter 4, the expected dark charge accumulated per pixel in a single 10 s science image is only 0.023 e^- . To be able to measure the dark current with any accuracy, I mean-averaged 21 300 s dark images to create the master dark frame, although I note that once read noise is factored in, this method is only capable of measuring the dark current in each (non-hot) pixel with a signal-to-noise ratio of around 1.2. Since long exposures are susceptible to cosmic-ray impacts, which liberate many electrons, the dark frames may be most effectively combined by median or clipped-mean stacking – the latter to retain sub-ADU precision.

An alternative and equally valid approach to dark current correction is to simply forgo a separate bias correction and create a combined dark and bias frame from a series of dark images with exposure times equal to those used for the science images. The advantages of this approach are that it skips a calibration step and is not susceptible to dark charge accumulation that is not linear with time during an exposure. The disadvantages are the loss of the flexibility to use different exposure times, and the fact that the resulting calibration frame does not give any indication of the true level of dark current present in the sensor. Finally, if the dark current level is truly negligible in a given context, a dark current correction can be skipped entirely.

2.1.3 Flat-Field Correction

The final astronomical calibration image is the flat-field frame. While the bias/dark corrections remove signal-independent effects from the science images, the flat-field frame corrects for linear, signal-dependent effects arising from both the camera properties and the optical geometry of the imaging system. Thus, to take flat-field images, the camera must be connected to the telescope in the same configuration that will be used for the science images, including any photometric filter to be used. Simply stated, the purpose of the flat-field calibration frame is to determine the relative response of each pixel to a constant level of illumination in space, given the optical geometry of the system. The flat-field response is sensitive to two different components. Firstly, to the photo-response non-uniformity of the camera – two pixels collecting identical numbers of photons will not produce the same pixel values, even after correcting for bias and dark current, and even in the absence of noise. This can be because the sizes of the pixel collecting areas are not identically uniform, the pixels have different probabilities of converting an incident photon to a stored electron, or because (in CMOS cameras) the on-pixel amplifiers have slightly different gain values. Secondly, even when the optical system is illuminated by a source of constant surface brightness, different regions of the sensor will receive differing levels of light. Non-uniform illumination of the sensor is a natural consequence of any plausible telescope geometry, but can also be impacted by physical obstructions in the optical path, such as dust or debris on any of the optical surfaces.

As with the other calibration frames, the master flat is constructed by averaging many individual images, but unlike the other frames, flat-field exposures actually collect light. Typically, they are taken in the evening, just after sunset, and/or in the morning, just before sunrise. The telescope is pointed at a region of sky with relatively constant surface brightness, near zenith, and many exposures are taken. Throughout the evening (morning) twilight, the sky will steadily get darker (brighter), so, ideally, the exposure times should vary to produce an approximately constant integrated brightness in the images; 50 per cent of the saturation level allows for a high signal-to-noise ratio while staying within the linear photo-response regime of most cameras.

Dome flats, where the telescope is pointed at a screen illuminated by lamps within the telescope dome, offer an alternative to sky flats. Dome flats may be more suitable when taking flats for many filters, given the limited twilight time to take sky flats. They may be especially useful when using narrowband filters or slow optics, in which case the sky flux may be low. An advantage of this approach is that flats can be taken at any time with constant exposure times and stacked together to achieve an arbitrarily high signal-to-noise ratio. The disadvantage of dome flats is that light coming from the inside of the dome does not follow the same optical path as light from the sky, which may be treated as originating from infinitely far away. Consequently, while high-frequency spatial noise may be well-corrected, dome flats may not accurately reproduce the vignetting pattern present in the science images. An incorrect vignetting pattern in the master flat would manifest itself as two identically bright sources appearing to have different brightnesses in the calibrated science images. Whether this is problematic depends on the specific science purpose of the imaging – if accurate relative brightnesses of sources in the field are needed, sky flats may be more appropriate.

I note that flat-field corrections are not always used. If the telescope tracking is very precise compared to the angular size of one pixel, and one is only interested in the brightness evolution of targets in the field without regard to their brightnesses relative to each other – as is the case with

exoplanet transit photometry – a flat-field correction may only increase noise and offer no benefit. This is the case for photometry from the Next Generation Transit Survey (NGTS; Wheatley et al. 2018). For similar reasons, flat-field corrections may also be avoided in transit spectroscopy (e.g. Claringbold et al. 2026).

Regardless of whether sky flats or dome flats are used, the stacking procedure follows broadly the same steps. Firstly, each flat image must be corrected for bias and dark current with the master bias and (scaled) master dark frames. The resulting images must then each be normalized to their mean or median pixel values – this is especially important for sky flats with varying mean signal levels. If sky flats are used, the resulting images must be combined in such a way as to remove the effects of star or satellite trails present under twilight conditions. As with the creation of the master dark, median or clipped-mean stacking serves this purpose. The resulting stacked image is then the *master flat* frame.

The final image calibration is as follows: the master bias and master dark frames are subtracted from each science image, and the result is divided by the master flat frame. At this stage, the images are *calibrated*, but this is not the end of the reduction pipeline.

2.1.4 Solving Astrometry

Once the images are calibrated, the next stage of the reduction is to astrometrically solve them. That is, to determine a map that translates pixel coordinates (x, y) to sky coordinates (RA, Dec). A popular tool for this is `ASTROMETRY.NET` (Lang et al., 2010). The problem this software aims to solve is: given an unknown image position, rotation, scale, and parity (reflection), match the sources present in an image to those in a catalogue of sources with known positions, and fit a map to translate all pixel positions to sky positions within the image.

To enable fast image solving, most of the computational work is front-loaded. Before any image is solved, the information held within the catalogue of sources is compressed into a highly searchable, indexable form. `ASTROMETRY.NET` does this by extracting quadrilateral features of stars within the catalogue, since the angles between the sources are invariant to changes in position, rotation, and scale. These ‘quads’ are then stored in index files, each covering a particular range of scales and region of the sky. At the point of solving, sources are extracted from the image with a tool like `SExtractor` (Bertin & Arnouts, 1996), and quad features are identified in the resulting source lists. The index files are then searched for matching features. Once two matches are identified, the solution is verified against all sources in the image. Finally, the source positions in both coordinate systems are used to fit the map between the two, which is generally non-linear. The Simple Image Polynomial (SIP; Shupe et al. 2005) standard provides a framework for fitting non-linear distortions in the map.

While this astrometry method can solve images blindly, prior information is often available. The pixel angular scale of an optical imaging system is determined by its focal length and pixel size, and the parity of the images it produces depends only on its optical geometry; both are known in advance and neither changes from image to image. Additionally, observatories use pointing models to aim the telescope at a particular region of the sky, so the approximate sky position of the image field is also known. These three pieces of information drastically reduce the parameter space to be searched and, by extension, the solving time, which can be much less than one second.

2.1.5 Background Subtraction and Photometry Extraction

Following calibration and astrometric solving, the images can now be used to determine object brightnesses in the field; we are ready to extract photometry. The simplest method of photometric extraction is aperture photometry, which, at its core, involves counting pixel values in the image around the targets of interest within a photometric aperture.

Aperture Photometry

In order to determine the brightness of any targets in the image field, it is usually necessary to separate flux coming from the object of interest from background flux coming from various sources. Once calibrated, images will often display an approximately uniform background level, similar to the bias level, but less spatially uniform and time-varying. For ground-based observations, this background level is usually dominated by atmospheric effects: scattered sunlight, either directly or reflected off the moon, scattered terrestrial light, and airglow. Outside the atmosphere, background sources include the zodiacal light (e.g. May 2008), nebulae, galaxies (including our own) and unresolved stars.

To zeroth order, the effect of background light can be corrected by subtracting the median pixel brightness level. Over a small field of view with little surface brightness variation, this may be sufficient, especially if the targets of interest are bright. For dimmer stars, the background flux forms a larger proportion of the total integrated flux at the location of the target. Thus, the importance of an accurate background measurement increases as the brightness of the target of interest decreases. To exceed the performance of a simple median subtraction, there are two broad options: local and global corrections. Local corrections aim to measure the background level directly in the vicinity of the target, while global corrections aim to map the background level across the entire image.

Local corrections typically estimate the background flux by averaging pixel values within a circular annulus centred on the target (Howell, 2000), surrounding a circular aperture measuring the target flux. Since this annulus may contain additional stars, a median or clipped-mean approach must again be used. Local estimation is well-suited to bright stars in sparse fields and can tolerate surface-brightness variations on scales as small as the outer edge of the annulus (but no smaller). The corrected target flux is calculated by subtracting the average background level, multiplied by the number of pixels in the target aperture, from the total target flux.

In denser environments, a local annulus may be overwhelmed by background stars and thus no longer represent an accurate measurement of the background level. Additional stars can be avoided to an extent by changing the inner and outer radii of the annulus, but there are some fundamental limits: the inner radius must be large enough to exclude light from the target, the outer radius must be large enough to allow enough pixels to fall within the annulus to make a precise measurement, and both radii must be small enough for the background measurement to be meaningfully local. For dim targets, the uncertainty in the annulus-derived background measurement may contribute a significant proportion of the uncertainty in the corrected flux, given other constraints on the annulus size.

To reduce the ground-based observations described in this work, I used global background subtraction, primarily because the reduction pipeline I developed for the TWIST telescope (Chapter 3) was built to service an observational campaign of a very dense field of dim stars centred on the globular cluster M3 (Chapter 4). The approach adopted used the tools offered by `PHOTUTILS` (Bradley et al.,

2023), first masking the visible stars, and then fitting a slowly varying function to the rest of each image. The individual stellar fluxes were calculated by summing the pixel values within a circular aperture centred on each target star. A global background map avoids the problem of contaminated annuli and large uncertainties in the background level that can arise from using a small number of pixels for estimation. However, this approach is neither perfect nor universally applicable; a global background model may not capture all the features in an image, for example, the model may not be able to distinguish between the PSF wings of bright stars and background surface brightness, or capture complex nebulosity. Moreover, both approaches measure the target flux with a finite-sized aperture, and neither method can fully account for contaminating sources that fall within these apertures.

For data from *TESS* (Ricker et al., 2015), contamination within the pixel apertures is estimated using the *Gaia* (Gaia Collaboration et al., 2023) catalogues, using the known positions and brightnesses of nearby sources to estimate their contribution to the aperture (Castro-González et al., 2024). This method can remove the average impact of contamination, but not any time-dependent variability.

There are two main alternatives to the simple aperture photometry approaches described above; both hinge on the ability to accurately model the shape of the point-spread function (PSF) – the shape an infinite point source projects on the detector – in each image. The approaches are *difference imaging* and *PSF-fitting* photometry, which I introduce briefly below.

Difference Imaging

Difference imaging, as the name suggests, involves subtracting one image from another and performing the measurements on the residual image. The complexity of this approach arises from the fact that any two images will differ in background level, PSF size and shape, and zeropoint magnitude – the magnitude an object would need to generate detector flux at a rate of $1 \text{ e}^- \text{ s}^{-1}$; all of these effects must be accounted for. The process starts by forming a pristine reference image: a high-signal-to-noise image with the smallest attainable PSF, perhaps generated by stacking many excellent-seeing images. For each science image, the difference imaging software must then fit a convolution kernel that transforms the reference PSFs into the PSFs of the science image. Because the convolution operation cannot reduce the PSF size of the reference image, the reference image must be at least as sharp as the sharpest science images. The software must also subtract the background levels of both images and scale the convolved reference PSFs to match the flux levels in each science image. Only then can the images be subtracted. An example implementation of this method is *HOTPANTS* (Becker, 2015).

The advantages of the difference image approach are that it works in the most crowded fields, can cope with complex nebulosity and background behaviour, and that it is not necessary to know *a priori* where a variable source may fall within a science image; its location will be obvious in the subtracted image(s). Difference imaging has been implemented in the *GOTO* pipeline (Lyman et al., 2026) to detect supernovae in distant galaxies, as well as in gravitational microlensing surveys (e.g. Albrow et al. 2009). A disadvantage is that the method returns only additive differential flux rather than multiplicative relative flux, so it requires knowledge of the initial brightness to infer relative changes. More importantly, its effective implementation depends on the simple modellability of the PSFs in the images, which cannot be taken for granted. If the convolution does not match the PSFs well, the subtracted image may contain both highly positive and highly negative artefacts around each source. These artefacts make the measured differential flux highly dependent on the photometric aperture size

used on the difference image; it becomes necessary to include the full structure of the PSF residual while avoiding artefacts from neighbouring sources, negating the benefit of the approach.

PSF Photometry

If, instead of detecting new variability at unknown locations within an image, one is interested in performing photometry on many targets in a crowded field, PSF-fitting photometry is an excellent approach, popularized by the DAOPHOT package (Stetson, 1987). If the PSF shape is known in advance (feasible for space missions, e.g. *PLATO*; Rauer et al. 2025), or can be deduced from the science images (e.g. Massari et al. 2020), a model including one PSF component for every star can be fit simultaneously to every calibrated science image. Typically, this method is implemented in two steps: a background subtraction, followed by a flux estimation (Espinosa et al., 2025).

The advantages of PSF-fitting photometry are twofold. Firstly, compared with other methods, PSF fitting has the highest SNR ceiling. For aperture photometry, an aperture size must be selected to balance capturing more flux (larger aperture) with reducing read noise, dark current noise, and background noise (smaller aperture). In theory, PSF-fitting automatically captures 100 per cent of the flux, while including less background noise than even a small circular aperture. For a Gaussian PSF, the noise-equivalent number of pixels in a PSF-fitting ‘aperture’ is¹

$$n_{\text{eff}} = \frac{\pi}{2 \ln 2} \text{FWHM}_{\text{pix}}^2, \quad (2.1)$$

where FWHM_{pix} is the full-width at half maximum of the PSF in pixels. Thus, the equivalent aperture radius for the purpose of calculating the contribution of background noise sources is

$$\frac{r_{\text{eff}}}{\text{FWHM}} = \frac{1}{\sqrt{2 \ln 2}} \approx 0.849. \quad (2.2)$$

Thus, under ideal assumptions, PSF photometry collects all of the flux, but only the background noise of a $0.849 \times \text{FWHM}$ radius circular aperture, which is much smaller than the often used $3 \times \text{FWHM}$ aperture radius recommended by Howell (2000) for aperture photometry. Equivalently, an aperture with the above $\frac{1}{\sqrt{2 \ln 2}} \text{FWHM}$ radius would only capture a fraction $1 - e^{-2} \approx 0.865$ of the total flux of a source under the same Gaussian assumptions. Only in the limiting case that the flux is infinite or the background noise is zero does the noise performance of aperture photometry approach that of PSF-fitting.

The other main advantage of PSF-fitting is the natural way it handles contaminants. While aperture photometry can be corrected for the average contribution of contaminants, it cannot account for varying contaminants – PSF-fitting photometry can. In principle, multiple variable sources with overlapping PSFs can be measured simultaneously, with the resulting fluxes including zero contamination. This quality is especially useful when applied to star clusters, which often contain many variable stars, especially in their denser central regions.

The disadvantages lie in the practicalities of implementing the method. It is not always trivial to determine the PSF shape from the science images, especially if there are few bright, isolated stars in the image or the PSF shape varies across the image plane. Additionally, the computational cost can be much higher than that of aperture photometry. For bright, isolated stars, the increase in precision may

¹See LSST Technical Note LSE-40

be negligible and insufficient to justify the added complexity. For these reasons, aperture photometry is more commonly used for exoplanet photometric studies (e.g. Borucki 2016; Ricker et al. 2015), which necessarily focus on bright stars, making them less likely to be severely contaminated. In contrast, the cost is often justified for dim stars in crowded fields (e.g. Massari et al. 2020).

Chapter 3

The TWenty Inch Survey Telescope (TWIST)

This chapter introduces and presents a full characterization of TWIST, a general-purpose robotic telescope that I helped commission in March 2023, and which is used in subsequent parts of this work. Part of this characterization uses observational data from a large survey of the globular cluster M3, which is detailed in Chapter 4. In addition to the survey, I used TWIST to support three *JWST* programs (Cycle 2 GO 4227, Cycle 4 GO 9077, and Cycle 4 DD 12507) in collaboration with their respective PIs. These observations are detailed respectively in Sections 3.2, 3.3, and 3.4.

3.1 TWIST

TWIST (shown in Fig. 3.1) is a general-purpose, fully robotic imaging telescope equipped with a full-frame, backside-illuminated sCMOS detector. Operated and maintained by the University of Warwick, it has been in operation since its installation in March 2023 at the Observatorio del Roque de los Muchachos on La Palma, Canary Islands. The telescope is housed in the former enclosure of the SuperWASP-N observatory, which uses a sliding roof design. The building is outfitted with a Vaisala¹ weather station and multiple rain detectors, as well as an uninterruptible power supply (UPS) to ensure the roof can close in the case of loss of power. Additional details about the enclosure and its design can be found in Pollacco et al. (2006).

3.1.1 Telescope

TWIST comprises a commercially available PlaneWave Instruments² CDK20 optical tube assembly (OTA) mounted equatorially on a PlaneWave Instruments L500 direct drive mount, which does not require a pier flip when crossing the meridian. The OTA is of the ‘Corrected Dall-Kirkham’ design, which uses an ellipsoidal primary mirror, a spherical secondary mirror, and two corrector lenses. PlaneWave Instruments claims this produces a flat focal plane and minimal off-axis coma and astigmatism up to a 52 mm imaging circle. This OTA uses an open truss design that permits light to be scattered onto the optics through the gaps of the trusses, so a light shroud was installed in

¹<https://www.vaisala.com/en>

²<https://planewave.com/>




Figure 3.1: The TTwenty-Inch Survey Telescope in its enclosure on La Palma, Canary Islands. This image was taken from the southern entrance of the building, facing North.

Table 3.1: Details and specifications of TWIST’s optical tube assembly, mount and camera. The camera read noise, full well capacity, dynamic range, gain, and linearity error were all measured in the camera’s ‘Photographic’ mode at gain setting 26 – the setup used for the survey in the present work. The readout time is applicable to the streaming frame transfer mode.

Quantity	Value
Telescope	
Latitude	28°45′36.″9 N
Longitude	17°52′45.″3 W
Altitude	2349 m
Primary diameter	508 mm
Secondary diameter	191 mm
Focal length	3454 mm
Focal ratio	$f/6.8$
Plate scale	59.7 arcsec mm ⁻¹
Slew speed	50 deg s ⁻¹
Camera	
Total pixel area	9600 × 6422
Exposed pixel area	9576 × 6388
Pixel size	3.76 μm
Pixel angular scale	0.225 arcsec pixel ⁻¹
Camera field of view	36.1 × 24.1 arcmin ²
Read noise (RMS)	2.579 e ⁻
Full well capacity	22200 e ⁻
Dynamic range	78.7 dB
Gain	0.414 e ⁻ ADU ⁻¹
Linearity error (EMVA 1288 v4.0*)	0.248 ‰
Dark current (0°C)	8.25 e ⁻ pixel ⁻¹ hour ⁻¹
Readout time	39.027 μs

*<https://www.emva.org/standards-technology/emva-1288/>

February 2024 to prevent stray light from entering the optical path. PlaneWave Instruments describes the CDK20 as a ‘diffraction-limited telescope’. The complete hardware specifications of the system are detailed in Table 3.1.

3.1.2 Camera

TWIST is equipped with a QHYCCD³ QHY600M Pro sCMOS camera, a motorized electronic focuser, and a five-slot filter wheel housing a blocking filter for taking bias and dark images, as well as Bessel B , V , and R filters; one slot remains open for white-light photometry. Combined with the optical tube, the $36 \times 24 \text{ mm}^2$ Sony IMX455 sensor of the QHY600M Pro produces an on-sky field of view of $36.1 \times 24.1 \text{ arcmin}^2$. A detailed characterization of the camera’s performance and the total system spectral response is presented in Section 3.1.8, and the key results from these analyses are summarized in the lower half of Table 3.1. Additionally, an analysis of the photometric performance of the system is presented in Section 3.1.9. The photometric performance is in line with expectations for the typical observed point spread function half-flux diameter of around 3 arcsec.

³<https://www.qhyccd.com/>

3.1.3 Software

The mount, camera, filter wheel, focuser, and roof are all controlled by an on-site computer running custom `ROCKIT`⁴ software on Rocky Linux⁵. The `ROCKIT` system is a modular observatory control framework designed to manage all aspects of robotic telescope operations, including hardware control, environmental monitoring, and data acquisition. It operates as a collection of lightweight services (daemons) that communicate through a remote procedure call interface, ensuring scalability and modularity. These services handle tasks such as aggregating weather data, managing dome and telescope hardware, automating observation schedules, and processing image metadata. In addition, a centralized web dashboard provides real-time system monitoring and access to weather, operational status, and video feeds.

This system automates flat-field imaging (morning and evening), focusing, target field acquisition, filter selection, and target observation. These operations are defined as *actions* within `ROCKIT` and can be scheduled nightly using an observing script file. The schedule is carried out for as long as the environment is considered safe. That is, the temperature, humidity, wind speed, and sun altitude are within an acceptable range, rain is not detected, the UPS is powered, network connectivity is maintained, and disk space is available. If these conditions are violated, the roof is automatically closed, and observations are suspended. These provisions negate the need for human monitoring of the observatory during its operation.

We perform autoguiding on the science images using `DONUTS` (McCormac et al., 2013), allowing us to forgo a separate guide telescope, which would add unnecessary complexity.

3.1.4 Visibility

The sliding roof design was initially chosen to facilitate rapid movements of the SuperWASP cameras. However, this configuration limits TWIST’s visibility, especially in the North, where altitudes $< 63^\circ$ are obstructed by the building. Nevertheless, all objects with declinations $-22^\circ \leq \delta \leq 56^\circ$ retain some visibility from TWIST. Additionally, the visibility is asymmetric in the East-West direction due to the L-mount design (see Fig. 3.1). The telescope hangs lower when pointing East, and is therefore more obstructed by the walls of the enclosure than when pointing West. As a result, the positive hour angle limit is greater than the negative hour angle limit in magnitude for a given declination. The full observing constraints of TWIST are illustrated in Fig. 3.2.

3.1.5 Dark Current

Unlike typical CCD cameras, the QHY600M Pro CMOS camera used by TWIST achieves a very low dark current with only modest cooling. We measured the dark current by taking ten dark frames with exposures ranging from 50 to 500 s under dark conditions with a blocking filter in place and the camera set to 0°C . To avoid the influence of dead or hot pixels, we analysed the mean pixel values from the non-overscan region of the images, excluding the smallest and largest 0.5 per cent of values. Fitting a linear model to the mean values, as shown in Fig. 3.3, and assuming a gain of $0.414 \text{ e}^- \text{ ADU}^{-1}$ (see Table 3.1), we calculated a dark current of $(8.25 \pm 0.05) \text{ e}^- \text{ pixel}^{-1} \text{ hour}^{-1}$.

⁴<https://rockit-astro.github.io/>

⁵<https://rockylinux.org/>




Figure 3.2: The observing limits of TWIST. **Top:** The minimum observable altitude (black line) measured to the nearest degree versus the azimuthal angle, sampled every 10 degrees. The white-shaded region shows the observable altitude/azimuth combinations. The dashed lines are loci of points of a fixed declination, illustrating the sidereal motion of an object of that declination. **Bottom:** The observable hour angles (shaded white) visible for a given declination and the total unobstructed time per night (black line) at that declination. This time does not include sun- or moon-induced visibility restrictions.




Figure 3.3: **Top:** Truncated mean signal versus exposure time for ten dark exposures ranging from 50 to 500 s shown as red points. The mean signal is computed from the central 99 per cent of pixel values, excluding the lowest and highest 0.5 per cent. The black line represents the least-squares linear fit to the data. The gradient of this line is the calculated dark current, and the vertical axis intercept is the mean bias level. **Bottom:** The residual signal calculated by subtracting the linear model from the data.

This dark current represents a negligible contribution to the overall noise budget, particularly during short (10-second) exposures, for which the expected dark current noise is just 0.151 e^- .

3.1.6 Read noise

With the same setup as above, we took 21 bias images with the shortest available exposure time ($39 \mu\text{s}$) to measure the frame-to-frame variability of individual pixels in the absence of signal – the read noise. Fig. 3.4 shows the distribution of the means and standard deviations of the pixel values.

We observe a clearly non-Gaussian distribution of read noise values, which can likely be attributed to random telegraph noise (also known as salt and pepper noise), which is characteristic of CMOS sensors and was found to be present in the QHY600M Pro by Alarcon et al. (2023). To account for the fact that photometric apertures contain many pixels, we take the overall read noise value to be the root mean of the variances: 2.579 e^- . Assuming the dark signal is uniform and follows a Poisson distribution, the read noise is dominant over the dark current noise for exposures shorter than around 48 minutes.

3.1.7 Photon Transfer Curve and Linearity

The Photon Transfer Curve (PTC) was introduced in Janesick et al. (1985) and can be used to extract valuable information about the camera sensor, including the gain, full-well capacity, read noise, and fixed pattern noise. The PTC has two main implementations – one made with single exposures per signal level and one made with the difference image of two exposures per signal level. The latter




Figure 3.4: Read noise and mean pixel value distribution from 21 consecutive bias frames. **Left:** A 2D histogram showing the temporal, frame-to-frame noise and the mean signal of each pixel in these bias frames. The colour bar indicates the number of pixels in each bin. **Right:** The black histogram shows the distribution of read noise values, and the red dotted line shows the root-mean-square read noise.

approach is more tolerant of less perfectly uniform sensor illumination, although it removes – and therefore cannot be used to measure – fixed pattern noise. We adopted this approach based on the testing equipment available to us.

We illuminated the sensor with a constant light source in a dark room with the camera off the telescope. We varied the exposure time to achieve different signal levels, taking two images per exposure time. After correcting for bias and dark current, we cropped each image to a 200×200 pixel central region and measured the mean signal of each pair of images. We then subtracted one image from the other and measured the variance of the difference image. For the PTC, we divide this variance by two to account for the noise contribution from both images in the pair and infer the single-image noise. The PTC is shown in Fig. 3.5. We fit a straight line to the linear region of the curve, and its inverse gradient yields a gain of $0.414 \text{ e}^- \text{ ADU}^{-1}$ – very close to the manufacturer’s quoted $0.412 \text{ e}^- \text{ ADU}^{-1}$. The full-well capacity is given by the signal level at which maximum variance occurs, around 22,200 electrons. This is much lower than the value of 27,045 electrons given by QHYCCD. In fact, the full-well capacity given by QHYCCD is not possible using the gain they quote, as 16-bit digital saturation would occur first.

With the same set of measurements used for the PTC, but instead using only one image per exposure level, we measured the linearity error of the sensor according to version 4.0 of the EMVA 1288 standard (Jähne, 2010). Therein, the linearity error is calculated by considering a linear regression of, with respect to time, the signal within a range of 5-95 per cent of the full-well capacity. The regression minimizes the *relative* residuals – the residuals divided by their respective best-fit values. The linearity error is then defined as the mean of the absolute values of these relative errors, in this case 0.248 per cent. The linearity curve is shown in Fig. 3.6.




Figure 3.5: Photon transfer curve for the QHY600M Pro camera in its ‘Photographic’ mode at gain setting 26. The best-fit line in the top plot spans only the data used in the fit. The residuals span the entire data range but quickly fall below the axis limits.




Figure 3.6: Linearity curve for the QHY600M Pro camera in its ‘Photographic’ mode at gain setting 26. The best-fit line in the top plot, which minimizes the absolute sum of the *relative* residuals, spans only the data used in the fit. The residuals span the entire data range but quickly fall below the axis limits. δ expresses the relative residuals as a percentage.




Figure 3.7: Optical throughput of TWIST and *Gaia*. The black line represents the modelled total throughput of TWIST, including quantum efficiency and camera lens transmission taken from Gill et al. (2022), atmospheric extinction at the Roque de Los Muchachos observatory at an airmass of 1 from King (1985), and assuming a constant transmission of 93 and 99.5 per cent for the mirrors and corrector lenses respectively. The green, blue and red lines show respectively the *Gaia* (E)DR3 G , G_{BP} , and G_{RP} passbands from Riello et al. (2021).

3.1.8 Spectral Response

Like most modern sCMOS detectors, the Sony IMX455 sensor of the QHY600M Pro is most sensitive to blue light, reaching a peak quantum efficiency of over 90 per cent between 450 and 510 nm according to QHYCCD. A model of the optical throughput of the entire system, compared to the *Gaia* DR3 passbands (Gaia Collaboration et al., 2023), is shown in Fig. 3.7. PlaneWave Instruments does not publicly share the wavelength-dependent transmission of the lenses and mirrors of the OTA. Hence, we approximate the throughputs of each mirror and each lens to be constant at 93 and 99.5 per cent, their averages over the interval 400–700 nm. We note that the transmission is likely lower outside that range.

Additionally, we examined the spectral and absolute response of the system observationally using the unblended, non-variable stars in the survey (Chapter 4). Fig. 3.8 shows the measured flux-averaged G -band zeropoint magnitudes of the stars against their $G_{BP} - G_{RP}$ colours. The trend is approximately linear over the colour range of the sample and shows a clear reduction in sensitivity towards redder stars. Applying a 3σ -clipped linear regression to the individual zeropoints yields an average colour-dependent G -band zeropoint,

$$ZP_G = 23.0640 - 0.2591 \times (G_{BP} - G_{RP}) \pm 0.0413. \quad (3.1)$$

The scatter about the fit is almost entirely not due to measurement uncertainties, but rather a systematic uncertainty arising from the complexity of the real stellar spectra that cannot be captured by the $G_{BP} - G_{RP}$ colour alone. We note that this calculation uses fluxes from 5-pixel radius apertures, which capture, on average, 70.6 per cent as much flux as 15-pixel radius apertures in our testing – equivalent to a 0.388 mag lower zeropoint.




Figure 3.8: Mean G -band zeropoint magnitudes (difference between G magnitudes and flux-averaged instrumental magnitudes for a 5-pixel aperture radius) versus $G_{BP} - G_{RP}$ colours for 148 stars brighter than $G = 18.5$. The stars chosen are not marked as variables in *Gaia* DR3. They are unblended, which we define as being at least 22.5 arcsec (50 binned pixels) from a source brighter than three magnitudes dimmer than that star. The data points are coloured according to the stars’ catalogued G magnitudes. A 3σ -clipped linear least squares fit to the data is overplotted.

3.1.9 Photometric precision

The photometric performance of the system was measured from observations of the globular cluster M3 (Chapter 4) and compared to a theoretical noise model. The noise model combines errors from the target, sky background, dark current, read noise, and scintillation. The total noise for a point source, N_{total} , is modeled by equation 3.2:

$$N_{\text{total}} = \sqrt{F_{\text{target}} + n_{\text{pixel}}f_{\text{sky}} + n_{\text{pixel}}f_{\text{dark}} + n_{\text{pixel}}\sigma_{\text{read}}^2 + \sigma_Y^2 F_{\text{target}}^2}, \quad (3.2)$$

where F_{target} is the total integrated target flux in electrons, f_{sky} and f_{dark} are the integrated sky and dark flux in electrons per pixel, n_{pixel} is the number of pixels in the photometric aperture, σ_{read} is the read noise of one pixel in electrons, and σ_Y is the fractional scintillation noise approximated by the modified Young formula (Young, 1967; Osborn et al., 2015):

$$\sigma_Y^2 = 10^{-5} C_Y^2 D^{-4/3} t^{-1} \sec^3 z \exp(-2h_{\text{obs}}/H), \quad (3.3)$$

where C_Y^2 is the empirical Young’s coefficient that depends on the level of atmospheric turbulence in the vicinity of the observatory site. D is the telescope aperture size, t is the exposure integration time, z is the zenith angle ($\sec z$ is the approximate airmass), h_{obs} is the height of the observatory above sea level, and H is the atmospheric scale height. The relevant values of the quantities described above for the TWIST/La Palma are listed in Table 3.2.

We found that the photometric aperture size chosen to maximize the signal-to-noise ratio for dim stars is not ideal for brighter stars. This is likely for several reasons. Firstly, averaged across all the images used for photometry, 5-pixel radius apertures captured only 70.6 per cent as much flux as

Table 3.2: Relevant quantities influencing the noise properties of the photometric system within our model. The dark current and read noise were measured in the present work using our chosen camera setup of 0°C and gain setting 26 in Photographic mode.

Quantity	Value	Reference
H	8000 m	Osborn et al. (2015)
C_Y	$1.30 \text{ m}^{2/3} \text{ s}^{1/2}$	Osborn et al. (2015)
f_{dark}	$0.00229 \text{ e}^- \text{ pixel}^{-1} \text{ s}^{-1}$	This work
σ_{read}	2.579 e^-	This work
D	0.508 m	
t	10 s	
h_{obs}	2349 m	

15-pixel radius apertures (the largest we extracted). As a result, the ‘signal’ component of the SNR is diminished. This is especially significant for sources whose precision is limited by photon shot noise. Secondly, the use of small photometric apertures exacerbates the impact of their imperfect placement; a difference in the placement of the aperture produces a larger difference in the collected flux for a smaller aperture. Finally, and most pertinent to the brightest stars, small apertures prevent spatial averaging over distortions of the stars’ PSFs due to atmospheric scintillation. As a result, the theoretical scintillation noise floor was not reached, except in our testing, where we used much larger apertures. Hence, we alter the model, replacing the scintillation component with the ~ 12 ppt empirical noise floor we observed. The model and its constituent components are shown alongside the measured noise-to-signal values in Fig. 3.9.

Even with the very small pixel angular scale present in our setup, the per-pixel read noise is low enough that sky background dominates at larger magnitudes. Photon shot noise is only dominant at around $G = 16$. As expected, redder stars have greater noise than bluer stars of the same G magnitude, owing to TWIST’s blue-biased spectral response.

For simplicity, the preceding analysis neglects non-uniform response across the field arising from vignetting. In the current optical configuration, the flat-field response at the dimmest corner of the detector is approximately 80.6 per cent of that at the brightest point near the image centre. To assess the impact of this variation, we generated two additional noise models – representing the brightest and dimmest regions of the field – by modulating the zeropoints and sky background levels of the nominal model according to the relative flat-field response. The resulting predictions indicate that bright stars already at the noise floor experience identical photometric precision across the field. For photon-limited stars, the expected increase in noise from the centre to the edge of the detector is 11.4 per cent, while for sky-background-limited stars, the increase rises asymptotically with magnitude towards 15.1 per cent.

Based on the results of the survey in Chapter 4 (Fig. 4.1), the median PSF half-flux diameter achieved by TWIST is around 3 arcsec. The best image achieved 1.7 arcsec. This is significantly higher than the typical atmospheric seeing from La Palma and the theoretical diffraction limit for a 0.5 m class telescope. It is also worse than the 1.3 arcsec achieved by pt5m (Hardy et al., 2015), another 50 cm telescope located just 215 metres away from TWIST. The larger-than-expected PSF has persisted despite attempts to recollimate the telescope and tune the mount; we suspect dome seeing, perhaps related to the slow release of heat from the concrete floor, to be a likely culprit. The poor




Figure 3.9: Photometric precision as a function of stellar magnitude predicted by our noise model and measured by our survey observations at a 10-second cadence. The solid line represents the total modelled noise, with contributions from photon shot noise, sky background noise, and read noise, and an empirical noise floor component. Dark current noise is negligible and would not be visible on this scale. The dots represent the measured scatter in the detrended light curves of 993 stars categorized in *Gaia* DR3 as non-variable and brighter than $G = 18.5$. The noises are scaled by the inverse square root of the proportion of images in which a given star is detected to portray the precision achievable in a fixed time interval. The points are coloured by the stars’ $G_{BP} - G_{RP}$ colours.

seeing has a significant, negative impact on the signal-to-noise ratio of dim stars due to the need for larger photometric apertures, capturing more read noise and sky-background noise. We would expect the SNR to double for stars in this regime if the typical seeing were instead around 1.5 arcsec.

To confirm whether scintillation-limited performance could be achieved, we performed an identical noise analysis on 10 s cadence test data from the globular cluster M15, obtained before the survey, since we had already extracted photometry over a wider range of aperture sizes. We found that the minimum aperture size needed to achieve scintillation-limited performance was a 20 pixel radius (9 arcsec; $\sim 3 \times$ the typical HFD). This is in line with the $3 \times \text{FWHM}$ recommendation from the Handbook of CCD Astronomy by Howell (2000)⁶. We found that scintillation-limited performance was reached in both bright and dark-sky environments with 20-pixel apertures (illustrated in Fig. 3.10), and that the corresponding noise for dimmer stars is much higher than with 5-pixel apertures. We observed much less scatter about the noise model with the larger apertures, likely due to greater spatial averaging of the PSFs and the addition of white noise, which drowned out any correlated noise.

⁶For a Gaussian profile, the FWHM and HFD are identical




Figure 3.10: Photometric precision as a function of stellar magnitude for both full-moon (left) and dark sky (right) conditions predicted by our noise model and measured by our observations, using 20-pixel radius apertures. The grey line represents the total modelled noise, with contributions from photon shot noise (green), sky background noise (pink), read noise (orange), dark current noise (blue), and scintillation noise (purple). The dots represent the measured scatter in the detrended light curves of 225 unblended stars with $10 < G < 18$ in a field centred on Messier 15 under dark and light observing conditions. We assume the median scintillation conditions at La Palma (empirical Young's coefficient, $C_Y=1.30$). The model predicts scintillation noise for each night using the actual zenith angles traversed over the course of the observations. The model flux is calculated using the derived mean zeropoint for each night, and the sky background levels are taken from the mean measured background level on each night.

Finally, on the same dark- and bright-sky 20-pixel aperture data, we tested for the presence of correlated noise by measuring the scatter in the detrended light curves of the brightest ($G < 13$) stars as a function of the number of data points binned (averaged) together and comparing to the $\sim 1/\sqrt{n}$ expectation under the assumption of white noise. The results are shown in Fig. 3.11. In both sky conditions, the noise appears to decrease as expected, at least up to 420 s, after which we stopped to ensure that each standard deviation measurement takes at least 20 data points. The results are similar to those of Apergis et al. (2026), who used a CMOS APS device for aperture photometry, and demonstrate that CMOS cameras can perform not only precise, but also stable photometry.




Figure 3.11: Noise to signal ratio versus exposure time of an artificial reference star composed of the average detrended and normalized light curves of the 21 unblended stars with $10 < G < 13$ mag in the field centred on Messier 15. The data points represent the measured standard deviation of the reference star’s light curve as the 10-second exposures are mean average binned together, effectively increasing the exposure length. The dashed lines show the expected $1/\sqrt{n}$ decrease in noise from the unbinned data, where n is the bin size. The grey points represent the data from 2023 July 1 (bright sky), and the black points represent the data from 2023 September 21 (dark sky).

3.2 Transit of TOI-1899 b

This section details the observation and data reduction of a transit of the temperate Jupiter, TOI-1899 b, across its early M-type main-sequence host star (see Lin et al. (2023) for details of the planetary system). The observations were made to support the *JWST* Cycle 2 GO Program 4227 (PI: Claringbold), which were planned for 2024 September 19–20, but actually occurred on 2025 May 10–11, using the NIRSpec instrument to characterize the atmosphere of TOI-1899 b. To refine the transit ephemeris and help plan these observations, I observed a transit of TOI-1899 b with TWIST on the night of 24 June 2024.

3.2.1 Observation Planning

We used the NASA Exoplanet Archive’s Transit and Ephemeris Service⁷ to identify suitable transits observable from the Roque de Los Muchachos Observatory. We propagated the ephemeris and period from Lin et al. (2023) and selected the night of 2024 June 24 as a good candidate for observation. However, we did not expect to observe the transit ingress. Fig. 3.12 shows the expected transit visibility on the night.




Figure 3.12: Predicted transit interval of TOI-1899 b generated by the NASA Exoplanet Archive Transit and Ephemeris Service. The solid line shows the altitudes traversed by the star during the night. The pink section of this line indicates the expected transit interval.

I calculated the visibility from TWIST using the calculator described in Section 4.2.1 and found TOI-1899 to be observable from 23:14 (coincidentally the predicted mid-transit time) until 05:14 UTC with a mean lunar illumination and separation of 95 per cent and 64° respectively. These timings correspond to 2.35 hours of in-transit observations and 3.65 hours of out-of-transit observations. The observability window is illustrated in Figure 3.13. From Section 3.1.9, the expected photometric noise for a $G = 13.54$ mag star measured with a 10-second cadence during bright sky is around 8.82 ppt. Assuming a transit depth of 2.82 per cent from the ratio of the radii of the planet and host star in Lin et al. (2023), by Equation 5.4, this gives an expected signal-to-noise ratio of around 93, which implies that the transit would be highly detectable.

⁷<https://exoplanetarchive.ipac.caltech.edu/cgi-bin/TransitView/nph-visibletbls>, accessed 03 July 2024




Figure 3.13: Visibility of TOI-1899 from TWIST on 24 June 2024. The dotted yellow and white lines show the sun and moon paths, respectively. The dark and light blue lines show the path of TOI-1899 while unobscured and obscured, respectively. From brightest to darkest, the shades of blue in the background indicate times of daylight, civil twilight, nautical twilight, astronomical twilight, and astronomical night.

3.2.2 Observations

On the night of 24 June 2024, the telescope roof was opened at 20:28 UTC. Ninety-four flat-field images (flats) were taken in the evening, starting at 20:30 UTC at solar altitudes ranging from -4.1° to -6.4° . Following the taking of flats, the roof remained open to allow the telescope to reach thermal equilibrium with its environment. Five minutes before the first visibility of TOI-1899, at 23:09 UTC, the telescope underwent its autofocus routine. At 23:14 UTC, the target field was acquired using the blind offset technique. Starting at 23:17 UTC, the target field was observed using 10-second exposures and 2×2 binning. 2,142 images were taken, corresponding to 5 hours and 57 minutes of exposure time. To maximize optical throughput, no photometric filter was used. In the morning, 109 flats were taken at solar altitudes ranging from -6.2° to -4.1° . For morning and evening flats, 2×2 binning was used and the exposure times were automatically adjusted to maintain approximately 50 per cent digital saturation, ensuring a high signal-to-noise ratio and good sensor linearity.

3.2.3 Image Reduction

The flat-field images were cropped to the active image regions, discarding the optically black, over-scan, and metadata rows and columns. The resulting flats were then corrected for bias and dark current and normalized by their medians. The normalized images were subsequently median stacked. The resulting image was detrended, and the master vignetting pattern was applied (discussed in Section 4.2.2), producing the final master flat. Existing master dark and bias images, as well as the hot-pixel map, were used in this analysis. Each science image had a bias, dark current, flat-field,

and hot-pixel correction applied. The images were then corrected for sky background and solved astrometrically.

3.2.4 Selecting Reference Stars

In order to perform relative photometry on the target star, at least one reference star was needed. I queried the *Gaia* DR3 source catalogue for stars visible within the field-of-view of TWIST, centred on TOI-1899. All proper motions were propagated from the *Gaia* DR3 epoch to the night of observations. To avoid blended sources within the apertures, I performed a cut on all stars with at least one source brighter than 3 magnitudes (in G) dimmer than the target within a separation of 20 arcsec. Since no photometric filter was used, it was vital to select reference stars of a similar colour to the target to avoid airmass trends in the photometry. A further cut was made to include only stars within 0.1 magnitude in $G_{BP} - G_{RP}$ of the target and brighter in G than the target. Finally, any stars flagged as variable in *Gaia* DR3 were excluded. This left six candidate reference stars, which are listed in Table 3.3.

Table 3.3: TOI-1899 and candidate reference stars properties from *Gaia* DR3. The first entry in the table is TOI-1899 itself. The right ascensions and declinations (in ICRS) were calculated by propagating the *Gaia* DR3 coordinates by their proper motions over the interval of time between the *Gaia* DR3 epoch and the night of observations.

<i>Gaia</i> DR3 ID	RA (deg)	Dec (deg)	G (mag)	$G_{BP} - G_{RP}$ (mag)
2073530190996615424	299.426963	40.143393	13.542	1.913
2075034387304990592	299.573343	40.325198	10.600	1.981
2072776368387032960	299.589693	40.018621	12.209	1.901
2072763655276345728	299.693359	39.961632	12.243	2.000
2074277515312077184	299.718406	40.112762	12.352	1.898
2073531904665125632	299.429601	40.184176	12.784	1.880
2074282669272738176	299.569910	40.213023	13.066	1.904

3.2.5 Photometry Extraction

Once the images were calibrated, background-subtracted and astrometrically solved, photometric apertures could be placed and flux measured. Using the Python package `PHOTUTILS`, forced photometry was extracted at the catalogued locations of the target and reference stars using a range of aperture radii – from 4 to 20 binned pixels at 2-pixel intervals.

3.2.6 Check for systematics

To ensure the stability of the reference stars, for each aperture size, the ratio of each reference star’s light curve to that of every other reference star was calculated. *Gaia* DR3 2073531904665125632 was found to be unstable with respect to the other reference stars, regardless of aperture size, so it was rejected. The remaining five reference stars, whose (normalized) 20-pixel flux ratios are shown in Fig. 3.14, were found to be stable with respect to each other.




Figure 3.14: Flux ratios of the five selected reference stars from Table 3.3, normalized by the median ratio. Each star is labelled in the plot titles by the final five digits of its *Gaia* DR3 ID. The red lines show the quadratic polynomial best fit to the ratios.

I inspected the raw fluxes of the remaining stars – shown in Fig. 3.15 alongside TOI-1899 – finding smooth, consistent variation of brightness with time. The slow trend is as expected, given the evolution of the airmass over time: it decreases initially and rises again after TOI-1899 passes the meridian.




Figure 3.15: Raw flux of TOI-1899 and reference stars from TWIST on the night of 24 June 2024. Each light curve is normalized by its median value and offset by an integer multiple of 0.1 for clarity. The error bars represent the 1σ uncertainties derived from the photon and background noise (composed of sky background and read noise).

Additionally, to check for systematic effects impacting all stars, I plotted the evolution of the median half-flux diameter (Fig. 3.16) and the sky background flux (Fig. 3.17). The jump in HFD between 00:00 and 01:00 UTC is reflected as instability in the photometry with the smaller apertures, but has little visible effect on the 20-pixel apertures. The sky background appears to vary smoothly, initially rising as the moon’s elevation increases, remaining roughly constant, then rapidly rising at the onset of nautical twilight. There is no sign of abrupt changes in sky brightness due to cloud cover or the headlamps of passing cars, which could interfere with the photometry and potentially require the masking of some data.




Figure 3.16: Median point-spread function half-flux diameters during observations of TOI-1899 with TWIST on the night of 24 June 2024.




Figure 3.17: Median sky background flux during observations of TOI-1899 with TWIST on the night of 24 June 2024.

Finally, I visually inspected the images in the vicinity of the target and reference stars, overlaying all *Gaia* sources to assess the impact of contaminating stars. Cutouts from one sample image are shown in Fig. 3.18. At the largest tested size, each aperture contains at least several stars brighter than $G = 18$ mag. However, none of the stars show any contaminants within 4 mag located within the apertures or close to the aperture boundaries. This confirmation provides the final layer of confidence in the choice of reference stars.




Figure 3.18: Apertures of the target star TOI-1899 and the selected reference stars in a sample image, labelled in the plot titles by their *Gaia* DR3 IDs. The catalogued locations of the target and reference stars are shown by the red circles. The catalogued locations of contaminating stars are shown by the cyan circles. Within each image, the area of the marker for each star is directly proportional to the expected flux received from that star.

3.2.7 Light curve normalization

The five final reference stars' light curves were normalized by their medians and mean-averaged together to produce a master reference star. The relative flux from TOI-1899 was then calculated by taking the ratio of the raw target flux to the master reference flux for each aperture size. To remove any remaining airmass trend in the light curve, a second-order polynomial was fit to the obviously out-of-transit portion, and the entire light curve was divided by this trend. The resulting light curves were then binned to 10-minute time intervals as in Lin et al. (2023). The out-of-transit scatter, the standard deviation of the normalized flux, was calculated from the *binned* data, and the lowest scatter was found from the 20-pixel radius aperture data. The binned data was used to expose systematic bumps in the light curve that would otherwise be hidden by white noise. Due to its noise profile, the 20-pixel radius aperture data was selected.

3.2.8 Time conversion and fitting

The mid-exposure time stamps were converted from UTC into Barycentric Julian Date (BJD) in the Barycentric Dynamical Time (TDB) scale using `ASTROPY`, given the coordinates of the target star and the location of the observatory. The final light curve is shown in Fig. 3.19. The final reduced photometry was passed on to a collaborator for fitting. Their analysis revealed that the transit egress occurred within one minute of the time expected from Lin et al. (2023). This allowed the observation window for the transit to be measured with *JWST* to be confidently handed over to STSci. The results of the atmospheric analysis will be shown in Claringbold et al. (in prep).




Figure 3.19: Final relative photometry of TOI-1899 from TWIST on the night of 24 June 2024. The grey points represent the unbinned data, and their error bars show the 1σ uncertainties in the normalized flux. The black points represent the data binned to 10 minutes, and their error bars represent the expected uncertainty assuming only white noise in the unbinned data.

3.3 Pulsations of the ZZ Ceti G29-38

As a further demonstration of the capability of TWIST in a different context, this section *briefly* overviews the results of a 12-night observation campaign of the dust cloud-surrounded white dwarf ZZ Ceti *g*-mode pulsator G29-38. These observations were carried out in support of the *JWST* Cycle 4 GO Program 9077 (PI: Cunningham), for which G29-38 was observed using the NIRSpec instrument on 2025 Nov 13. The observations from TWIST formed part of a large ground-based campaign aimed at providing continuous photometric coverage of G29-38 over a three-week period centred on the *JWST* observations. The purpose of the ground-based campaign was to provide a wider baseline to constrain the active stellar pulsation modes around the time of the observations, since the *JWST* observing time was limited to only 11 hours on NIRSpec. Moreover, G29-38 is known to have a power spectrum that evolves over periods of weeks to years (Kleinman et al., 1994), so maximum coverage over a timescale of relative stability was essential.

Since the data from TWIST were to be combined with data from other instruments, parts of the observations and reduction were standardized across observatories. Namely, the *V* filter was utilized, and a single comparison star, TYC 584-1035-1 was used for reference. Aside from these externally imposed choices, I used a 20 s exposure time, which is double that used in the previous section, despite G29-38 having a similar brightness to TOI-1899. This exposure time was chosen to account for the flux loss induced by the placement of the *V* filter and to lower data rates, while still allowing around 10 samples per period of the highest frequency of interest (~ 5 mHz or $1/200$ s $^{-1}$).

Additionally, to further lower the data rate, I used 3×3 binning, increasing the effective pixel angular scale to 0.675 arcsec pixel $^{-1}$, which still satisfactorily samples the typical 3 arcsec HFD PSFs of TWIST observations. Naturally, this required the use of the same binning for the dark, bias, and nightly flat-field images. The remaining observing and reduction steps proceeded as in the previous section: flats were taken in the evening and morning where possible (in the *V* filter), and the target was acquired at the beginning of astronomical twilight and observed until it reached TWIST’s visibility limit at an altitude of around 28° above the horizon. The images were calibrated as before, correcting for bias, dark, and flat-field, before undergoing background subtraction and being solved astrometrically. Forced photometry was extracted using a wide range of aperture sizes at the proper motion-propagated positions of all the $10 < G < 16$ mag stars in the field, which included the target and reference star. These flux measurements were used to generate a mean *G*-band zeropoint for each image at each aperture size. Since there is no equivalent out-of-transit constant baseline region for G29-38, the optimum aperture size was selected by visually inspecting the light curves, balancing statistical scatter and systematic non-astrophysical trends, rather than by an explicit calculation of NSR. An aperture radius of 20 unbinned pixels (6.67 binned pixels) was chosen on this basis.

Unlike the case for the transit of TOI-1899 b, the weather conditions were not always amenable to observations for the campaign of G29-38. Of the possible 21 nights, I was only able to observe at all on 12 due to adverse weather conditions. On the nights of Nov 6, 8, 16, 24, and 25, the control software deemed the conditions unsafe at some point due to high humidity, but the observations could be restarted later in the night. The nights of Nov 14, 23, and 24 saw partial cloud cover, as evidenced by abrupt drops in the zeropoint magnitude, which affected the photometry at various points. Images with a mean *G*-band zeropoint magnitude less than 21.45 mag were rejected, and this threshold was chosen as the minimum required to prevent cloud-induced photometric spikes. The final photometry

was passed on to the program PI to be combined with the data from the other observatories and is shown in Fig. 3.20. The full analysis will be found in Cunningham et al. (in prep).




Figure 3.20: Final relative photometry of G29-38 from TWIST for 12 nights between the nights of 2025 Nov 05–25. The points are unbinned, and the error bars showing the 1σ uncertainties in the normalized flux are too small to see.

3.4 Planetary debris transiting the white dwarf ZTF J0328–1219

As a final demonstration of TWIST’s capabilities outside the main survey discussed in the next chapter, I present the photometry from another 12-night observing campaign, this time on a dimmer target, ZTF J0328–1219 ($G = 16.6$ mag). These observations were carried out at short notice – I started the observing run on the same day that the request to observe was received – in support of the *JWST* Cycle 4 DD Program 12507 (PI: Xu).

ZTF J0328–1219 is one of only around a dozen white dwarfs known to show deep photometric transits due to tidally disrupted planetary debris. Evidence of planetary debris transits was first detected at the Zwicky Transient Facility (ZTF), and follow-up photometry from the McDonald Observatory showed these transits to be regular (Guidry et al., 2021). ZTF J0328–1219 was later studied in depth by Vanderbosch et al. (2021) using photometry from six observatories, as well as spectroscopy from the Southern Astrophysical and Magellan Telescopes. They found a helium-dominated atmosphere and strong circumstellar calcium absorption. They isolated two periods (9.937 and 11.2 hr) from the photometry, which they interpret as the orbital periods of different debris clumps. They note further that the transit morphologies evolve on nightly, weekly, and monthly timescales, with some observations showing no transits at all.

ZTF J0328–1219 was observed as part of the XCOV31 Whole Earth Telescope (WET) campaign, which took place between 2025 Oct 13 and Nov 6, in support of the *HST* Cycle 32 GO Program 17789 (PI: Guidry), which took place on 2025 Oct 25–26. As part of the WET campaign, I also observed ZTF J0328–1219 using the Warwick 1m telescope for five nights – 2025 Oct 18, 20, 22, 24, and 26. The WET and *HST* observations demonstrated large photometric activity, with transits blocking up to 35 per cent of the white dwarf flux – the system had entered a highly active phase. The *JWST* program investigators used this heightened activity as the basis for the application for director’s discretionary time, to observe the chemical absorption and thermal emission of the newly produced dust.

The purpose of my observations was to generate a baseline over the 12 nights leading up to and including the night of the *JWST* run, to achieve full phase coverage of the debris period, and to facilitate understanding of the evolution of the individual transits. Despite the mid-winter conditions in La Palma, I was able to observe on 12 out of a possible 13 nights, including simultaneously with *JWST*, and including the day after the *JWST* run. In addition to my observations, a team at the University of Delaware took three nights of data using the SARA Observatory 0.6m telescope on Cerro Tololo, Chile. For consistency with their observations, I again used the Bessel V filter. As with the G29-38 observations, I used 3×3 camera binning and 20 s exposures.

By late January, ZTF J0328–1219 was already past its most observable season, visible for only four hours in the evening. As with the previous set of observations, the target was acquired at the onset of astronomical twilight and observed until TWIST’s visibility limit was reached – in this case, at an altitude of 25° . Evening and morning flats were taken where possible. In total, 6,920 on-target frames were collected, corresponding to 38.44 hr of exposure time. The data reduction proceeded in exactly the same manner as for the G29-38 observations, with a single reference star, *Gaia* DR3 5161804400570962048 – the same as was used for the WET campaign.

I ran a PDM periodogram on the reduced light curve, finding a period of 11.36 d, which is consistent with the behaviour observed at the time of the *JWST* proposal writing, but which represents

an evolution from the state observed by Vanderbosch et al. (2021). The observations achieved full phase coverage with sufficient precision to identify several evolving transiting clumps. The final TWIST phase-folded light curve is shown in Fig. 3.21, and the white light NIRSpec light curve is shown in Fig. 3.22. The full analysis of the combined data will be found in Rogers et al. (in prep).




Figure 3.21: Final relative photometry of ZTF J0328–1219 from TWIST for 12 nights between the nights of 2026 Jan 27 – Feb 8. The grey points represent the unbinned data, and their error bars show the 1σ uncertainties in the normalized flux. The black points represent the data binned to 200 seconds, and their error bars represent the expected uncertainty assuming only white noise in the unbinned data. The *JWST* observing window is shown in red.




Figure 3.22: White light NIRSpec light curve of ZTF J0328–1219 from 2026 Feb 7–8. The error bars show the 1σ uncertainties in the normalized flux. This figure was provided by Siyi Xu.

Chapter 4

Photometric Study of Messier 3

We present the first scientific observations from the newly installed TWIST observatory – a 50 cm telescope equipped with an sCMOS camera providing a 36.1×24.1 arcmin² field of view – housed in the former SuperWASP-North enclosure. We conducted a 67-night, 199-day baseline white-light monitoring campaign centred on the globular cluster M3 aimed at characterizing stellar variability within the cluster while also assessing the photometric performance of the newly commissioned system. We report the discovery of four new SX Phoenicis variables (V301–304), confirm their cluster membership, and identify fundamental-mode pulsation in one, allowing an independent period–luminosity-based distance estimate to M3. We revisited 231 previously known RR Lyrae stars, providing updated period measurements for 203 and white-light amplitudes for 198. We detected Blazhko-like modulation in 53 stars and characterized the modulation parameters for 28. Notably, we measure periods and amplitudes for the unclassified variables V286 and V287 for the first time. We also identify three foreground flaring M dwarfs, and assess the feasibility of detecting microlensing events in M3, concluding that expected rates are negligible. Alongside the scientific results, we introduce a new correction technique for flat-field images affected by scattered light and present a full characterization of the observatory’s photometric capabilities. These results demonstrate the scientific utility of TWIST for high-cadence time-domain surveys using modest-aperture instrumentation.

4.1 Introduction

The globular cluster M3 (NGC 5272) is among the most extensively studied in the Milky Way. Located at a distance of 10.2 kpc (Harris, 2010) at high Galactic latitude ($b = 78.7^\circ$), M3 contains one of the richest populations of variable stars in any Galactic cluster, second only to ω Cen. Since the discovery of the first periodic variable in a globular cluster by Pickering (1890), M3 has served as a benchmark for studies of RR Lyrae variables and horizontal branch evolution (e.g. Catelan, 2004). The most recent update to the catalogue by Clement et al. (2001) in March 2019 lists 273 confirmed variables, including 241 RR Lyrae (RRL) and 9 SX Phoenicis (SXP) stars.

M3 is the prototypical Oosterhoff type I cluster (Oosterhoff, 1939), exhibiting relatively short-period RR Lyrae variables and a greater predominance of fundamental-mode (RRab) pulsators. While once considered a classic example of a simple stellar population, M3 is now known to host multiple populations (e.g. Massari et al., 2016; Lee & Sneden, 2021), possibly the result of a merger. The cluster also hosts the first discovered blue straggler stars (Sandage, 1953), of which its SXP variables

are examples. The radial distribution of M3’s blue stragglers shows a bimodal structure (e.g. Ferraro et al., 1993, 1997), and while early surveys found relatively few SXPs (Kaluzny et al., 1998; Hartman et al., 2005), this may reflect selection biases in observing cadence and sensitivity.

M3 has remained the focus of ground-based variability studies for the last several decades across a wavelength range spanning the near-ultraviolet to the near-infrared. Early modern near-infrared observations by Longmore et al. (1990) explored the K -band period–luminosity relation of RR Lyrae variables in M3 and other clusters, demonstrating a correlation between $\log$ (period) and K magnitude almost free of systematic scatter. Subsequent CCD variability surveys were undertaken by Kaluzny et al. (1998) in the V -band, Corwin & Carney (2001) in the BV -bands and Hartman et al. (2005); Benkő et al. (2006) in the BVI -bands. They obtained high-quality light curves for hundreds of variables, including newly identified RRL and SXP stars. These works presented 9, 14, 19 and 9 nights of data, respectively, and had minimum observing cadences ranging from around 5 to 20 minutes.

A series of follow-up investigations by Jurcsik et al. (2012, 2015, 2017); Jurcsik (2019) further examined the long-term stability, mode switching, Blazhko modulation, and physical properties of M3’s RR Lyrae population through extended photometric and spectroscopic monitoring. In the near-infrared, Bhardwaj et al. (2020) obtained time-series JHK_s photometry for 233 RR Lyrae stars, establishing precise mean magnitudes and updated period–luminosity relations. Most recently, Kumar et al. (2024) presented a multiwavelength ($UBVR IJHK_s$) analysis based on an extensive 35-year dataset, deriving improved periods, Fourier-based light-curve parameters, and inferring corresponding physical properties using an artificial neural network for a large sample of RR Lyrae stars.

The aforementioned variability studies have all used CCD imaging with cadences of minutes, which may limit sensitivity to short-period pulsators and brief transients. Recent advances in scientific complementary metal-oxide-semiconductor (sCMOS) technology now offer an alternative to CCDs for time-series photometry. With improved quantum efficiency, lower read noise, and orders of magnitude shorter readout times, sCMOS cameras are increasingly adopted in both ground-based projects (e.g., Licandro et al., 2022; Law et al., 2022; Airey et al., 2025) and proposed space missions (e.g., Malbet et al., 2022; Ge et al., 2022). Large-format sCMOS sensors are now readily available at a lower cost than many traditional CCDs. These capabilities make sCMOS cameras particularly well-suited for high-cadence, wide-field photometry – ideal for monitoring dense stellar fields such as globular clusters. However, their performance in long-term, time-domain studies remains relatively underexplored.

In this paper, we present the first results from TWIST (the TWenty Inch Survey Telescope), a new 50 cm telescope equipped with an sCMOS camera, housed in the former enclosure of the SuperWASP-North facility. We used TWIST to carry out a 10-second cadence, 67-night, 199-day baseline white-light photometric campaign of M3, aimed at validating the system’s photometric capabilities and evaluating its scientific return.

We describe the observations and data reduction, including a new method for correcting scattered-light contamination in flat-field images in Section 4.2. We report new discoveries and updated properties for known variable stars in M3 in Section 4.3, and evaluate the probability of observing microlensing events originating from the cluster in Section 4.4.




Figure 4.1: Overview of the M3 survey observations. The red bars represent the total number of 10-second exposure, on-target images captured each night, and the blue points indicate the median half-flux diameters of all images taken on that night. Nights with a lower-than-expected number of images typically reflect either delayed starts to observations or early terminations due to environmental triggers. The two large gaps in the data are the result of deliberate non-observing periods, not adverse weather or other external factors. The month labels reflect the beginning of the month in the year 2024.

4.2 Observations and Data Reduction

Once TWIST became operational, we undertook a dedicated survey of the large, variable-rich galactic globular cluster M3 as its first science target.

4.2.1 Observations

We conducted initial test observations of M3 on 2023 December 13–14, when nightly visibility was approximately one hour. Regular monitoring commenced on 2024 January 22, with observations carried out on every clear night. We produced custom code to calculate the observable window each night, folding in the complex observing constraints of TWIST. This code also automatically generated the nightly observing script. When environmental constraints allowed, we observed each night for the entire visible window and took morning and evening flats. Operations were paused between 2024 March 18 and 2024 April 24 for system maintenance, and the survey culminated on 2024 June 29 when nightly visibility fell below three hours. The survey observations are overviewed in Fig. 4.1. In total, 130,074 science images were captured, corresponding to over 15 days of exposure time.

M3 has a tidal radius of 28.7 arcmin (Harris, 2010), comparable to the 36.1×24.1 arcmin² field of view of TWIST. Thus, the overwhelming majority of the cluster’s stars are encompassed within a single pointing. The field was centred on the cluster’s core to maximize the number of member stars included. The full field is shown in Fig. 4.2.

To avoid excessive data rates, we binned the pixels to 2×2 bins, creating an effective angular pixel scale of 0.45 arcsec pixel⁻¹, which still oversamples the typical seeing. To ensure the efficacy of the autoguiding algorithm, we limited exposure times to 10 s in this work. We opted to forgo a photometric filter in order to maximize the amount of light collected – given TWIST’s relatively small

aperture primary mirror – in order to observe as deep as possible. We ran the camera in its streaming frame transfer mode, where the readout time was only $39.027\ \mu\text{s}$, corresponding to a fraction of time spent actively exposing (duty cycle) of over 99.9996 per cent. Finally, we operated the camera in its ‘Photographic’ mode with gain setting 26 – a configuration that moderately favours lower read noise at the expense of reduced full-well capacity. This compromise was chosen to balance noise performance and dynamic range, given the faintness of our targets – the apparent *Gaia* DR3 (Gaia Collaboration et al., 2023) *G*-band magnitude of the horizontal branch is around 15.5 mag, while the main-sequence turn-off occurs at around 19 mag.




Figure 4.2: The 36.1×24.1 arcmin² survey field centred on M3. The locations of known variables from Clement et al. (2001) that have a corresponding *Gaia* source are shown. RRL, SXP, semi-regular, eclipsing binary, W Virginis and uncategorized stars are shown with open circles, squares, diamonds, triangles, stars, and pentagons, respectively. The newly discovered SXP-type variables V301–304 and the unclassified variables V286–287 are labelled. North is up and East to the left. The image was created by aligning with *SPALIPY* (Lyman, 2021) and stacking 360 calibrated frames (totalling one hour of exposure) with measured half-flux diameters less than 2 arcsec taken during dark sky environments. The 5σ limiting magnitude of the stacked image for a 5-pixel aperture radius is around $G = 22.1$.

When conditions allowed, around 80 flat-field images were taken under the automated control of the ROCKIT system per morning/evening twilight. Images were taken at 75° altitude and anti-solar azimuth at sun altitudes less than -4° . The exposure times were automatically adjusted to maintain ~ 50 per cent digital saturation up to a maximum exposure time of 10 s. At the end of each night, the camera was automatically shut down, and the telescope was parked facing West at an altitude of 20° to both prevent debris from falling on the telescope optics and to ensure the telescope cannot point at the sun in case the roof fails to close.

4.2.2 Data Reduction

Master Dark and Bias Frames and the Bad Pixel Map

With the same camera settings used for the survey and the blocking filter in place, we took 21 dark images, each with a 300 s exposure, and 21 bias images with the minimum $39 \mu\text{s}$ exposure. We mean-averaged the bias images to create the master bias frame. We mean-averaged the dark frames and subtracted the master bias to create the master dark frame used in the subsequent data reduction. From the master dark image, we created a hot pixel map of pixels with greater than 1000 ADU, roughly corresponding to the transition from read-noise-dominated to dark-current-dominated noise. This method identified 40986 pixels (0.27 per cent of the total) as hot according to this definition. In addition, we used CCDPROC's (Craig et al., 2015) CCDMASK function on an evening of flat-field images to identify pixels with a non-linear photo response. This identified 2054 pixels, of which 1011 were already present in the hot pixel map. The union of these maps resulted in the combined bad-pixel map.

Flat Field Corrections

To account for the regular appearance and disappearance of dust grains on the telescope's optical components (most visibly on the camera window), we created a new master flat frame for every night on which valid (greater than 30° lunar separation) flat field images were taken. The master flat frames were constructed by median stacking the individual dark- and bias-subtracted flat images which were themselves normalized by their median pixel values. With this method, we noted significant variations in the vignetting pattern of the master flat from night to night and especially from dusk to dawn, which we attributed to scattered light entering the optical system. In our original analysis, these variations manifested as jumps in the photometry from one night to the next at the ~ 1 per cent level.

While we expect the pattern of dust and debris in the flat images to evolve with time, the overall vignetting pattern is simply a product of the fixed optical geometry of the system – barring minute changes arising from focus shifts or temperature changes – and thus should remain constant. Thus, we devised a method whereby we fit and detrend the vignetting pattern of the master flats, while leaving the dust spots and the photo response non-uniformity (PRNU) information contained in the flats. This way, we can apply the same vignetting pattern to every master flat, preventing night-to-night variations while being sensitive to the appearance and disappearance of dust spots and a potentially evolving PRNU. The correcting process is illustrated in Fig. 4.3.




Figure 4.3: Overview of the flat-field correcting process. **a)** The untouched master flat from 2023 December 13 – the first night of observations. **b)** The flat image detrended by a bivariate spline fit. **c)** Cross-correlation map of the spline-detrended flat with an example dust doughnut. **d)** The spline-detrended flat overlaid with the locations of the detected doughnuts. **e)** The fit mask; masked areas are shown in white. **f)** The RBF-detrended master flat. **g)** The master vignetting pattern used for the whole survey. **h)** The corrected master flat image.

Dust and debris on the camera window create doughnut shapes in the flat field images with lower counts than the surrounding area. Masking these doughnuts from the fit was essential to capturing the underlying vignetting trend. For each night, we first detrended the master flat with a rectangular bivariate spline using `SCIPY` (Virtanen et al., 2020). We then cross-correlated the resulting image with a sharp, uncontaminated example of a dust doughnut to produce a cross-correlation map. The locations of the dust grains were then inferred from the positions of the peaks on the map. This approach leverages the fact that all the doughnuts are similar in size and shape. The dust locations were used to generate a fit mask, conservatively approximating the doughnuts as the region between two concentric circles with radii of 50 and 150 pixels, respectively. We then used the `RBF` package of Hines (2016) to interpolate the masked master flat with radial basis functions (RBFs). The final trend was calculated as the result of mean-averaging 30 interpolants, each generated using 1,000 samples of the unmasked region of the image.

After each night’s master flat was detrended, we restored the vignetting by multiplying by the trend from the flats of 2024 April 24 – the first night after which the camera window had been cleaned during the survey, and by which time the stray light-blocking truss sleeve had been installed – we refer to this as the master vignetting pattern. For uniformity, this master trend was used for all nights in the dataset. Where no valid flats were taken, the master flat frame from the closest available night was used.

The value of the RBF approach to detrending over a simple spline fit is demonstrated in panels **b)** and **f)** of Fig. 4.3 by the lack of artefacts around the doughnuts in the RBF-detrended flat image. In most cases, there is a visible artefact in the top left corner of the RBF-detrended flats, which is typical of RBF interpolation, but this did not impact any of the stars studied in this work.

Image Calibration

All science images were calibrated in five steps. First, the master bias was subtracted. Next, the master dark was subtracted, correcting for the difference in exposure times between the science and dark images. The resulting image was then divided by the relevant corrected master flat frame for that night. Fourth, a bad pixel correction was applied, in which each bad pixel value was replaced with the median of its (up to) eight surrounding good pixels.

Finally, for each calibrated science image, a background map was produced. This was generated by masking the centre of the cluster and any remaining stars in the image and fitting the remaining region. Star detection and background fitting were carried out using `PHOTUTILS` (Bradley et al., 2023). The background maps were subtracted from each of the calibrated images. This process, combined with the flat-field correction, minimized the impact of scattered light both during twilight flats and night-time science observations.

Astrometry

We derived an initial World Coordinate System (WCS) solution for each calibrated image by extracting sources with `SEP` (Bertin & Arnouts, 1996; Barbary, 2016) and providing the resulting source table to `ASTROMETRY.NET` (Lang et al., 2010). To model image distortion, we used a Simple Image Polynomial (SIP; Shupe et al. 2005) of degree 3.

To refine the initial `ASTROMETRY.NET` fit, we re-fitted the source positions against a catalogue of stars from *Gaia* DR3. The catalogued positions of the stars were adjusted by propagating their proper motions from the J2016.0 epoch to the time of the observations. We selected stars with magnitudes in the range $10 < G < 16$. Stars with neighbours brighter than three magnitudes dimmer ($\Delta G < 3$) within 22.5 arcseconds (50 binned pixels) were excluded to avoid centroid contamination. After applying these criteria, 110 *Gaia* stars in the M3 field were available for use in the astrometric fit.

The pixel coordinates of stars extracted by `SEP` were transformed into equatorial coordinates using the initial WCS solution and then cross-matched with the *Gaia* catalogue. Matches were refined by excluding those with catalogued-to-measured separations exceeding 10 pixels or zeropoint magnitudes deviating by more than 3σ from the mean. A degree 3 polynomial was then re-fitted to the filtered subset of matched stars to produce the final astrometric fit. This method for deriving the WCS solution consistently outperformed using `ASTROMETRY.NET` alone. To quantify this improvement, we analysed the images obtained on 2023 December 13 (the first night of observations). For these images, between 77 and 98 matched stars per frame were used in the final fits. Across all images, the mean and standard deviation of the median residuals between measured stellar positions and those predicted by the fitted WCS were 0.088 ± 0.009 arcsec using `ASTROMETRY.NET` alone, and 0.061 ± 0.007 arcsec using the improved method.

Additionally, at this stage in the pipeline, the non-variable (according to *Gaia* DR3), uncontaminated matched stars were used to measure a final colour-dependent zeropoint magnitude of the whole image by applying a linear fit to the zeropoint (catalogued minus instrumental) magnitudes of the individual stars in *G* against their $G_{BP} - G_{RP}$ colours.

Aperture Photometry

To facilitate the detection of transient objects and leverage existing information from *Gaia*, we combined a catalogue-driven and a non-catalogue-driven approach. Following image calibration, the sources were extracted with `SEP`. The signal-to-noise ratio (SNR) of each source in each image was calculated using the measured flux and the image-wide background RMS. Sources with $\text{SNR} < 5$ were rejected. To avoid blended and non-stellar sources, we rejected sources with ellipticity greater than 0.5 measured by `SEP`. We measured the half-flux diameters (HFDs) of the remaining sources and rejected images with median HFDs greater than 4 arcsec.

Each extracted source was cross-matched with the brightest *Gaia* source within 0.5 arcsec. Unmatched sources were retained for manual inspection as potential transients. However, upon review, all unmatched sources were identified as galaxies, asteroids, or camera defects.

Given the very high density of stars within the cluster, we approximated the level of contamination from other sources in the image by integrating the predicted contaminating flux over the photometric aperture of neighbouring *Gaia* sources given their magnitudes, colours and the colour-dependent zeropoint of the image, assuming Gaussian-shaped point spread functions (PSFs) with the measured median HFD. Sources with a predicted flux leakage of more than 1 per cent from other stars into the aperture were flagged as contaminated in that image. The contamination status of a source thus varied with the HFD of each image. In practice, the actual PSFs exhibited axial asymmetry and broader tails than a Gaussian, meaning this approach provided only an approximate indication of contamination.

Initially, we explored a difference imaging method using `HOTPANTS` (Becker, 2015), similar to

the approach taken by McCormac et al. (2014) in their survey of M71. However, TWIST’s significant susceptibility to wind shake produced complex PSFs that the software could not effectively handle. Very large photometric apertures would have been needed, increasing the noise and the impact of contamination, negating the benefit of the approach. Consequently, we employed simple aperture photometry for our analysis.

We tested a range of aperture radii from 2 to 15 (binned) pixels and measured the resulting SNR in the light curves. Larger apertures were beneficial for brighter stars, while smaller apertures yielded better results for dimmer stars and in bright conditions. Ultimately, we adopted a 5-pixel aperture radius for the entire survey, which was optimal or near-optimal for stars with $14 < G < 19$. This consistent aperture size also ensured homogeneous data processing.

The flux error for each star in each image was calculated based on the background RMS and Poisson shot noise. The error model did not account for other noise sources, such as atmospheric scintillation and HFD-dependent contamination.

In total, 4,199 stars were detected in at least one valid image, 2,694 in at least 1 per cent of valid images, 1,038 in at least 50 per cent and 372 in at least 99 per cent. Stars detected in fewer than 50 per cent of images were generally dimmer than around $G = 18.7$. The variation in the number of detections primarily reflects changes in seeing, sky transparency, and sky brightness, which affected the limiting magnitude between exposures.

We converted the fluxes of all the matched sources to G magnitudes using the colour-dependent zeropoint magnitude of each image. This conversion does not account for the star’s full spectral energy distribution, intrinsic colour variations, or non-linearity in the zeropoint-colour relation, but it offers a practical approach for handling large stellar samples where a precise absolute photometric calibration is not required.

Additionally, we ran forced photometry with the same 5-pixel aperture size on the proper motion-propagated positions of all the known RRL variables within the field from Clement et al. (2001) after cross-matching with *Gaia* DR3. We excluded stars for which we could not find a cross-match in *Gaia* DR3, specifically V187, V244, V248, V251, and V265, as *Gaia* photometric parameters were required for the subsequent analysis. These excluded sources lie 11–32 arcsec from the cluster centre in regions of severe crowding. We also excluded V297, identified as a foreground non-RRL variable (Bhardwaj et al., 2020). Finally, four variables (V123, V205, V206, and V299) fell outside the survey field of view, leaving 231 stars in the final sample.

4.3 Periodic Variable Analysis

This section describes the search for and analysis of periodic variables in the direction of the cluster.

4.3.1 Periodicity Search

On all of the corrected light curves, we performed a Lomb-Scargle (LS – Lomb, 1976; Scargle, 1982) periodogram using the implementation of *ASTROPY* (Astropy Collaboration et al., 2013, 2018, 2022). We followed the prescription of VanderPlas (2018) and set a frequency grid with spacing, $\Delta f = 1/n_o T = 0.0005 \text{ d}^{-1}$ where T is the length of the survey, $\sim 200 \text{ d}$, and $n_o = 10$ is the number of test frequencies per peak, set to the same value as in Richards et al. (2012) in their work on

classifying variable stars. We set a minimum frequency of 0.0005 d^{-1} and a maximum frequency of 60 d^{-1} , allowing the detection of SXP stars' pulsations, including higher harmonics. In most cases, some residual signal remained from the detrending with frequencies of integer multiples of 1 d^{-1} . We masked out frequencies within 5 per cent of 1 d^{-1} and within 2 per cent of 2, 3 and 4 d^{-1} . This resulted in masking out the primary frequencies of some RRL stars. However, the large amplitudes of their non-sinusoidal pulsations allowed detection at higher harmonics. We additionally masked frequencies within 5 per cent of the lunar synodic frequency, $1/29.53 \text{ d}^{-1}$.

We visually inspected all stars with a full amplitude at their peak-power frequency of greater than 0.01 mag and rejected obviously spurious signals. The remaining periodic variables were all matched to currently known variable stars in the catalogue of Clement et al. (2001), with the exception of four. We classified these stars as SXP variables and designate them V301–304, from brightest to dimmest in *G*, following on from the discovery of the long-period variable V300 by Siegel et al. (2015). Finding charts for these stars can be found in Fig. 4.4.




Figure 4.4: Finding charts showing the locations of the newly discovered SXP-type variables V301–304. Each chart is 2×2 arcmin across. The images were created by aligning with SPALPY (Lyman, 2021) and stacking 360 calibrated frames (totalling one hour of exposure) with measured half-flux diameters less than 2 arcsec taken during dark sky environments. North is up and East is to the left.

Table 4.1: Relevant photometric and astrometric properties of the four newly discovered SXP type variables from *Gaia* DR3. Positions and proper motions are with respect to the ICRS reference frame at the J2016.0 epoch. $\mu_\alpha \cos \delta$ is the proper motion in right ascension corrected for declination, μ_δ is the proper motion in declination. The V band magnitude is derived from the sources’ G and $G_{BP} - G_{RP}$ magnitudes using the relation from Chapter 5.5.1 of version 1.3 of the *Gaia* DR3 documentation (Busso et al., 2022).

V	<i>Gaia</i> DR3 ID	RA (ICRS)	Dec (ICRS)	$\mu_\alpha \cos \delta$ (mas/yr)	μ_δ (mas/yr)	Parallax (mas)	G (mag)	$G_{BP} - G_{RP}$ (mag)	V (mag)
301	1454778304460743936	$13^{\text{h}}41^{\text{m}}54.^{\text{s}}85$	$+28^\circ12'28.''6$	-0.24 ± 0.11	-2.51 ± 0.07	0.05 ± 0.10	17.61	0.39	17.66 ± 0.03
302	1454828912055208192	$13^{\text{h}}43^{\text{m}}11.^{\text{s}}42$	$+28^\circ22'32.''2$	-0.38 ± 0.16	-2.77 ± 0.10	0.18 ± 0.13	18.20	0.32	18.24 ± 0.03
303	1454730746283291264	$13^{\text{h}}42^{\text{m}}20.^{\text{s}}67$	$+28^\circ11'34.''8$	-0.19 ± 0.16	-2.69 ± 0.09	0.01 ± 0.14	18.26	0.31	18.30 ± 0.03
304	1454782595128645376	$13^{\text{h}}41^{\text{m}}22.^{\text{s}}86$	$+28^\circ16'32.''5$	-0.18 ± 0.17	-2.56 ± 0.11	-0.10 ± 0.16	18.27	0.36	18.32 ± 0.03

4.3.2 SX Phoenicis Variables

SXP variables form a metal-poor subclass of δ Scuti variables. They lie on or above the zero-age main sequence in the instability strip of the Hertzsprung–Russell (HR) diagram with typically A or F spectral types. They follow a period–luminosity relation (e.g. McNamara 2011, Cohen & Sarajedini 2012) which has been used to estimate distances to globular clusters and nearby galaxies. Compared to the more metal-rich δ Scuti variables, SXP stars are less massive, less luminous, and have, on average, shorter pulsation periods.

In this section, we demonstrate that the variables V301–304 are cluster members, that they are members of the SXP class, and, using their pulsation frequencies, provide an independent estimate of the distance to M3.

We cross-matched the four new SXP variables with their corresponding *Gaia* DR3 sources and show their associated source IDs and astrometric and photometric properties in Table 4.1. Each has a parallax consistent with zero to 2σ , indicating they are not nearby field stars. Additionally, V301–304 have angular separations from the cluster centre of 10.8, 13.2, 11.2, and 12.3 arcmin, all well within M3’s 28.7 arcmin tidal radius (Harris, 2010), indicating that they are coincident with the cluster.

According to Libralato et al. (2022), the absolute proper motion of M3 is $(\mu_\alpha \cos \delta, \mu_\delta) = (-0.261 \pm 0.055, -2.674 \pm 0.040)$ mas/yr. To assess the consistency of each star’s motion with that of the cluster, we computed the two-dimensional Mahalanobis distance (Mahalanobis, 1936) between the star and cluster proper motions, treating the right ascension and declination components as statistically independent and combining the star and cluster uncertainties in quadrature.

Given that the squared Mahalanobis distance follows a χ^2 distribution with two degrees of freedom, we calculated the probability (p -value) of observing larger differences in proper motion under the assumption that each star shares M3’s motion. For V301–304, we found $p = 0.11, 0.53, 0.89$, and 0.56 respectively. In all cases, these probabilities indicate that the stars’ proper motions are statistically consistent with that of M3, given the measurement uncertainties.

On the basis that each star is coincident with the cluster within the sky-projected tidal radius, has a proper motion consistent with the cluster, is not nearby, and is located far from the galactic plane, away from large numbers of field stars, we verify the membership of M3 of variables V301–304.

With their cluster membership confirmed, we place the variables on the colour–magnitude diagram (CMD; shown in Fig. 4.5), constructed from the *Gaia* sources in the survey field with proper motions within three times the central velocity dispersion of M3’s proper motion (taken from Libralato et al. 2022). The δ Scuti instability region of the CMD was obtained by transforming the empirical Hertzsprung–Russell (temperature–luminosity; $T_{\text{eff}}-L$) diagram boundaries of Murphy et al. (2019):

$$\log(L/L_\odot) = -0.000811 T_{\text{eff}} + 6.672 \quad (\text{red edge}), \quad (4.1)$$

$$\log(L/L_\odot) = -0.001000 T_{\text{eff}} + 10.500 \quad (\text{blue edge}), \quad (4.2)$$

into *Gaia* photometric coordinates – firstly into intrinsic absolute magnitude in the G passband, M_G , against intrinsic colour, $(G_{\text{BP}} - G_{\text{RP}})_0$, and then into apparent G magnitude against apparent $G_{\text{BP}} - G_{\text{RP}}$ colour.

The relations from Murphy et al. (2019) were evaluated over $0.5 \leq \log(L/L_\odot) \leq 1.4$, ensuring

the instability region boundaries fall within the temperature range over which the empirical boundaries were calibrated (6,500–10,000 K). Effective temperatures were converted to intrinsic $(G_{\text{BP}} - G_{\text{RP}})_0$ colours using a solar-metallicity PARSEC model grid (Bressan et al., 2012) in the *Gaia* EDR3 photometric system (Riello et al., 2021), adopting $\log g = 4.0$ and $[M/H] = 0.0$ as representative of the Population I δ Scuti stars used to define the boundary. Luminosities were converted to absolute *Gaia* magnitudes using

$$M_G = M_{\text{bol}} - BC_G = 4.74 - 2.5 \log(L/L_\odot) - BC_G, \quad (4.3)$$

where M_{bol} is the absolute bolometric magnitude (whose zeropoint is defined by Mamajek et al. 2015) and BC_G is the bolometric correction in the G band. The intrinsic colours and G -band bolometric corrections were interpolated from the PARSEC model grid at the effective temperature of each point along the instability-strip boundaries.

Finally, the boundaries were shifted into the observed CMD. M_G was modified by the true distance modulus, $(m - M)_0 = 15.04$, from Bhardwaj et al. (2020), and the G -band extinction, A_G , to produce the apparent G magnitude:

$$G = M_G + (m - M)_0 + A_G, \quad (4.4)$$

where, from Casagrande et al. (2021), $A_G = R_G E(B - V)$, with

$$R_G = 3.071 - 0.511(G_{\text{BP}} - G_{\text{RP}})_0 + 0.058(G_{\text{BP}} - G_{\text{RP}})_0^2. \quad (4.5)$$

For this relation, we adopt $E(B - V) = 0.013$ from VandenBerg et al. (2016). The apparent $G_{\text{BP}} - G_{\text{RP}}$ colours were calculated by correcting the intrinsic $(G_{\text{BP}} - G_{\text{RP}})_0$ colours for the expected reddening, $E(G_{\text{BP}} - G_{\text{RP}})$:

$$G_{\text{BP}} - G_{\text{RP}} = (G_{\text{BP}} - G_{\text{RP}})_0 + E(G_{\text{BP}} - G_{\text{RP}}). \quad (4.6)$$

The *Gaia* colour excess, $E(G_{\text{BP}} - G_{\text{RP}})$, was calculated from the difference between the extinction in the two passbands,

$$E(G_{\text{BP}} - G_{\text{RP}}) = A_{\text{BP}} - A_{\text{RP}} = (R_{\text{BP}} - R_{\text{RP}})E(B - V), \quad (4.7)$$

where the colour-dependent coefficients R_{BP} and R_{RP} were taken from Table B1 of Casagrande et al. (2021) and evaluated at the intrinsic colour of each point along the instability-strip boundaries. The enclosed area in the CMD is an approximate transformation of the empirical δ Scuti instability region, which itself is not a sharply defined boundary.

V301–304 are located brighter and bluer than the main-sequence turn-off in the blue straggler region, indicating that they are themselves blue straggler stars. As expected for SXP stars (e.g. Balona & Nemec 2012), they fall within or immediately adjacent to the δ Scuti instability strip, and consistent with the results of Murphy et al. (2019) for δ Scuti types, they fall closer to the red boundary.

The four new variables have detected peaks in their LS periodograms significantly greater than the 0.1 per cent false alarm probability (FAP) threshold at frequencies that could have no apparent systematic, non-astrophysical cause. The frequencies of the highest amplitude oscillations of each




Figure 4.5: Colour–magnitude diagram of the *Gaia* sources in the survey field with proper motions within three times the central velocity dispersion of the mean proper motion of the cluster (values taken from Libralato et al. 2022), which we take to signify cluster membership in the absence of a more detailed analysis. M_G is the absolute G magnitude assuming a distance modulus of $(m - M)_V = 15.07$ per Harris (2010) and zero colour excess, $E(V - G)$. The locations of the newly discovered SXP-type variables are shown as open squares. Other variables are indicated as in Fig. 4.2. The δ Scuti instability region is taken from Murphy et al. (2019), evaluated in the interval $0.5 \leq \log(L/L_\odot) \leq 1.4$ and transformed into *Gaia* (E)DR3 photometric coordinates per Section 4.3.2. The opacity of each data point is given by the proportion of images in which its corresponding star is detected with at most 1 per cent contamination and at least at 5σ .

star fall in the interval given by McNamara (2011) for SXP stars -7.9 – 29.5 d^{-1} . V301 falls in the overlapping region of the period ranges for SXP and δ Scuti stars, but with a metallicity, $[\text{Fe}/\text{H}]$, of -2.04 dex, derived from *Gaia* DR3’s Astrophysical parameters inference system (Apsis; Creevey et al., 2023), and the simple statement that they are members of a globular cluster, the SXP classification is strongly preferred. Since the four stars are blue straggler cluster members with pulsation periods typical of SXP stars, we confirm their variable classification.

We searched for evidence of radial mode pulsation in the light curves of the SXP stars by considering the frequencies present. We fit an n -term Fourier sine series of the form:

$$m = m_0 + \frac{1}{2} \sum_{i=1}^n A_i \sin(2\pi i f t + \phi_i) \quad (4.8)$$

to the light curve (in magnitudes) of each star at its peak LS power frequency, f . m_0 is the constant offset, A_i are the (full) amplitudes, ϕ_i are the phases, and t is barycentric time. The derived Fourier parameters are then given by $\phi_{ij} = j\phi_i - i\phi_j$ and $R_{ij} = A_i/A_j$.

We performed an MCMC fit using `EMCEE` (Foreman-Mackey et al., 2013), using unbounded uniform priors and allowing the frequency to vary as a free parameter. We used as many sine terms in the fit as had an amplitude detected with SNR of at least 2. We then pre-whitened the light curves by the Fourier fit and computed new LS periodograms. We repeated this process until the highest remaining peak was due to residual signal from the largest amplitude pulsation (V301–302) or there were no remaining peaks above the 0.1 per cent FAP threshold (V303–304). The Fourier fits and periodograms are shown in Fig. 4.6 and 4.7, and the extracted Fourier parameters are shown in Table 4.2.




Figure 4.6: Lomb-Scargle periodograms (left panels) and phased light curves (right panels) of the newly discovered SX Phoenicis pulsators V301 and 304, showing all of the extracted frequencies for each star. After each frequency is detected in the periodogram, the data are pre-whitened by subtracting a best-fit (multiple harmonic where possible) sinusoidal model of that frequency. The periodogram is recalculated after each pre-whitening step to extract additional frequencies. Vertical dashed lines in the periodograms mark the extracted frequencies, and in cases of multiple harmonic fits, the higher-order harmonics are also marked with dashed vertical lines. The horizontal dashed lines in the periodograms represent the 0.1 per cent false-alarm probability threshold. The phased light curves (right) show the data folded over each extracted frequency, with the best-fit models overlaid in red. For visibility, each light curve is binned by phase into 100 bins because the uncertainty of each measurement is, in most cases, considerably larger than the amplitude of the oscillation. Each light curve is normalized to the peak brightness of the fit.




Figure 4.7: The same as Fig 4.6 for V302 and V303.

Table 4.2: Fourier properties of the newly discovered SXP stars. Magnitudes are in the instrumental system of TWIST. Amplitudes are the full peak-to-peak variations. Where only a single harmonic is detected, the total amplitude is identical to A_1 .

V	Label	Frequency (d^{-1})	Total amplitude (mag)	A_1 (mag)	ϕ_{21} (rad)	R_{21}	ϕ_{31} (rad)	R_{31}
301	f_1	19.27057 ± 0.00001	0.1487 ± 0.0006	0.1452 ± 0.0006	2.30 ± 0.03	0.155 ± 0.004		
	f_2	19.39537 ± 0.00004	0.0436 ± 0.0006	0.0430 ± 0.0006	2.35 ± 0.12	0.112 ± 0.013		
	f_3	19.04511 ± 0.00009		0.0217 ± 0.0006				
	f_4	38.66771 ± 0.00016		0.0135 ± 0.0006				
	f_5	19.62343 ± 0.00016		0.0112 ± 0.0006				
	f_6	32.75288 ± 0.00021		0.0098 ± 0.0006				
302	f_1	26.13173 ± 0.00002	0.1918 ± 0.0013	0.1865 ± 0.0009	2.04 ± 0.03	0.178 ± 0.005	4.16 ± 0.15	0.034 ± 0.005
	f_2	27.07126 ± 0.00006	0.0531 ± 0.0009	0.0530 ± 0.0009	1.87 ± 0.16	0.113 ± 0.018		
	f_3	53.20279 ± 0.00019		0.0158 ± 0.0009				
	f_4	34.60693 ± 0.00020		0.0147 ± 0.0009				
	f_5	33.29425 ± 0.00028		0.0110 ± 0.0009				
303	f_1	26.28969 ± 0.00026		0.0107 ± 0.0010				
	f_2	27.04032 ± 0.00031		0.0096 ± 0.0009				
	f_3	50.76065 ± 0.00039		0.0091 ± 0.0011				
304	f_1	24.13592 ± 0.00016		0.0207 ± 0.0010				

In all periodograms, we observe frequency aliases at multiples of 1 d^{-1} separations from the highest peaks, which can be attributed to the observing window function. In many cases, this makes the identification of the correct peak uncertain. The 10-second survey cadence improves the definition of the true pulsation peak by reducing high-frequency alias contamination and providing denser sampling of each cycle, although daily 1 d^{-1} sidelobes from the window function remain unavoidable and complicate peak selection in several cases. With the exception of f_1 and f_2 of V301 and V302, whose higher harmonics constrain the correct peaks, the extracted frequencies may differ from the intrinsic frequencies by 1 or 2 d^{-1} , which would be a much larger error than the fit uncertainties. This is particularly relevant to V303, which has three other peaks with at least 90 per cent of the LS power of the most significant peak.

From Olech et al. (2005), we expect the period ratios of the first few adjacent radial modes to be in the vicinity of the interval from 0.77 to 0.85 (see their Fig. 6). We identified two pairs of frequencies with (inverse) ratios in this interval. Firstly, f_4 and f_6 of V301 have a ratio of 0.847, which matches the third to second overtone period ratio according to the models of Olech et al. (2005). However, f_4 itself appears to be almost exactly equal to $f_1 + f_2$, meaning it is not an independent mode (and especially not a radial mode) but instead arises from non-linear interactions between f_1 and f_2 .

In contrast, f_1 and f_5 of V302 have a ratio of 0.785, which is characteristic of the fundamental and first overtone radial modes. Given that V302 also has a relatively large amplitude at f_1 , we conclude that its largest amplitude variability is due to pulsation in the fundamental radial mode. Knowing this, we can use a period–luminosity relation for fundamental mode pulsators to estimate the distance modulus to V302. Using the relation from Chapter 5.5.1 of version 1.3 of the *Gaia* DR3 documentation (Busso et al., 2022), we estimate from the *Gaia* photometry the apparent V magnitude of the variables (shown in Table 4.1), and using equation 5b of McNamara (2011) for the metal-poor fundamental mode period–luminosity relation, we calculate the predicted absolute magnitudes in the V band. Propagating the statistical and systematic uncertainties, we estimate a distance modulus $(m - M)_V$ to V302 (and therefore M3) of 15.02 ± 0.10 , in excellent agreement with the literature value of 15.07 from the catalogue of Harris (2010).

From VandenBerg et al. (2016), we adopt a reddening of $E(B - V) = 0.013$ and assume the canonical diffuse-ISM value $R_V = 3.1$ (Cardelli et al., 1989). This yields a V -band extinction, $A_V = 0.040$ and a true distance modulus of $(m - M)_0 = 14.98 \pm 0.10$, which is in excellent agreement with the more modern true distance modulus estimate of $\mu = 15.041 \pm 0.017$ (statistical) ± 0.036 (systematic) of Bhardwaj et al. (2020). The uncertainties in the slope and intercept of the period–luminosity relation dominate the uncertainty in the distance moduli. Although Pych et al. (2001) and Olech et al. (2005) have demonstrated that radial pulsations dominate for SXP stars even at low amplitudes, we refrain from assuming radial pulsations for estimating the distances to the other stars when evidence from the frequency ratios is not available.

We note that variables V301 and V302 are flagged as variable stars in *Gaia* DR3 and are broadly classified as belonging to one of the following types: γ Doradus, δ Scuti, or SX Phoenicis by the classifier of Rimoldini et al. (2023). However, their classifier does not distinguish between these specific classes for these sources. Within the variability tables of *Gaia* DR3 (Gaia Collaboration, 2022), a primary frequency is listed for V301 ($19.27084 \text{ day}^{-1}$, with no stated uncertainty), which is in agreement with the value derived in this work ($19.27057 \pm 0.00001 \text{ day}^{-1}$). No frequency is

reported for V302. The newly identified variables V303 and V304 are not flagged as variable in *Gaia* DR3, likely owing to their faintness and the low amplitude of their variability, demonstrating that the TWIST observations are sensitive to sources below the *Gaia* variability detection threshold.

V301–304 are the first new SXP stars to be discovered in M3 since the discovery of six by Hartman et al. (2005). The present study increases the number of known SXP variables in the cluster from nine to thirteen.

4.3.3 RR Lyrae Variables

RR Lyrae stars are low-mass metal-poor radial pulsators on the horizontal branch of the HR diagram in the classical instability strip. They are commonly found in globular clusters, and M3 hosts the most known RRL stars of any globular cluster in the Milky Way (Clement et al., 2001).

Period and Amplitude Determination

Incidental to our search for new variability in the cluster, we ran forced photometry on 231 previously discovered RRL stars within the survey field with cross-matches to *Gaia* sources. We studied these stars for evidence of period changes, the appearance or disappearance of pulsation modes, and Blazhko-like amplitude modulation. Given the non-sinusoidal shape of RRL pulsations, we used the phase dispersion minimization technique (PDM; Stellingwerf 1978) to identify periodicity in the extracted light curves. Specifically, we used the PDM2B data-rich implementation, which attempts to measure period and amplitude changes.

For each star, we ran the PDM analysis in two different modes: first, without segmenting the data to obtain a best estimate of the mean period and amplitude, and second, segmenting the data into overlapping windows of 7,500 contiguous points, sliding each time by 500 points to measure period and amplitude modulation. Full phase coverage within the segmented intervals was essential for resolving time-dependent amplitude variations, and this was substantially aided by the 10-second cadence of the TWIST observations. In both the segmented and non-segmented cases, we tested frequencies within 0.02 d^{-1} of the highest peak of an initial LS search. Each star had around 106,000 data points passing basic quality checks – less than four arcsec HFD and mean zeropoint magnitude greater than 22.7.

We derive the periods from the lowest θ statistic period of the whole-data PDM runs. We used bootstrap resampling (with replacement) to generate 500 simulated light curves and PDM period samples for each star. We derive the uncertainties in the periods from the standard deviation of the distribution of extracted periods.

Overall, the photometry was of sufficient quality such that 203 of the 231 RRL stars had one visually distinct peak in the PDM periodogram over the tested frequency range and an associated $\text{FAP} < 0.001$ according to the statistic of Schwarzenberg-Czerny (1997). The periods of these stars are shown in Table 4.3. The remaining 28 stars, located 9.5–116 arcsec from the cluster centre, were strongly affected by crowding and blending, to the extent that reliable period determination was not possible.

Since the majority of the variable stars were heavily contaminated in the dense field, to measure the amplitudes accurately, for each star we subtracted whatever constant flux was needed to make its zeropoint corrected flux-averaged magnitude identical to its catalogued *Gaia* DR3 *G* magnitude.

Table 4.3: Best fitting periods and amplitudes for the 203 observed RRL variables. The Blazhko flag indicates a strong positive detection of Blazhko-like modulation *in our data*. Amplitudes are in the instrumental system of TWIST. The full table is available as supplementary material in machine-readable format in Mitchell et al. (2026a)

V	Class	Bl.	Period (d)	Amplitude (mag)
1	<i>ab</i>	N	0.5205813 ± 0.0000014	1.13 ± 0.04
3	<i>ab</i>	N	0.558206 ± 0.000003	1.14 ± 0.04
4n	<i>ab</i>	N	0.585065 ± 0.000005	1.03 ± 0.04
4s	<i>ab</i>	N	0.593010 ± 0.000009	1.05 ± 0.04
5	<i>ab</i>	Y	0.505856 ± 0.000004	0.76 ± 0.03
...	...	...	...	...
261	<i>c</i>	N	0.444723 ± 0.000012	0.41 ± 0.02
266	<i>c</i>	N	0.342354 ± 0.000012	0.53 ± 0.03
270n	<i>ab</i>	N	0.62593 ± 0.00002	0.57 ± 0.03
270s	<i>ab</i>	N	0.690186 ± 0.000006	0.87 ± 0.03
290	<i>d</i>	N	0.240398 ± 0.000003	0.0387 ± 0.0019

This introduces a systematic uncertainty arising from the uncertainty in the absolute photometric calibration (the scatter about a linear fit to the data in Fig. 3.8 – 0.0413 mag), but permits amplitude measurements of highly contaminated stars, including 16 with no published amplitudes. The resulting systematic uncertainty is typically around ~ 3 -5 per cent of the amplitude in magnitudes and is the dominant source of uncertainty. We were able to measure the amplitudes of 198 of the stars with firmly detected periods; the remaining five have periods very close to one half or one third of a day, which either precluded complete phase coverage or artificially inflated the extracted amplitude owing to residual, nightly observing-induced periodicity with frequencies multiples of 1 d^{-1} . Representative phase-folded light curves for the RR Lyrae stars with successfully determined periods and amplitudes are presented in Fig. 4.8.

Comparison with Catalogued Properties

The extracted oscillation amplitudes are compared to those in the catalogue of Clement et al. (2001) in Fig. 4.9. A 3σ sigma-clipped linear fit to the data gives a slope of ~ 0.873 , a vertical intercept of -0.013 mag, and negligible formal fit uncertainties. We observe smaller amplitudes in our instrumental band than the catalogued V amplitudes, likely owing to TWIST having greater red coverage than the V band. We do not expect a perfect one-to-one correspondence between the two sets of amplitudes, since intrinsic factors such as colour, temperature, and evolutionary phase within the instability strip introduce natural scatter. The observed correlation nonetheless confirms that the TWIST amplitudes are physically sensible and can therefore be used as a basis to estimate the V -band amplitudes of stars lacking published values, including the 16 stars for which we measured amplitudes but no V -band amplitude is available.

Of the six most significant outliers (labelled in Fig. 4.9), four of them (V35, V44, V106, and V130) show clear and significant Blazhko modulation, which contributes to the discrepancies. The remaining outliers, V246 and V254, are more likely to be statistical in nature, as their derived periods are consistent with the catalogued values. The discrepancies are likely due instead to larger measurement uncertainties caused by significant blending in the crowded central field.




Figure 4.8: Phased light curves for representative RR Lyrae stars observed with TWIST, illustrating the variety of pulsation behaviours. The upper panels show two fundamental-mode (RRab) stars, one stable and one showing Blazhko modulation. The lower panels show a first-overtone (RRc) and a double-mode (RRd) star. These examples include V13, V44, and V125, whose classifications are discussed in Section 4.3.3. Phase is shown modulo one pulsation cycle, and relative brightness is plotted in magnitudes relative to the minimum magnitude of the whole-data PDM_{2B} fit. The data are binned to 0.007 d (~ 10 -minute) bins prior to folding and the error bars are too small to see.




Figure 4.9: Comparison between the catalogued V amplitudes of the studied RRL stars from Clement et al. (2001), with the extracted amplitudes of this study in the instrumental bandpass of TWIST. Outliers with residuals greater than 0.35 mag are labelled with their SAW IDs.

Extracted periods generally matched the catalogued values of Clement et al. (2001) to the precision with which they are quoted therein, and in all cases were within 1 per cent of the catalogued period. The variables V29, V41, V106, and V270n (all *RRab* type) stand out as deviating from the catalogue values by more than one minute. The period for V29 in Clement et al. (2001) is listed as 0.47 d on account of the discrepancy between the measurements of Corwin & Carney (2001) and Benkő et al. (2006). Our measured period is inconsistent with Corwin & Carney (2001) by 52.7 seconds (~ 7.6 times our measurement uncertainty) but consistent with Benkő et al. (2006) to within 3.3 seconds ($\sim 0.5\sigma$). Similarly, the period of V270n is listed to only two significant figures, but our measurement is consistent with Jurcsik et al. (2017) to 6.9 seconds ($\sim 4\sigma$). For V41, our measured period is inconsistent with Benkő et al. (2006) by 72 seconds ($\sim 21\sigma$), but consistent with the value given in Jurcsik et al. (2017) to 5.2 seconds ($\sim 1.5\sigma$).

For V106, Jurcsik et al. (2017) report a fundamental pulsation period, $P_J = 0.54689$ d, while we extract $P_{\text{our}} = 0.542424$ d – ~ 385 s shorter. While period changes are expected in RRL stars, such a reduction on a timescale of roughly a decade is physically implausible (Jurcsik et al., 2012). We note that the difference between the two extracted frequencies, $f_{\text{our}} - f_J = 0.0151 \text{ d}^{-1}$, is within a few per cent of the Blazhko frequency given in Jurcsik (2019), $f_{\text{Bl}} = 0.0157 \text{ d}^{-1}$, and therefore we interpret our extracted period as a Blazhko side peak rather than a genuine decrease in the fundamental pulsation period.

Following visual inspection of the phase-folded light curves, we agree with the RRL classifications given in Clement et al. (2001), except for V44 and V125, both of which they label as *RRd*-type, and which we classify as *RRab* and *RRc* respectively. In the case of V44, we agree with the reclassification of Jurcsik (2019) on the basis that no multiple periodicity is seen and that the characteristic sawtooth shape of an *RRab* star is present.

In the case of V125, we disagree with its *RRd* classification in the catalogue, which is taken from the work of Jurcsik et al. (2015), and find no reclassification in any study since then. We note that Bhardwaj et al. (2020) observed V125 as part of their near-infrared survey of M3’s RR Lyrae population in 2019 May and retained the multi-mode classification. In our data, only a singular periodicity is visible in the light curve, with the characteristic sinusoidal shape of an *RRc* pulsator. We searched for evidence of fundamental-mode pulsation in the LS periodogram of its first overtone-mode prewhitened light curve at the expected frequency given in Jurcsik et al. (2015), but we found none. Additionally, in the same reference, a second-overtone mode frequency is tenuously provided for the star V13, which, after pre-whitening by the fundamental and first-overtone modes, also cannot be detected in our data. This is despite the densely sampled and high-precision photometry. Thus, we consider that the fundamental mode of V125 and the second overtone mode of V13 (if ever present) have disappeared. Phase-folded light curves for V13, V44, and V125 are all shown in Fig. 4.8.

Blazhko Modulation

First described by Blažko (1907), the quasi-periodic phase and amplitude modulation of variable pulsators is known as the Blazhko effect. This effect, although poorly understood, has been observed in around half of the *RRab* stars in M3 (Jurcsik et al., 2012), which is a typical fraction for Oosterhof Type I clusters (e.g. Jurcsik et al. 2011; Smolec et al. 2017). With our limited photometric precision for contaminated stars, we confidently recovered 53 stars showing visible, unambiguous amplitude

modulation. In most cases, Jurcsik et al. (2012) lists these stars as positively identified modulators, with the exception of the RRc stars V126 and V203, which were later identified by Jurcsik et al. (2015).

To characterize the modulation of the Blazhko stars, we used the results of the sliding-window PDM2B analysis, which returns a period and amplitude for one 7500 data point window of the light curve at a time, sliding by every 500 points between measurements. We rejected measurements that spanned a range of more than 15 days to ensure the measurements better describe the instantaneous properties of the light curve. We additionally rejected data points for which the largest gap in the pulsation period-folded light curve was greater than 0.05 to reduce the probability of spuriously low amplitude measurements arising from incomplete phase coverage. We were not generally able to measure period modulation, and hence, we omit any such analysis. We fit the amplitude modulations each to a single harmonic sine wave and visually inspected the fits. This resulted in good fits for 28 stars, whose Blazhko phase-folded light curves are shown in Fig. 4.10 and whose fit parameters are listed in Table 4.4. For the other stars, either there were too few amplitude measurements on account of the phase coverage requirement, the modulation period was too long relative to the survey baseline, or the fits did not convincingly pass visual inspection. In the case of V35, the very large amplitude modulation (~ 0.75 mag) could not be approximated as sinusoidal.

All of the Blazhko stars for which we could fit the amplitude modulation are of type RRab and have published modulation periods in Jurcsik (2019), which are shown in Table 4.4. Although no uncertainties are given for the periods in that reference, we are generally in agreement for the stars with periods less than around 100 d. Although the longer period modulators may have evolved, the discrepancies are likely due to a lack of Blazhko phase coverage and an over-reliance on the not perfectly applicable sinusoidal model, resulting in systematic uncertainties that are not captured by the fit uncertainties. The modulators for which we detect a significantly different period from the literature values are V24, V66, and V79. The detected period for V24 is very close to half the literature value, and the light curve is similarly well-fit by a 50.4 d period. Due to the period aliasing, we cannot confidently rule out either period, although the 50.4 d-folded light curve visually appears to have a greater amplitude, making it more likely. Our observations of V66 and V79 are clearly incompatible with the periods reported by Jurcsik (2019), indicating that their modulation periods have changed since that study. Such changes are consistent with the quasi-periodic nature of the Blazhko effect and have been observed in RR Lyr itself (Borgne et al., 2014).

We note that Skarka et al. (2020) introduced a morphological classification to the Blazhko-phase folded light curves in the *I* band; however, we refrain from classifying the observed Blazhko modulators according to this system due to the significantly different bandpass used in the present survey.

4.3.4 Unclassified Variables

Two of the variable stars recovered in our survey are the unclassified variables V286 and V287. These stars were initially identified by Hartman et al. (2005), who had insufficient phase coverage to determine reliable periods or amplitudes. Both variables were later assessed as highly probable cluster members in the membership study by Prudil & Arellano Ferro (2024). Their positions on the colour-magnitude diagram are shown in Fig. 4.5, to the right of the giant branch.




Figure 4.10: Blazhko phase-folded light curves of the modulators listed in Table 4.4. The folded light curves are centred on the maximum pulsation amplitude of the sine fit, and the data are binned to 0.007 d (~ 10 -minute) bins prior to folding. The SAW IDs and the best-fit modulation periods are shown in each plot.

Table 4.4: Fit parameters for a sinusoidal fit to the amplitude modulation of the constrained Blazhko stars. The uncertainties given are the 1σ scatter in the MCMC fits. The modulation amplitude is the full peak-to-peak range of pulsation amplitudes. For comparison, we show the 2012 modulation periods from Jurcsik (2019) (labelled J19) for these stars. Multiple modulation periods are separated by forward slashes, and lower bounds are denoted with colons.

V	Mod. Period (d)	J19 Mod. Period (d)	Mod. Amp (mag)
10	175 ± 10	191	0.120 ± 0.007
14	61.0 ± 0.3	60.6	0.317 ± 0.006
23	115.6 ± 0.6	125	0.405 ± 0.006
24	25.1 ± 0.7	49.8	0.22 ± 0.02
28	64.3 ± 0.3	65.9	0.623 ± 0.013
33	41.18 ± 0.19	40.5/34.8:	0.59 ± 0.02
38	51.9 ± 0.4	51.8	0.60 ± 0.03
41	120 ± 5	190	0.147 ± 0.010
43	98.7 ± 1.0	103	0.387 ± 0.015
45	26.88 ± 0.11	26.6	0.426 ± 0.017
47	45.5 ± 0.3	69/53.3/44.6	0.50 ± 0.03
48	122.8 ± 1.8	143	0.121 ± 0.005
50	178 ± 12	178:	0.204 ± 0.011
59	62.5 ± 0.3	62.1	0.238 ± 0.008
61	167 ± 4	134:	0.505 ± 0.010
63	58.9 ± 0.3	58.7	0.390 ± 0.009
66	58.6 ± 0.5	65.9/49.3	0.283 ± 0.009
67	104.4 ± 0.4	103	0.618 ± 0.005
71	45.3 ± 0.2	46.7	0.355 ± 0.017
73	50.7 ± 0.7	54.3	0.073 ± 0.005
78	39.7 ± 0.2	39.5	0.387 ± 0.015
79	87.1 ± 0.6	160/62.9/18.3:	0.530 ± 0.019
104	119 ± 2	110	0.131 ± 0.005
110	45.50 ± 0.17	44.9	0.347 ± 0.012
114	55.9 ± 0.3	54.5	0.078 ± 0.002
117	46.8 ± 0.2	47.2	0.314 ± 0.008
119	56.5 ± 0.4	56.9	0.439 ± 0.010
121	66.6 ± 0.6	65.1	0.45 ± 0.02




Figure 4.11: Light curves of the unclassified variables V286 and V287 relative to their peak brightnesses in the instrumental bandpass of TWIST. The data are binned to 14.4 min (0.01 d) bins.

V287 was subsequently observed in the near-UV by Siegel et al. (2015) using the Ultra-Violet and Optical Telescope (UVOT) on *Swift*; however, no significant variability was detected in those observations. More recently, both V286 and V287 were included in the *Gaia* DR3 epoch photometry. We applied LS periodograms to the *G* band photometry, yielding most likely periods of 16.5 ± 0.2 d for V286 and 4.773 ± 0.011 d for V287. These uncertainties were estimated from the half-width half-maxima of the respective periodogram peaks. Although both peaks exceed the 0.1 per cent FAP threshold, visual inspection of the phase folded light curves does not reveal convincing singular periodicity in either star.

Our observations (shown in Fig. 4.11) achieved full phase coverage over several cycles for each star. In both cases, the average variability is well fit by a single two-harmonic sine series, indicating a dominant single period. However, both light curves exhibit amplitude modulation on timescales longer than the survey baseline. As described in Section 4.3.2, we performed MCMC fits using a two-harmonic sine model.

The fit results are summarized in Table 4.5. We extract a period for V286 that is consistent with the value from *Gaia*. However, the period we measured for V287 does not match the *Gaia* epoch photometry at all. Although we constrain the period and amplitude and have reliable estimates of the stars’ intrinsic brightnesses and colours, we are unable to robustly classify either variable based solely on photometric data. The variability characteristics do not immediately correspond to any known type in the General Catalogue of Variable Stars (Samus’ et al., 2017).

However, we may posit a hypothesis based on their positions on the CMD. According to the definition of Geller et al. (2017), these objects are *red stragglers* (RSs), which are thought to be related to the more numerous and less luminous population of sub-subgiant (SSG) stars, differing only in evolutionary state; both populations of stars defy standard models of single-star evolution. Geller et al. (2017) identified seven RSs as probable members of open or globular clusters. Of these seven, three are photometric or radial-velocity variables, and four are X-ray sources, consistent with

Table 4.5: Properties of the observed unclassified variables. The amplitudes given are in the instrumental bandpass of TWIST and represent the average peak-to-peak amplitude over the observed window. The uncertainties are therefore concerning the *average* amplitudes. The variations from one oscillation to the next are larger (see Fig. 4.11).

V	<i>Gaia</i> DR3 ID	Period (d)	Amplitude (mag)
286	1454782118391718528	16.3515 ± 0.0011	0.2463 ± 0.0005
287	1454783320982452992	22.865 ± 0.003	0.1351 ± 0.0003

the proportions seen among SSGs and with the picture that these objects are (occasionally active) binary systems, related to the RS Canum Venaticorum (RS CVn) type.

We were unable to identify any previously known X-ray sources associated with either V286 or V287, although existing X-ray surveys of M3 (e.g., Stacey et al. 2011; Zhao et al. 2019) have focused on its core, with fields of view that exclude these two variables. We also queried the *Chandra* Source Catalog Release 2.1 (Evans et al., 2024) within 1 arcmin of each star, but this returned no sources and only upper limits for both stars. Low levels of X-ray activity may be expected, however, given the relatively long periods of photometric variability. Additionally, we searched the catalogue of individual globular cluster star radial velocities from Baumgardt & Hilker (2018), similarly finding no data for either star.

In the absence of both radial velocity and X-ray data, we can make no definite determination of the variable types, but the RS CVn classification remains compelling. RS CVn stars have active chromospheres, which can cause large spots and quasiperiodic photometric variability with amplitudes similar to those seen in V286 and V287. With periods longer than two weeks, both stars would fall into the long-period class defined by Hall (1976). Rotation-modulated spot-induced variability could explain the lack of variability seen in V287 by *Swift* and the non-detection of the 22.9 d period by *Gaia*. Spectroscopic follow-up would provide an effective test of the RS CVn hypothesis, as such data could detect variable Ca II H and K emission lines indicative of an active chromosphere, and radial velocities from these measurements could also test for a binary period similar to the photometric variability period.

In any case, we suggest that, given their proximity on the colour–magnitude diagram, and the similarity of their periods and amplitudes, V286 and V287 are most likely of the same variable type.

4.4 Non-Periodic Variability

4.4.1 Microlensing

Globular clusters are, in theory, ideal locations to observe microlensing events against background stars (Paczynski, 1994). Since the potential lenses have common and constrained distances and proper motions, resolving the lens at all is unnecessary. Additionally, knowing the lens distance and proper motion drastically reduces the complexity of often degenerate microlensing systems. If the background source’s distance and proper motion can also be determined (or otherwise reasonably assumed), the problem of determining the lens mass reduces to measuring the Einstein crossing time (Gould, 2000), which is trivial using photometry.

We considered the possibility of observing microlensing events from low-mass, sub-stellar lenses

in our data by analytically predicting event rates under two scenarios – firstly, with a lens anywhere in the cluster and the source far in the background (background lensing), and secondly, with a lens in the halo of the cluster near the tidal radius and the source in the centre of the cluster (self-lensing).

Self-lensing

Due to dynamical mass segregation in globular clusters, lower-mass objects are expected to migrate toward the outer regions over time, while higher-mass stars sink towards the centre (Spitzer, 1969). Hence, we assume for our model that potential lenses are located in a shell at a distance much greater than the cluster’s stars. In this regime, the relative source-lens proper motion is dominated by the motion of the source in the cluster centre, so we do not consider the lens’s motion. We take a microlensing ‘event’ to be an alignment of a source and lens within the Einstein radius, θ_E , which corresponds to a source amplification of around 1.34 times. The event rate, Γ , is then the area swept out by twice the Einstein radius over lens-source proper motion, multiplied by the surface density of source stars. This approach is inspired by that of Kiroğlu et al. (2022). Formulaically,

$$\Gamma = \frac{v_d N_s N_l \sqrt{G} \overline{\sqrt{M_l}}}{\pi r^{3/2} c}, \quad (4.9)$$

where N_s is the number of sources, N_l is the number of lenses, r is the radius of the shell of lenses, and $\overline{\sqrt{M_l}}$ is the mean value of the square roots of the lens masses. v_d is the velocity dispersion of the cluster, which we use as a characteristic velocity of the sources in this approximation.

Using the central velocity dispersion (5.5 km/s) from Harris (2010), the limiting radius (103 pc), from Trager et al. (1995), the present-day mass of M3 (450,000 $M_\odot$) from Massari et al. (2016), and assuming an average stellar mass of 0.35 $M_\odot$, we predict an *inverse* event rate,

$$\Gamma_{M3}^{-1} \approx 52300 \frac{N_s}{N_l} \frac{\overline{\sqrt{M_l}}}{\overline{\sqrt{M_l}}} \left(\frac{r}{r_L} \right)^{3/2} \text{ yr}, \quad (4.10)$$

where r_L is the limiting radius. Hence, microlensing by sub-stellar lenses of the central stars would be extremely rare, on the order of once per tens of thousands of years, even if a large population of such lenses existed. Additionally, the sky-projected Einstein radius due to a Jupiter-mass lens at the limiting radius would be only $\sim 6.1 R_\odot$, meaning finite source effects would be present, reducing the maximum amplification and detectability.

Background Lensing

Similarly, we created an analytic model to predict the event rate in the scenario that a lens anywhere in the cluster amplifies a stellar source in the distant background. For simplicity, we assume the source is located infinitely far away and stationary, and that the relative source-lens proper motion is equal to the mean proper motion of the cluster.

In this case, the event rate is given by

$$\Gamma = 2N_l \sigma_s \mu \sqrt{\frac{4G}{c^2} \frac{1}{d} \overline{\sqrt{M_l}}}, \quad (4.11)$$

where σ_s is the angular area density of sources, μ is the magnitude of the cluster proper motion,

and d is the distance to the cluster. Note that the product of the square-rooted terms (the mean Einstein radius) should be interpreted as an angle in radians, thus σ_s must be given in sr^{-1} and μ in radians per unit time. Taking the stellar mass of the galactic halo beyond the cluster to be $3 \times 10^8 M_\odot$ – the value for 10–45 kpc from Bland-Hawthorn & Gerhard (2016) – and the mean stellar mass to again be $0.35 M_\odot$, we crudely estimate the source density behind the cluster to be $6.8 \times 10^7 \text{ sr}^{-1}$, assuming a purely radial halo distribution and ignoring the difference in line of sight of the cluster between the sun and the galactic centre. We take the distance to the cluster to be 10.2 kpc from Bhardwaj et al. (2020) and again use the cluster proper motion from Libralato et al. (2022). This results in an inverse event rate of

$$\Gamma_{\text{M3}}^{-1} \approx 3270 \frac{N_{\text{stars}}}{N_1} \frac{\sqrt{M_J}}{\sqrt{M_1}} \text{yr}, \quad (4.12)$$

where N_{stars} is the number of stars in the cluster and hence N_1/N_{stars} is the number lenses per star. This result indicates that such microlensing events are still too rare to hope to observe: once every few millennia for the scenario of one Jupiter-mass lens per star. This does not even consider that most halo stars are too dim to be detected except by the largest space telescopes. If we consider instead microlensing by the cluster stars themselves, the event rate would be around an order of magnitude larger, but still far too rare to consider.

Microlensing Event-Rate Discussion

The predicted rarity is perhaps not surprising. Past globular cluster microlensing searches have focused on the clusters M22 and 47 Tuc, which each have far richer stellar backgrounds than M3 – the galactic bulge and the Small Magellanic Cloud, respectively, and these searches have borne little.

47 Tuc was observed over a baseline of almost five years, with a mean cadence between 2 and 3 days, in field SMC140 of the third iteration of the Optical Gravitational Lensing Experiment (OGLE-III, Wyrzykowski et al. 2011). No microlensing events were detected.

Sahu et al. (2001) observed M22 with the *Hubble Space Telescope* over a baseline of 113 days at a typical cadence of 3 days and reported one time-resolved and six time-unresolved microlensing events. However, all were later disputed (Sahu et al., 2002; Anderson et al., 2003). Pietrukowicz et al. (2005) reported the probable detection of a microlensing event in M22 with the 1-metre Swope telescope at Las Campanas Observatory, which was later confirmed by Pietrukowicz et al. (2011), making it the first confirmed microlens in a globular cluster. To our knowledge, this remains the only confirmed example.

Kiroğlu et al. (2022) predicted microlensing event rates by M dwarf, white dwarf, neutron star, and black hole lenses for 47 Tuc and M22 using simulated motions of the cluster constituents, finding negligible event rates for self-lensing. Background lensing rates were dominated by white dwarfs and M dwarfs, and were at a rate of around a few per year. Given M3’s much poorer stellar background, we would not expect to observe any microlensing events over the course of our observations due to background lensing. Self-lensing rates would likely occur at similar rates to those of M22 and 47 Tuc, and thus would also not be expected to be found in our data.




Figure 4.12: Light curves of the three suspected M dwarf flares, approximately centred on their times of peak brightness. *Gaia* DR3 *G* magnitudes are approximated per the reduction in Section 4.2.2.

Search for transient brightening

Despite the predicted rarity of microlensing events, we searched for evidence of non-periodic transient brightenings in all the stars observed in the survey. In the reduced light curves, we checked for three consecutive points more than 5 times the median absolute deviation (in mag) brighter than the median brightness. Apart from obvious systematics, this identified three candidates with brightenings of 0.5-0.8 mag, lasting from around 7 to 30 minutes. Their *Gaia* DR3 IDs are 1454876878250136576, 1454801733503276672, and 1454799087803284608. For these three, we verified that the brightenings were visible in both the raw flux and the inferred magnitudes and that the centroids remained stable over the event durations. Their light curves are shown in Figure 4.12.

The three stars have $G_{BP} - G_{RP}$ colours – coupled with the conversion table of Pecaut & Mamajek (2013) – indicative of spectral types M3–M4.5 and absolute brightnesses consistent with the main sequence. Additionally, the variations were asymmetric, exhibiting a rapid rise and a more gradual decay. This leads us to conclude that the rises we saw were nothing more than standard M dwarf flares in foreground stars. Thus, as expected, we did not detect any microlensing events, despite

demonstrating sensitivity to short-duration transient brightenings.

4.5 Conclusion

We have carried out a 67-night photometric monitoring campaign of M3 with the newly commissioned TWIST observatory. The survey has yielded both technical validation of the system and astrophysical results that build on more than a century of variability studies in this cluster. Among the new findings are four SX Phoenicis stars, including one fundamental-mode pulsator that provides an independent distance consistent with published values. The detection of these short-period SXP variables particularly highlights the advantage conferred by TWIST's high cadence capability and the large data volume of this survey, despite the relatively small 50 cm aperture. Earlier searches (e.g. Kaluzny et al. 1998; Corwin & Carney 2001; Hartman et al. 2005; Benkő et al. 2006), which operated at cadences of order 5–20 minutes and with far fewer data points, were inherently less sensitive to the ~ 1 –3-hour pulsations characteristic of this class.

For the extensive RR Lyrae population, we measured updated periods and amplitudes for the majority of stars, and identified Blazhko-like modulation in over fifty cases. The long timescale of this variability motivated the need for stable photometry over timescales of hundreds of days, which was enabled by the development of a new correction method for scattered-light contamination in flat-field frames.

Notably, amplitudes were established for sixteen RR Lyrae stars that had no published values, improving the completeness of the literature catalogue. We also obtained the first period and amplitude measurements for the unclassified variables V286 and V287, and identified several flaring foreground M dwarfs. A search for microlensing events supported the prediction that the expected rate in M3 is negligible, providing an empirical benchmark for the detectability of these rare events with small-aperture instrumentation.

Although the analysis presented here uses white-light imaging, TWIST is equipped with a complement of *BVR* filters. Filtered observations were not employed for this programme because the associated reduction in photon flux would have limited our sensitivity to the faintest and lowest-amplitude variables detected in this study. TWIST operates as a general-purpose facility, and since these observations it has been used for a range of projects, including exoplanet transit monitoring, studies of white dwarf variability, and testing of new sCMOS detectors. The circumstances under which the filters will be used in future will depend on the requirements of those programmes, offering flexibility for a wide range of forthcoming projects.

Chapter 5

PLATO Timing Predictions for the *Kepler* Field

The European Space Agency is set to launch *PLATO*, the third medium-class mission of its Cosmic Vision programme, in early 2027. Using the transit method, *PLATO* is expected to detect thousands of exoplanets orbiting bright, nearby stars of spectral types F5–K7. Although the mission is primarily designed to enable mass measurements via radial velocities through its dedicated Ground-based Observation Program (GOP), its precise photometry and long observational baselines may also permit the detection of transit timing variations (TTVs), which can provide complementary dynamical constraints in multi-planet systems. One possible *PLATO* observing scenario involves a two-year-long observation of a Northern field that may partially or fully overlap with the original *Kepler* field, creating an opportunity to revisit known multi-planet systems with a photometric baseline exceeding 20 years. We simulate *PLATO* observations of 152 *Kepler* host stars containing at least one planet with previously detected TTVs, yielding a sample of 361 transiting planets. Our CCD-level simulations incorporate realistic stellar variability and employ both aperture and point spread function (PSF)-fitting photometry, accounting for each target’s real photometric contaminants. While the extended temporal baseline offers the potential for improved dynamical constraints in favourable cases, our simulations show that this potential is strongest for carefully selected systems, as *PLATO*’s smaller collecting area and larger pixel scale limit the achievable per-transit precision relative to *Kepler*. We identify a subset of systems most likely to benefit from complementary dynamical constraints through *PLATO* observations.

5.1 Introduction

Following in the footsteps of *Kepler* (Borucki, 2016), PLANetary Transits and Oscillations of stars (*PLATO*; Rauer et al. 2025) will use high-precision photometry to detect and measure the transits of thousands of exoplanets. A fraction of these planets will exhibit non-linear ephemerides due to gravitational interactions with other bodies in their systems (Agol et al., 2005; Holman & Murray, 2005). These deviations from a linear ephemeris, known as transit timing variations (TTVs), most commonly arise from planet–planet interactions, but can in principle also be induced by additional companions such as exomoons (Kipping, 2009). TTVs can therefore be used to detect non-transiting

companions and to place dynamical constraints on planetary systems (e.g. Lithwick et al. 2012; Hadden & Lithwick 2017).

According to the NASA Exoplanet Archive (Christiansen et al., 2025), as of February 2026, 379 of the 477 planets with detected TTVs were discovered by *Kepler* in its nominal mission. This follows from several aspects of the *Kepler* design, including its long-duration stare of a single field, high photometric precision, and short-cadence observations of many known planet-hosting stars. Many transits are typically required to detect measurable deviations from a linear ephemeris (Kane et al., 2019), and more still to sample a substantial portion of a TTV super-period – the long-timescale ephemeris modulation defined by Lithwick et al. (2012), whose duration depends on the proximity of a planetary pair to resonance. Sampling the super-period is important for constraining the properties of perturbing companions, and Agol & Fabrycky (2025) note that two well-sampled super-periods are preferable for a robust analysis. Although not all periodic TTV signals arise from near-resonant super-periods, Kane et al. (2019) showed that strongly periodic TTV signals from confirmed *Kepler* planets commonly span hundreds to more than 1000 days.

The four-year observing baseline of *Kepler*’s primary mission enabled the detection of TTVs on a range of timescales in hundreds of systems. In contrast, the Transiting Exoplanet Survey Satellite (*TESS*; Ricker et al. 2015), with its 27-day observing sectors, has had more limited success with this method. Only 60 planets discovered by *TESS* have been identified as showing TTVs.

The precision with which transit timings can be measured depends on both the photometric transit signal-to-noise ratio (SNR) and the observing cadence. Price & Rogers (2014) showed that the uncertainty in mid-transit times is several times larger for light curves with 30-minute integrations compared to those with instantaneous sampling, and that this uncertainty scales approximately inversely with total transit signal-to-noise ratio (SNR).

The possibility of detecting TTVs with *PLATO* is motivated by its expected multi-year stare durations. One current proposed observing strategy includes two two-year stares, one in each ecliptic hemisphere, although longer stares remain under consideration (Rauer et al., 2025). Kane et al. (2019) estimate that around 30 per cent of strong TTV signals from *Kepler* remain detectable with just one year of data, suggesting that the *PLATO* baseline will be sufficient to recover TTVs in favourable cases.

PLATO will produce 25-second cadence imagerettes for its ~7,500 bright P1 sample targets in each field (see Rauer et al. 2025 for an overview of the *PLATO* samples), while most of the fainter P5 sample will be observed at ten-minute cadence. Previously known exoplanet hosts may also be allocated faster cadence data products. The effect of cadence on mid-transit timing precision is therefore a relevant consideration. In addition, *PLATO*’s photometric precision will vary across its target list with the number of observing cameras and due to factors such as aperture contamination and intrinsic stellar properties.

In this study, we use PlatoSim (Jannsen et al., 2024) to simulate two years of time-series photometry of 152 *Kepler* host stars containing 361 transiting planets, where at least one planet per system has previously detected TTVs. In addition to the instrument and detector modelling provided by PlatoSim, the end-to-end simulations incorporate simplified prototype implementations of the onboard and on-ground data processing pipelines. These prototype pipeline components, developed at LIRA Paris Observatory, approximate the expected photometric extraction and processing steps,




Figure 5.1: Period-radius diagram of the 361 sample planets (coloured by host brightness) and the remaining 2,331 transiting exoplanets discovered by the *Kepler* satellite (shown in grey). The data were taken from the NASA Exoplanet Archive.

although they do not represent the final mission pipelines.

The goal of this work is to assess whether *PLATO* can recover TTV signals from these systems. The results identify which *Kepler* systems are likely to benefit from additional photometric coverage and provide a benchmark for *PLATO*’s expected performance in this context.

In Section 5.2, we overview the sample properties and simulation inputs. In Section 5.3, we discuss the noise characteristics of the extracted light curves. In Section 5.4, we detail the transit fitting process and introduce a timing constraint classification scheme. In Section 5.5, we discuss the timing performance across the sample and highlight systems likely to benefit from *PLATO* follow-up.

5.2 Simulations

5.2.1 Target selection

We first used the NASA Exoplanet Archive to identify all of the systems hosting at least one planet exhibiting transit timing variations and at least one planet discovered by the prime *Kepler* mission. This identified 178 potential hosts. For simplicity in simulation, we limited the sample to single-star hosts, reducing the number of targets to 165. Finally, we used the host stars’ temperatures and radii to exclude likely sub-giant hosts, as our stellar variability simulator cannot handle these stars. For this, we used the cutoff described in Nascimbeni et al. (2022), which is based on the work of Pecaute & Mamajek (2013). This cut reduced the number of hosts to a final sample of 152 stars, which, between them, host 361 known transiting exoplanets. We include all of these 361 planets in our sample regardless of whether they show transit timing variations. We show the sample planets’ periods and radii in the context of all *Kepler* discoveries in Fig. 5.1. The host brightnesses span the interval $9.77 \leq V \leq 16.69$, with a median brightness of $V = 14.67$. This is marginally brighter than the median brightness of all *Kepler* hosts at $V = 14.83$.

We note that *PLATO*’s primary P1 and P5 samples include only stars brighter than $V = 11$ and $V = 13$, respectively. Within our sample, only four stars are brighter than $V = 11$, and twenty are in the range $11 < V \leq 13$. Additionally, Rauer et al. (2025) lists the photometric range of *PLATO* to be $4 \leq V \leq 16$, while 11 of our targets are dimmer than $V = 16$. Hence, we naturally expect the majority of hosts in our sample to be challenging targets for *PLATO*.

With the sample chosen, we used a modified version of the `PICSIM` module in `PlatoSim` to generate target and contaminant catalogues for the targets. This method queries `Simbad` (Wenger et al., 2000) to identify the named target star and find its coordinates, and then *Gaia* DR3 (Gaia Collaboration et al., 2023) to find the target’s DR3 photometric properties as well as those of contaminating sources in its vicinity. The *Gaia* DR3 G , G_{BP} , and G_{RP} magnitudes are used to estimate the *PLATO* passband P magnitudes using Eq. 9 of Marchiori et al. (2019). For each target, we opted to include all contaminants brighter than 1 per cent of the target brightness ($\Delta\text{mag} < 5$) within 45 arcsec (3 pixels) of the target. This choice struck a balance between including the effect of contamination while avoiding an excessive increase in photometry extraction time, which is strongly dependent on the number of contaminants.

The prime *Kepler* field is encompassed in its entirety by the provisional northern *PLATO* field, LOPN1 (Nascimbeni et al. 2022, shown with *Kepler* field and the sample stars in Fig. 5.2), and spans the full range of the number of co-observing, co-pointing camera groups, from one to four. Assuming LOPN1 pointing, the number of host stars in our sample observed by each possible number of normal cameras (N-CAMs) is: 6 cams – 11, 12 cams – 85, 18 cams – 23, and 24 cams – 33. For all our simulations, we assume the LOPN1 pointing, with the understanding that the final pointing may change. We note that any change is likely to be small, given the constraints on the minimum ecliptic latitude and the need to be close enough to the galactic plane to cover the required number of sample targets. Thus, it is unlikely that any proposed change would result in missing the *Kepler* field entirely.

We chose to simulate all 152 systems to achieve a statistically powerful sample of TTV planet hosts, despite the considerable computational cost (discussed in Section 5.2.5). We note that the similarity between the sample host and planet properties and those of the overall *Kepler* population additionally permits meaningful inferences of the performance of *PLATO* in following up *Kepler* targets generally.

5.2.2 Stellar variability

We used the `VARSIM` module in `PlatoSim` to generate a realistic intrinsic stellar variability signal for each star. This module takes as inputs the stellar mass, radius, temperature, surface gravity, and $[\text{Fe}/\text{H}]$ metallicity and is only valid for main-sequence dwarfs. Where stellar parameters were not available on the NASA Exoplanet Archive, we referred to the TICv8.2 catalogue (Paegert et al., 2021). This approach retains the often greater accuracy of spectroscopically derived parameters where possible. Where metallicity data were not available, we assumed the solar value. With these parameters, `VARSIM` generates a realistic stellar variability signal due to granulation based on the method of Corsaro et al. (2013), pulsation from Kallinger et al. (2014), and the effect of appearing and disappearing star spots modulated by stellar rotation from Aigrain et al. (2015).

A notable caveat of this approach is that for any given host star, the simulated variability may or may not accurately reflect its true variability. One natural choice may be to attempt to derive the stellar




Figure 5.2: The locations of the sample targets within the LOPN1 *PLATO* and *Kepler* fields. The provisional *PLATO* LOPN1 field is shown in blue, with the darkness of the shaded regions indicating the number of co-pointing cameras, ranging from 6 in the lighter outer regions to 24 in the darker centre. The long-term pointing field of *Kepler* is shown in red. The positions of the 152 sample stars are shown as circles coloured by their apparent V-band brightness.

variability from the *Kepler* data; however, this approach would have its own issues. For example, the problem of separating intrinsic stellar variability from instrumental noise, and that most stars will be in a different part of their magnetic cycles by the time they are observed by *PLATO* compared to the point at which they were observed by *Kepler*. Additionally, rather than attempting to generate some modelled intrinsic variability for all the contaminating stars, we modelled them as constant sources. We expect PSF-fitting photometry to mitigate much of the impact of varying contaminants in real data, although targets with aperture photometry will likely be more severely affected.

This method is a practical approach for simulating large numbers of stars, which, in addition to enabling statistical inferences, also enables meaningful predictions for individual systems.

5.2.3 Transit injection

We again used *VARSIM* and the NASA Exoplanet Archive to simulate the noise-free transit light curves for each star. The transit modeller within *VARSIM* takes as inputs the ephemeris time, orbital period, inclination, eccentricity, argument of periastron, planet mass, and planet radius in addition to the stellar parameters listed in Section 5.2.2. It implements *BATMAN* (Kreidberg, 2015) for primary transits, *SPIDERMAN* (Louden & Kreidberg, 2018) for secondary occultations, and *LDTK* (Parviainen & Aigrain, 2015) for generating *PLATO* bandpass-specific limb darkening coefficients. Doppler beaming and ellipsoidal variation are implemented in *PYASTRONOMY* (Czesla et al., 2019), but since the masses of *Kepler* planets are usually unknown, we assumed zero mass for all planets, thereby removing the impact of these (already minor) effects.

We adopted planetary periods, inclinations, eccentricities, and radii from the default parameter

sets provided by the NASA Exoplanet Archive. In cases where the argument of periastron was not listed – typically due to the difficulty of measuring this parameter at low eccentricity – we assumed a constant value of 90° . When eccentricities or inclinations were absent, values of zero and 90° were assumed, respectively. For Kepler-103 c, the default stellar and planetary parameters listed in the archive, drawn from Bonomo et al. (2023), yield a transit impact parameter greater than 3, implying a non-transiting geometry. The inclination reported by Bonomo et al. (2023), $87.704^{+0.12}_{-0.055}^\circ$, appears to be a transcription error from the source values reported by Dubber et al. (2019), differing by exactly 2° . This inconsistency was identified only after the simulations and subsequent analysis had been completed, and Kepler-103 c was therefore excluded from the final sample, reducing the total to 360 planets.

Since it is not currently known exactly when *PLATO* will begin observations of its Northern field (if at all), and a large fraction of our sample planets don’t have ephemeris times listed on the NASA Exoplanet Archive, we uniformly sampled random ephemeris times between zero and each planet’s orbital period.

We combined the (potentially multiple) noiseless transit light curves for each star with their stellar variability light curves by multiplying them together to produce two years of intrinsic photometric variability for each system.

5.2.4 Payload simulation

The final ingredient needed to generate CCD imagettes is the *PAYLOAD* module, which manages the telescope systematics. This module injects alignment errors between the cameras and the optical bench, pointing repeatability errors for each observing quarter, thermoelastic distortion of the optics, a realistic distribution of observational downtime, and gain-thermal transients that increase the CCD temperature after quarterly rolls. We used all of the default parameters of this module (see Janssen et al. 2024) to generate eight quarters of payload systematics.

5.2.5 Photometry extraction and simulation runs

We used the built-in *PYTHON* wrapper of PlatoSim, *PLATONIUM*, to run the simulations. This combines the outputs of the *PICSIM*, *VARSIM*, and *PAYLOAD* modules, runs the simulation for one star over one quarter (of a year) as observed by one camera, and then applies the chosen photometry pipeline.

We do not yet know precisely which data product will be produced for these targets; therefore, we adopted a mixed approach to photometry extraction. Firstly, we simulated the use of the 25-second cadence imagettes for on-ground PSF-fitting photometry via the L1 photometry pipeline. This is the treatment that the P1, P2, P4 and ~ 3.7 per cent of the P5 sample targets will receive. Secondly, on the same imagettes, we simulated onboard aperture photometry using the method of Marchiori et al. (2019), which constructs photometric precision-maximizing apertures composed of whole pixels that are updated periodically throughout the quarter. This method is analogous to the expected treatment of the majority of the P5 sample stars for which bandwidth limitations preclude imagette downlinking.

In both cases, the treatment of the PSF is important. For stars subject to PSF fitting, point spread functions with subpixel resolution are obtained via a microscanning approach (see Samadi et al. 2019 for details). This method involves acquiring a series of imagettes of the target that are

shifted by subpixel displacements following an Archimedean spiral, and combining these data using a regularized least-squares approach with a positivity constraint on the solution. In our simulations, PLATONIUM orchestrates the spiral observations within PlatoSim at the beginning of each observing quarter, while the L1 proto-pipeline performs the PSF inversion and fitting.

For stars for which only aperture photometry is performed, knowledge of the PSF is still required – in order to define the photometric apertures. In the nominal processing, PSFs are reconstructed for reference stars across the field of view and interpolated to estimate the PSF at a given focal-plane position. In our simulations, however, we do not perform this reconstruction step; instead, we use the theoretical PSFs provided by PlatoSim.

We initially ran 29,184 simulations – 152 stars observed by 24 cameras over 8 quarters. We ran simulations from mission quarters 9 to 16 in PLATONIUM to account for the expected instrumental degradation between the beginning of the mission and when the Northern field is observed. Since we did not know in advance which or how many cameras the targets would be visible to, we ran the simulations for all 24 cameras. Simulations for which the target was outside the field of view terminated early. These simulations took approximately seven weeks spread across three systems, each equipped with two 32-thread Intel Xeon Silver 4514Y processors, for a total of around 200,000 CPU hours. The execution time was dominated by the computational cost of the L1 PSF-fitting photometry proto-pipeline.

Additionally, we ran an identical set of simulations using only aperture photometry, with all contaminants removed, to quantify the mean photometric contamination level per simulation and per mask update. These simulations took around three days using one of the above-described machines for a total of around 4,000 CPU hours. For each mask update, we subtract the mean contamination level from the contaminant-present aperture photometry. This produces light curves corrected for contamination in the mean signal, while retaining the noise impact of photometric contaminants. In the real data, contamination will be addressed via the spacecraft jitter correction, which re-scales the measured flux to the level expected with only the target. These corrections can be validated by measuring the dilution of exoplanet transits relative to the catalogued transit depths. We did not utilize the jitter or long-term drift corrections in the present work.

The PSF-fitting photometry failed for one eighth of the targets – 19 stars hosting 50 planets. These 19 stars all have $V \geq 14.79$, and the mean and standard deviation of the failed PSF-fitting host brightnesses in V are 15.79 and 0.53 mag, respectively. Failed extraction in the current pipeline implementation appears to be associated with a combination of low target brightness and high levels of photometric contamination. We note, however, that PSF-fitting photometry is generally expected to outperform aperture photometry in crowded fields and for faint sources. This has been demonstrated for *Kepler* – which used similarly undersampled PSFs – by Libralato et al. (2016a,b). We therefore expect comparable improvements to be achievable. The PSF-fitting method used in the present work was primarily developed for and tested against stars with $V \leq 11$, which may contribute to the reduced performance observed for dimmer targets. However, given that the faint ($V \leq 16$) P4 M-dwarf sample stars will receive the same 25-second imagette data product as the bright P1/P2 samples, we expect the final implementation of the PSF-fitting routine to more effectively handle dim, crowded targets. Despite the limitations of the current PSF-fitting photometry implementation, aperture photometry was successfully extracted for the entire sample.

Each of the four CCDs in a *PLATO* camera is read out with a 6.25-second offset, allowing an effective cadence finer than 25 seconds. In principle, this enables a combined light curve with 6.25-second resolution when data from all four camera groups are used. However, only transits with very high signal-to-noise ratios would significantly benefit from such high temporal resolution, which was not the case for the overwhelming majority of our targets. Additionally, each camera group produces an independent light curve that would require separate detrending before combination, with the risk of introducing additional noise. Thus, we chose to average the light curves from the four groups to produce a final light curve with a uniform 25-second cadence, which remains superior to the shortest cadence of the *Kepler* data (58.85 seconds).

5.2.6 Detrending

We chose to use the biweight detrending method implemented in *WOTAN* (Hippke et al., 2019) with a window length of one day based on its good performance in Canocchi et al. (2023). Since we knew from the simulation inputs the transit times and durations, we masked out all of the transits before detrending. This work focuses on the measurement and detectability of known transits, making this approach readily applicable to real data. For each star, we applied this method to both the aperture photometry and, where available, the PSF-fitting photometry 25-second cadence light curves.

5.3 Photometric performance

We measured the standard deviations of the 3σ -clipped out-of-transit portions of the 25-second cadence detrended light curves derived from both photometry methods. We then binned the data to a 1-hour cadence by taking the mean average of every consecutive $3600/25 = 144$ points, weighted by their inverse variances and additionally measured the scatter in the binned light curves.

5.3.1 Comparison of PSF and aperture photometry

Substantial differences were observed between the noise-to-signal ratios (NSRs) obtained from PSF-fitting and aperture photometry. Across our sample, the ratio $\sigma_{\text{PSF}}/\sigma_{\text{aper}}$ spans more than two orders of magnitude, ranging from 0.23 (PSF lower noise) to 76 (aperture lower noise), as shown in panel (a) of Fig. 5.3.

The relative performance of the two methods varies from star to star and reflects a combination of target brightness and local crowding (see panels (b) and (c) of Fig. 5.3 showing the variation of the number of contaminants and the relative amounts of contaminating flux with host brightness). Because the L1 PSF-fitting pipeline is under active development, we do not attempt to interpret these differences in detail, but we note that relative performance is most strongly correlated with the number of contaminants, rather than the contamination ratio. For each target, we adopt the photometry method that yields the lower 1-hour noise. This ensures that our timing forecasts reflect the best-performing extraction method available for each star within the current simulation framework. In total, 28 stars achieve better precision with PSF-fitting photometry and the remaining 124 stars achieve better precision with aperture photometry.




Figure 5.3: Plots of, as a function of host brightness: **(a)** The ratio of the PSF-photometry-derived to aperture-photometry-derived 1-hour photometric noise. **(b)** The number of contaminating stars within 45 arcsec brighter than 5 magnitudes dimmer than the targets. **(c)** The mean ratio of contaminant to target flux within the photometric apertures. **(d)** The mean size in pixels of the Marchiori et al. (2019) derived photometric apertures. **(e)** The noise redness factor $\beta_{1\text{hr}}$ as defined in Eq. 5.1 for each photometry method. For each star, the value associated with the photometry method that yields the lower 1-hour noise is shown with greater opacity.




Figure 5.4: Noise-magnitude plot for the sample of 152 stars. The noise is the measured 3σ -clipped scatter in the detrended and then 1-hour binned out-of-transit light curves. Of the two noise-to-signal ratios calculated from the light curves of the two photometry extraction methods, the lowest noise result is shown for each star. The stars observed by 6, 12, 18, and 24 cameras are denoted by triangles, circles, squares, and diamonds, respectively. The scatter in the real *Kepler* data is calculated in the same way, shown as pentagons with lines connecting the real *Kepler* and simulated *PLATO* scatters for the same stars. The modelled (Cabrera et al., 2026) *PLATO* performances for 6, 12, 18, and 24 cameras are shown as solid lines. The 50 ppm threshold for Earth-analogue detection is indicated by a horizontal dashed line.

5.3.2 Comparison with modelled and *Kepler* noise

In Fig. 5.4, we show the 1-hour noise-to-signal ratio (NSR) as a function of host brightness and camera number, along with the beginning-of-life (BOL) condition ‘Quick noise model’ of Cabrera et al. (2026)¹. Additionally, the ratio of observed to modelled noise is shown against contamination ratio in the upper panel of Fig. 5.5.

Across our sample, we observe reasonable agreement with the model predictions across the range of host brightnesses. The median ratio of observed-to-predicted noise is 1.061, indicating that the model marginally underpredicts noise levels on average. This noise excess is not generally due to stellar variability. By applying the same detrending approach as in Section 5.2.6 to the raw, injected stellar signals, we find that the levels of intrinsic stellar variability on one-hour timescales fall in the interval $6.2 \text{ ppm} \leq \sigma_{\star, 1\text{hr}} \leq 32.4 \text{ ppm}$. Subtracting in quadrature the stellar noise from the total one-hour noise for each star reduces the median noise ratio to 1.059. Thus, the stellar component makes a negligible contribution to the noise of all but the brightest stars in our sample. As shown by Cabrera et al. (2026), the noise properties of *PLATO* are expected to degrade between its BOL and end-of-life (EOL) states. This evolution is modelled in PlatoSim and may explain the remaining

¹The noise model of Cabrera et al. (2026) is parameterized in *PLATO* P magnitudes. To maintain consistency with the use of V magnitudes throughout this work, we adopt the V magnitude to electron-flux conversion given in Appendix A, Eq. 1 of Rauer et al. (2025).




Figure 5.5: **Top:** Ratio of the measured 1-hour noise-to-signal ratios to those predicted by the model of Cabrera et al. (2026), shown as a function of the contamination ratio. The stars observed by 6, 12, 18, and 24 cameras are denoted by triangles, circles, squares, and diamonds, respectively. **Bottom:** Noise redness, β_{1hr} , as defined in Eq. 5.1, shown as a function of the contamination ratio (the mean ratio of contaminant to target flux in the photometric apertures).




Figure 5.6: Distribution of the relative noise performance of *PLATO* and *Kepler* for the sample stars by number of co-pointing cameras. Vertical dashed lines indicate the median noise ratio for each number of cameras (numerically shown in the legend).

noise excess.

Two notable outliers are the dimmest stars in *V*, which exhibit substantially lower noise than predicted by the model. These systems, Kepler-1624 and Kepler-267, also have the reddest $V - P$ colours in the sample, with values of 1.40 and 1.66 mag respectively, compared to the sample median of 0.44 mag. Their large $V - P$ colours imply that they are significantly brighter in the *PLATO* bandpass than their *V* magnitudes would suggest, naturally leading to lower photometric noise.

Conversely, a notable fraction of stars show significantly larger noise levels than predicted. This effect appears to be correlated with the flux contamination ratios (shown in the upper panel of Fig. 5.5) with a turn-on at ratios $\gtrsim 0.5$.

We note that the quick noise model of Cabrera et al. (2026) assumes a fixed photometric aperture size. However, as shown in panel (d) of Fig. 5.3, the photometric-precision-optimizing aperture selection method of Marchiori et al. (2019) naturally selects smaller apertures for dimmer stars. One might therefore expect the model to overestimate the noise levels of faint targets, since smaller apertures include smaller contributions from sky background and read noise. In practice, however, this effect is not clearly observed, likely because any reduction in read noise and sky-background noise is offset by lower integrated flux and the generally higher contamination levels associated with fainter stars.

Fig. 5.6 shows the distribution of the relative noise performance of *PLATO* and *Kepler* for the sample of host stars. The *Kepler* noise levels were derived in largely the same manner as for the *PLATO* light curves: by detrending, where possible, the entire available short-cadence data and measuring the 3σ -clipped 1-hour scatter. Long-cadence data were used when short-cadence data were unavailable. The median relative noise performance scales as expected with the inverse square root of the number of cameras, ranging from 2.68 times the *Kepler* noise for 24 cameras to 5.14 times the noise for six cameras.

5.3.3 Correlated noise

Under the assumption of white noise, binning by a factor n , or mean-averaging every consecutive n data points, lowers the root-mean-square (RMS) scatter by a factor $\sqrt{n}$. As defined by Winn et al. (2007), we take as a measure of correlation in the noise the ratio of the measured noise in the 1-hour binned out-of-transit light curve to the level expected by reducing the 25-second noise by a factor $\sqrt{n} = 12$. Formulaically,

$$\beta_{1\text{hr}} = \frac{12\sigma_{1\text{hr}}}{\sigma_{25\text{s}}}, \quad (5.1)$$

where a value $\beta_{1\text{hr}} = 1$ is consistent with the expectation for white-noise. A plot of $\beta_{1\text{hr}}$ against host brightness for each star and for each photometry method is shown in panel (e) of Fig. 5.3.

The median $\beta_{1\text{hr}}$ across the sample is 1.063 and the 90th percentile value is 1.297, suggesting that most stars are dominated by white noise up to 1-hr timescales, which implies that the out-of-transit variability is slow enough to be captured by the 1-day window detrending.

There are some significant outliers (nine stars with $\beta_{1\text{hr}} > 1.5$). We find that the best predictor of non-white noise (a large $\beta_{1\text{hr}}$) is simply the flux contamination ratio. This is illustrated in the lower panel of Fig. 5.5. Similar to the noise excess, significant red noise appears to occur at contamination ratios $\gtrsim 0.5$, although substantial variation in $\beta_{1\text{hr}}$ remains even at large contamination ratios.

Additionally, we investigated the impact of stellar variability on the measured noise redness. Using the detrended, injected intrinsic variability described in Section 5.3.2, we calculated a *stellar-only* $\beta_{1\text{hr}}$ value for each star. The resulting 16th, 50th, and 84th percentiles over the stellar sample are [3.30, 3.64, 4.39], indicating that the intrinsic stellar variability is strongly red. However, for this redness to measurably affect the extracted photometry, the stellar contribution must also constitute a significant fraction of the total noise budget.

To estimate the contribution from the instrument alone, we derived an *instrumental-only* $\beta_{1\text{hr}}$ for each star by subtracting the stellar noise in quadrature from the total noise at both the 25-second and 1-hour cadences. We find that the instrumental-only $\beta_{1\text{hr}}$ is only significantly closer to unity than the total $\beta_{1\text{hr}}$ – which we define as $\beta_{1\text{hr}} - 1$ decreasing by at least a factor of two – for four stars, all brighter than $V = 12$. The median instrumental-only redness is 1.059, only marginally smaller than the median value for the total noise. This suggests that, although stellar variability is intrinsically highly red, it contributes only weakly to the overall red noise budget for most stars in the sample. Consequently, the observed redness of the extracted photometry is dominated primarily by instrumental systematics except for the brightest targets, where the stellar contribution becomes appreciable.

5.4 Transit Fitting

In the following section, we describe the procedure used to infer mid-transit times from simulated *PLATO* light curves. Our aim is not to detect transits independently, but instead to quantify the precision with which individual transit epochs can be measured given strong prior knowledge of the systems from *Kepler* or from combining transits from *PLATO*. Accordingly, we treat the transit shape and ephemeris as known to a good approximation, and focus on the conditional inference of the mid-transit time for each observed transit.

5.4.1 Light curve preparation

For each planet, we prepared up to ten individual transit light curves for fitting, balancing statistical completeness against computational cost. Using the predicted mid-transit times, we extracted segments from the detrended 25 s cadence light curves spanning ± 1.5 transit durations around each expected transit centre, for all transits occurring within the two-year observing window. This window size provides a sufficient out-of-transit baseline for normalization while avoiding the inclusion of unnecessary stellar variability. For each star, we adopted the photometric extraction method that yielded the lowest 1-hour noise metric.

Within each extracted window, we masked any transits of other known planets in the same system to avoid contamination. We then constructed a vetted sample of transit windows by requiring at least 95 per cent temporal coverage of the transit itself and at least 50 per cent temporal coverage of both the pre- and post-transit baseline regions. From this vetted set, we selected up to ten transits per planet, approximately uniformly distributed over the observing baseline. For planets with fewer than ten suitable transits, all vetted transits were retained. Finally, we applied a 5σ clip to the injected-model-normalized flux to correct for the effect of cosmic rays, which is particularly impactful on aperture photometry. We note that this is a simplified, ad hoc approach to outlier-detection that does not reflect the most up-to-date methods in the photometry pipelines.

5.4.2 Transit model and likelihood

For each extracted transit window, we model the observed flux using a fixed-shape limb-darkened transit light curve computed with the `BATMAN` package, adopting a quadratic limb-darkening law. All orbital and geometric parameters are held fixed at their injected values, including the orbital period, scaled semi-major axis, scaled planet radius, inclination, eccentricity, argument of periastron, and limb-darkening coefficients. This approach isolates the inference of the mid-transit time while avoiding degeneracies with poorly constrained shape parameters.

The free parameters in the fit are the mid-transit time t_0 , a multiplicative scaling factor applied to the transit depth d , a multiplicative flux baseline B , and an additive jitter term s that accounts for potential underestimation of the nominal photometric uncertainties. Allowing for a depth scaling accommodates small mismatches between the injected and observed transit depths due to residual uncorrected contamination.

Assuming Gaussian-distributed photometric errors, the log-likelihood ($\ln \mathcal{L}$) for a set of flux measurements f_i at times t_i is given by

$$\ln \mathcal{L} = -\frac{1}{2} \sum_i \left[\frac{(f_i - m_i)^2}{\sigma_i^2 + s^2} + \ln \left(2\pi(\sigma_i^2 + s^2) \right) \right], \quad (5.2)$$

where $m_i \equiv m(t_i; t_0, d, B)$ is the model flux and σ_i are the nominal per-point uncertainties. In the case of the PSF-fitting photometry, the per-point uncertainties are formally estimated by – and a direct output of – the L1 proto-pipeline. Conversely, uncertainties were estimated for the aperture photometry derived using the method of Marchiori et al. (2019) by calculating, for each mask update, the standard deviation of the detrended light curves, since formal uncertainties were not provided. The additional variance term s^2 is fitted simultaneously to ensure that timing uncertainties are not

artificially underestimated in the presence of imperfect noise estimation.

Priors are imposed as follows. The mid-transit time t_0 is assigned a uniform prior within a window of $\pm$ duration around the predicted ephemeris. This choice reflects reasonable prior knowledge obtained either through the propagation of *Kepler* ephemerides or the stacking of *PLATO* transits while allowing the posterior to be dominated by the likelihood function. The jitter parameter is sampled uniformly in $\ln s$ over the range $[-20, 1]$. The depth scaling factor d follows a Gaussian prior centred close to unity, with an exact mean and width calibrated empirically from high signal-to-noise simulated transits and separately for the PSF-fitting and aperture photometry. This strong prior on the depth reflects confidence in both the knowledge of the transit shape and the accuracy of the contamination correction, and is necessary to prevent the sampler from focusing on high-likelihood but physically implausible solutions. The baseline parameter B is assigned a uniform prior over the interval $[0.9, 1.1]$.

5.4.3 Parameter estimation and posterior sampling

Posterior distributions for the model parameters were explored using Markov Chain Monte Carlo (MCMC) sampling. For each transit window, we employed the ensemble sampler implemented in the *emcee* package (Foreman-Mackey et al., 2013). The sampler was initialized with 32 walkers distributed around an initial parameter estimate corresponding to the injected mid-transit time, unit depth scaling, unit baseline, and a jitter term set by the median photometric uncertainty in the window.

Each chain was evolved for 10,000 steps, with the first 1,000 steps discarded as burn-in. Convergence was assessed by verifying that the autocorrelation lengths of all parameters were at most 1/50 the retained chain length (satisfied by ~ 99 per cent of fits).

For each parameter, posterior summaries are reported using the median and central credible intervals. In particular, the uncertainty on the mid-transit time is quantified using half the interval between the 16th and 84th percentiles of the marginal posterior distribution. We retain the full posterior samples for subsequent analysis, including the assessment of posterior structure and model support described in the following sections.

5.4.4 Classification of timing constraints

Not all transit fits yield mid-transit time estimates that are equally informative. In particular, low signal-to-noise transits may produce posterior distributions that appear well behaved despite being dominated by the imposed priors or by chance alignment with noise. To distinguish genuinely informative timing constraints from such cases, we classify each transit fit based on the structure of its posterior distribution and the degree to which the transit model is preferred to a flat, no-transit model.

Our classification scheme proceeds hierarchically, and each transit is assigned to one of four mutually exclusive classes, as described below.

Prior-limited

Fits are classified as *prior-limited* when the posterior distribution of the mid-transit time t_0 is strongly influenced by the bounds of the uniform timing prior. Specifically, if the posterior median is consistent

with the lower or upper prior boundary to within three times the lower or upper timing uncertainty, respectively, the timing constraint is considered prior-dominated and the transit is assigned to this class. This criterion captures cases where the apparent constraint arises primarily from the imposed timing window rather than from information in the data and is often associated with a mid-transit timing error much smaller than the calculated timing uncertainty.

Low significance.

For fits that are not prior-limited, we assess whether the data provide meaningful support for a transit-shaped signal within the transit window. This is quantified using the difference in maximum log-likelihood, $\Delta \ln \mathcal{L}$, between the transit model and a null (constant-flux) model. In the null model, we fix the jitter term to the best-estimate value from the transit model fit. This prevents high-likelihood null model fits from being achieved through unrealistic inflation of per-point uncertainties, which can consume the significance of the transit model relative to the null model. Transits with $\Delta \ln \mathcal{L} < 5$ are classified as *low significance*, indicating that the transit model is not strongly preferred over noise within the fitted window.

Asymmetric high significance and symmetric high significance.

Transits with $\Delta \ln \mathcal{L} \geq 5$ are considered to have significant support for a transit-shaped signal. Among these, we further distinguish cases based on the symmetry of the marginal posterior distribution of t_0 . To quantify this, we define an asymmetry metric based on posterior quantiles,

$$A = \frac{q_{84} + q_{16} - 2q_{50}}{q_{84} - q_{16}}, \quad (5.3)$$

where q_{16} , q_{50} , and q_{84} denote the 16th, 50th, and 84th percentiles of the posterior, respectively. This quantity provides a measure of skewness in the interval $[-1, 1]$.

Transits with $|A| > 0.3$ are classified as *asymmetric high significance*, typically corresponding to cases where the timing constraint is strongly one-sided or multi-modal. Transits with $|A| \leq 0.3$ are classified as *symmetric high significance* and represent the most robust per-transit timing measurements, exhibiting approximately unimodal and symmetric posterior distributions for t_0 .

This classification forms the basis for the population-level analysis of transit timing performance presented in the following section.

5.5 Transit timing performance and classification across SNR

In the following section, we refer to the transit signal-to-noise ratio (SNR), which we define in the canonical, albeit simplistic (Kipping, 2023), fashion:

$$\text{SNR} = \frac{d_{\text{transit}} \sqrt{t_{\text{duration}} / 1\text{hr}}}{\sigma_{1\text{hr}}}, \quad (5.4)$$

where d_{transit} is the maximum transit depth, t_{duration} is the total (first to fourth contact) transit duration, and $\sigma_{1\text{hr}}$ is the photometric noise-to-signal ratio in one hour.

5.5.1 Transit SNRs of the sample planets

The distribution of expected SNRs across the planet sample is shown in panel (a) of Fig. 5.7. In the *Kepler* mission’s Transiting Planet Search, folded transit features with detection statistics greater than 7.1σ were designated ‘Threshold Crossing Events’ (Jenkins et al., 2010). Although the simple transit SNR does not exactly map to their detection statistic, we use $\text{SNR} \geq 7.1$ as a heuristic benchmark for whether a transit is independently detectable.

Of the 360 sample planets, which are broadly representative of *Kepler* planets overall, only 66 have an expected individual $\text{SNR} \geq 7.1$. When combining all the expected transits within *PLATO*’s expected two-year stare, and assuming white-noise scaling, the number of $\text{SNR} \geq 7.1$ planets rises to 267. We emphasize again that the goal of this work is not to determine whether *PLATO* can detect these planets independent of prior knowledge from *Kepler*, but instead to assess the achievable timing precision for known systems.

5.5.2 Timing performance across SNR

Panel (b) of Fig. 5.7 illustrates how the achievable timing precision for individual transits depends on SNR. As expected, higher SNR transits tend to yield tighter timing constraints. However, at fixed SNR there remains a substantial spread in the inferred timing precision, even for transits of the same planet, which reflects that the quality of an individual timing is stochastic in nature. A best-fit log-log regression to members of the two high-significance classes yields the relation along with the 1σ residual scatter:

$$\log_{10} \left(\frac{\sigma_{t_0}}{t_{\text{duration}}} \right) = -1.1122 \log_{10} (\text{SNR}) - 0.3962 \pm 0.1578. \quad (5.5)$$

We note that using the analytic result of Carter et al. (2008) or linear transformations thereof to predict timing precisions from SNR, ingress time, and transit duration yields a lower proportion of explained variance in the measured uncertainties than the above relation. This is likely a result of idealized noise assumptions that are not realized in the present SNR regime. Hence, we prefer our linear fit. In total, of the 3,399 transit fits, 1,716 fall into the prior-limited class, 594 into the low-significance class, 142 into the asymmetric significant class, and 947 into the symmetric high-significance class.

Panel (c) of Fig. 5.7 shows the fraction of transit fits in each classification as a function of SNR. At low SNR, most transits are either prior-limited or fail to achieve significant support over a flat model, whereas at higher SNRs the population becomes increasingly dominated by high-significance solutions. The dashed vertical lines indicate characteristic SNR values at which at least 50 per cent of transits reach a given minimum classification. Specifically, half of all transits achieve at least a low-significance but not prior limited timing constraint at $\text{SNR} \approx 2.1$ (50.83 per cent of sample planets have at least this SNR), half yield at least asymmetric high-significance solutions by $\text{SNR} \approx 3.8$ (33.61 per cent of planets), and symmetric high-significance timing dominates beyond $\text{SNR} \approx 4.3$ (28.61 per cent of planets). These thresholds indicate that, given realistic prior knowledge from *Kepler*, transits can be timed with significantly less significance than is required for independent detection.




Figure 5.7: **(a)** Histogram distribution of the expected transit signal-to-noise ratios within the planet sample, both from one single transit and also when combining all the transits expected in a two-year window. **(b)** The ratio of the mean timing uncertainty (half the 16th-84th percentile range) to the transit duration for each individual transit fit, coloured by fit classification as a function of transit SNR. A linear (log-log) fit to the two high-significance classes is shown as a solid line. **(c)** The (non-cumulative) fraction of transit fits falling into each class as a function of transit SNR. The SNR at which 50 per cent of transit fits achieve at least a given class are shown as vertical dashed lines. **(d)** The ratio of the timing precisions arising from 600- and 25-second sampled and 50- and 25-second sampled data.

5.5.3 Validation of timing measurements

For each class, we investigated the relationship between the injected and extracted best-fit transit times to search for evidence of bias or underestimation of uncertainties. For each transit fit, we calculated the ratio of the difference between the injected and extracted times to the relevant uncertainty, resulting in a Z-score for each transit mid-time. Under the assumption of Gaussian posteriors, a distribution, $Z_{t_0} \sim \mathcal{N}(0, 1)$ is expected.

Across all timing classes, the mean values of Z_{t_0} were found to be consistent with zero to 2σ , indicating no significant bias in the recovered transit times. The dispersion of Z_{t_0} depends systematically on classification: prior-limited fits exhibit sub-unity dispersion ($\sigma_{Z_{t_0}} \simeq 0.7$), reflecting conservatively large uncertainties imposed by the prior boundaries, while low-significance fits are consistent with unit variance. High-significance fits show a marginally inflated dispersion ($\sigma_{Z_{t_0}} \simeq 1.1$), attributable to residual non-Gaussianity and posterior asymmetry. We conclude that for all but the prior-limited fits, the extracted uncertainties reasonably reflect the true timing uncertainties.

5.5.4 Performance variation with sampling cadence

It has been shown, e.g. by Price & Rogers (2014), that finite light curve sampling leads to greater timing uncertainty than instantaneous sampling. In the preceding analysis, we analysed *PLATO*'s native 25-second cadence output. Per Rauer et al. (2025), the P5 $V \leq 13$ sample stars will predominantly receive 10-minute cadence data product, with a subset receiving 50-second cadence data, and an even smaller subset receiving 25-second data. To investigate the impact of realistic lower-frequency sampling, we repeated the fitting process of Section 5.4 on the same data, but mean-average binned to 50-second and 10-minute cadence. We supersampled the *BATMAN* fits to the original underlying 25-second cadence to account for smearing from finite integration times and to avoid introducing timing bias.

Panel (d) of Fig. 5.7 shows the ratio of retrieved precisions arising from data of each sampling cadence. Across the range of SNRs, we find little difference between the uncertainties in the 25- and 50-second cadence data. The 16th, 50th, and 84th percentiles of this ratio are [0.980, 1.000, 1.021]. Similarly, the 10-minute cadence data achieves near-identical performance to the 25-second cadence data on average, although the uncertainty ratio shows more scatter. The corresponding 16th, 50th, and 84th percentiles are [0.947, 1.013, 1.104]. We repeated the analysis of Section 5.5.3 on the slower cadence data to search for evidence of additional uncertainty underestimation and, in both cases, found almost identical distributions of Z_{t_0} to the 25-second cadence analysis. Furthermore, the number of fits falling into each classification does not deviate in either case by more than 1 per cent of the total number of fits compared to the 25-second cadence analysis. Thus, we conclude that neither 50-second nor 10-minute cadence data seriously degrades performance relative to native cadence in the studied SNR regime.

5.5.5 Transit timing yield for known *Kepler* systems

A majority of the planets' transits in our sample cannot be timed individually at a high level of significance. Under the $\text{SNR} \geq 3.8$ high-significance criterion adopted in Section 5.5.2, 121 of the 360 sample planets exceed the threshold for well-constrained single-transit timing measurements.

This does not imply that the remaining systems are uninformative. Leleu et al. (2021) developed a deep-learning method to detect and measure TTVs from river diagrams that remains effective for planets with low per-transit SNR, provided the TTV amplitude is sufficiently large. They demonstrate their approach on Kepler-1705 b and c, which each have a transit SNR ≈ 1 . In our sample, 272 of the 360 planets have an individual SNR ≥ 1 .

Additionally, in multi-planet systems exhibiting TTVs, even modest additional timing constraints on one planet can improve the dynamical characterization of the entire system through coupled orbital interactions. Of the 360 planets, 189 reside in systems where at least one planet has an expected transit SNR ≥ 3.8 , and 332 reside in systems where at least one planet has an expected SNR ≥ 1 . The degree to which timing information can be recovered from lower-SNR transits depends sensitively on the amplitude and morphology of the TTV signal (see e.g. Holczer et al. 2016) and requires system-specific dynamical modelling beyond the scope of this work. However, these results suggest that a significantly larger fraction of the planets in these systems could receive additional dynamical constraints than the fraction whose transits can be individually timed.

Another alternative to individual transit fitting is the photodynamical approach, introduced by Ragozzine & Holman (2010), and later applied to a variety of often multiple stellar systems, including KOI-126 (Carter et al., 2011), Kepler-16 (Doyle et al., 2011), and Kepler-47 (Orosz et al., 2012). In contrast to the conventional two-step procedure – in which individual transit times are first measured and subsequently fit with an independent N -body model – photodynamical analyses simultaneously model the system dynamics and photometric observations within a single self-consistent model. This allows the full information content of the light curve to be utilized directly, rather than compressing the observations into a set of measured transit times.

A key advantage of the photodynamical approach is that it can remain effective even when individual transits have low signal-to-noise (see Langford & Agol 2025). By using the entirety of the data, information from many transits can be combined to constrain the orbital architectures and transit parameters, even when individual events are difficult to identify or characterize in isolation. Thus, the opportunity of this approach sparks further optimism that a much larger fraction of the planet sample remains tractable than their transit SNRs would initially suggest.

5.5.6 Individual systems

Jontof-Hutter et al. (2021) used forward dynamical modelling of known *Kepler* TTV systems to identify 26 systems hosting at least one planet whose projected TTVs diverge by more than 90 minutes by BJD 2460900 (corresponding to approximately 2025 August 12, 12:00 UTC), which are therefore prime targets for follow-up transit photometry. From this set, two systems were excluded from our sample because their host stars are likely subgiants (Kepler-87 and Kepler-223), and a further nine systems because none of their planets were flagged as exhibiting TTVs in the NASA Exoplanet Archive (Kepler-149, Kepler-184, Kepler-254, Kepler-266, Kepler-372, Kepler-382, Kepler-662, Kepler-822, and Kepler-1802). The remaining 15 systems, hosting 42 transiting planets, were included in our simulations. These planets, with their periods, calculated *PLATO* transit SNRs, and projected TTV uncertainties at BJD 2460900, are shown in Table 5.1.

Among the 42 planets in these 15 systems, we identify five planets (KOI-1783.01, KOI-1783.02, Kepler-36 c, Kepler-177 c, and Kepler-324 c) with expected transit SNRs sufficient for individual

transit timing measurements ($\text{SNR} \geq 3.8$). These planets are distributed across four systems hosting 10 transiting planets in total. In addition, we identify 27 planets in 13 systems with expected transit SNRs ≥ 1 ; these systems collectively host 38 transiting planets. Of the 15 systems considered, only KOI-1599 and KOI-3503 host no planets with expected $\text{SNR} \geq 1$.

The overlap between the systems identified by Jontof-Hutter et al. (2021) as promising TTV targets and those for which *PLATO* is expected to deliver useful timing information demonstrates that *PLATO* has the potential to meaningfully extend the temporal baseline for a subset of known *Kepler* TTV systems. Even modestly precise timing measurements, when separated by the long baseline between the *Kepler* and *PLATO* missions, can provide valuable information on long-term dynamical evolution.

5.6 Conclusions

In this work, we have assessed the extent to which the *PLATO* mission can deliver useful transit timing measurements for known transiting exoplanets in the *Kepler* field. While *PLATO* will not achieve the per-transit photometric precision of *Kepler* for these targets, its long-duration stare and temporal separation from the *Kepler* era provide an opportunity to extend transit timing baselines by almost two decades in favourable cases.

Using end-to-end simulated *PLATO* photometry, we investigated the relationship between target brightness and photometric precision, including for dim *Kepler* hosts considered out of scope for the prime *PLATO* mission. We found that for all but the most contaminated hosts, noise is mostly white and well modelled by the simple noise-magnitude relation of Cabrera et al. (2026). Moreover, we found that for our sample of *Kepler* hosts, *PLATO* data typically exhibits 2.68–5.14 times the photometric noise of *Kepler* on one-hour timescales, with the ratio varying with the number of co-pointing cameras.

We quantified the precision with which individual mid-transit times can be recovered as a function of transit signal-to-noise ratio. We introduced a posterior-based classification scheme to distinguish between prior-limited, low-significance, and well-constrained timing measurements, and demonstrated that the transition between these regimes occurs gradually over the range $\text{SNR} \simeq 2.1$ –4.3. Additionally, we found that timing precision is only weakly dependent on observational cadence within the range of expected data products to be delivered by the *PLATO* mission. Together, we provide a framework for using catalogue-level quantities to predict photometric precision and, by extension, transit SNR. From these, the predicted mid-transit timing precision can be computed using the relation calibrated in this work.

While most individual transits fall below the threshold for independent detection, a non-negligible fraction nonetheless yield statistically powerful timing constraints when prior knowledge of the ephemeris is available. When individual transits are weak, combining multiple transits using a river diagram or photodynamical approach can provide a plausible way to recover timing information, particularly for systems with large-amplitude TTVs.

We find that *PLATO* is well positioned to contribute additional timing constraints for a subset of known TTV systems, thereby strengthening dynamical inferences through an extended temporal baseline. *PLATO* can play a valuable complementary role in refining the long-term evolution of

Table 5.1: Properties of the 42 planets in the 15 systems from the sample identified by Jontof-Hutter et al. (2021) that were included in our analysis, selected as systems in which the projected TTVs of at least one planet diverge by more than 90 minutes by BJD 2460900. Orbital periods are the default values reported in the NASA Exoplanet Archive. Transit SNRs were calculated in this work, and the projected transit-timing uncertainties for the first transit occurring after BJD 2460900 (δ TTV) are taken from Jontof-Hutter et al. (2021). Transit SNRs exceeding the individual timing threshold of 3.8 are shown in bold.

Planet	Period (d)	Transit SNR	δ TTV (min)
KOI-1599.01	20.44	0.81	1434
KOI-1599.02	13.61	0.61	1132
KOI-1783.01	134.46	18.96	95
KOI-1783.02	284.21	28.76	562
KOI-3503 b	21.18	0.40	390
KOI-3503 c	31.83	0.44	432
Kepler-29 b	10.34	1.42	238
Kepler-29 c	13.29	1.34	310
Kepler-36 b	13.87	2.26	144
Kepler-36 c	16.22	13.86	256
Kepler-49 b	7.20	2.52	18
Kepler-49 c	10.91	2.41	33
Kepler-49 d	2.58	0.81	24
Kepler-49 e	18.60	1.07	278
Kepler-54 b	8.01	1.14	76
Kepler-54 c	12.07	0.42	131
Kepler-54 d	21.00	0.71	49
Kepler-60 b	7.13	1.00	143
Kepler-60 c	8.92	1.28	53
Kepler-60 d	11.90	1.47	230
Kepler-81 b	5.96	2.73	74
Kepler-81 c	12.04	2.94	31
Kepler-81 d	20.84	0.83	364
Kepler-176 b	5.43	0.47	18
Kepler-176 c	12.76	1.81	61
Kepler-176 d	25.75	1.90	>1 day
Kepler-176 e	51.17	0.71	>1 day
Kepler-177 b	36.86	3.02	52
Kepler-177 c	49.41	10.65	97
Kepler-307 b	10.42	2.06	10
Kepler-307 c	13.07	1.75	19
Kepler-324 b	4.39	0.61	152
Kepler-324 c	51.81	7.05	41
Kepler-324 d	34.17	0.84	218
Kepler-324 e	13.98	2.36	34
Kepler-342 b	15.17	2.14	19
Kepler-342 c	26.23	1.77	711
Kepler-342 d	39.46	3.08	28
Kepler-342 e	1.64	0.23	299
Kepler-359 b	25.56	0.71	>1 day
Kepler-359 c	57.69	1.21	759
Kepler-359 d	77.10	1.10	>1 day

carefully selected systems identified from *Kepler* data.

Chapter 6

Conclusions



*“You wanna know how I did it? This is how I did it, Anton.
I never saved anything for the swim back.”*

‘Gattaca’, Andrew Niccol, 1997

6.1 Thesis Summary

The work presented in this thesis has explored several aspects of a common theme: the acquisition, characterization, and interpretation of high-precision photometric time series.

The first part of this thesis focused on the commissioning and characterization of the TWenty Inch Survey Telescope (TWIST), a new robotic observatory developed at the University of Warwick. Measurements of dark current, read noise, gain, and photometric precision showed that the system performs broadly as expected from simple noise modelling. The resulting facility is capable of producing high-cadence, high-duty-cycle photometry across a wide range of target brightnesses, showing stability on months-long timescales. The results of this investigation add to the growing body of evidence that modern scientific CMOS detectors can deliver photometric performance comparable to CCDs at a fraction of the cost and power requirements.

The scientific observations presented in Chapter 3 demonstrated the facility’s versatility. TWIST has been used to probe a range of astrophysical phenomena, including exoplanet transits, multi-periodic pulsations in white dwarfs, and short-duration transits caused by planetary debris orbiting white dwarfs, showing the breadth of science accessible to relatively modest-aperture robotic facilities. The results reinforce the importance of small observatories operating in coordination with major facilities and space missions. While flagship observatories provide exceptional sensitivity, flexible ground-based telescopes are essential for long-baseline monitoring and rapid-response observations.

In Chapter 4, TWIST was used to conduct a 67-night survey of the globular cluster Messier 3. Despite M3 being one of the most extensively studied globular clusters in the Galaxy, the survey revealed four previously unknown SX Phoenicis variables, including two with variability undetectable by *Gaia*, highlighting the value of dense temporal sampling and long-baseline observations even in systems that have been observed for more than a century with photographic plates and CCDs. The survey provided updated periods and amplitudes for a large fraction of the cluster’s known RR Lyrae population, as well as the Blazhko modulation properties for more than 50 stars, contributing to the further cataloguing of these stars’ evolution on timescales longer than a human lifetime, which may be useful for future theoretical investigations. The development of a scattered-light correction procedure for flat-field calibration was also essential to achieving the long-term photometric stability required for this work.

The final scientific component of this thesis considered the future capabilities of ESA’s *PLATO* mission. Using end-to-end simulations incorporating realistic detector behaviour, contamination, stellar variability, and transit fitting, I investigated the expected transit-timing performance for known *Kepler* planetary systems if the mission observes the *Kepler* field. The simulations indicate that *PLATO* would generally deliver lower photometric precision than *Kepler* for these targets. Nevertheless, the decades-long temporal baseline separating the two missions creates an opportunity to probe long-term dynamical evolution in planetary systems that may otherwise remain inaccessible.

The analysis demonstrated that transit timing precision is strongly linked to transit signal-to-noise ratio, and that useful timing measurements can be recovered even in cases where individual transits would not be independently detectable. Several systems were identified for which *PLATO* observations could significantly extend existing transit timing baselines and improve constraints on dynamical models. These predictions for the scientific return of an observing field overlapping with *Kepler*’s directly inform the mission design for the second half of *PLATO*’s nominal mission, which

at the present time is undecided.

The characterization of TWIST shows that relatively small, modern observatories can achieve photometric performance sufficient for competitive time-domain science. The survey of M3 illustrates how new discoveries remain possible even within well-studied stellar systems when observations combine long baselines with high-cadence sampling and modern detector technology. The *PLATO* simulations demonstrate how future space missions can exploit decades-long temporal baselines to address questions about planetary system evolution.

6.2 Future Work

Several avenues for future work naturally arise from these results. For TWIST, the installation and characterization of next-generation CMOS detectors will provide an opportunity to assess how evolving detector technologies influence precision photometry in practice. Multi-band observations using the telescope’s filter set would additionally allow colour-dependent variability studies and improved classification of variable stars.

The M3 survey could be extended in several ways. Additional observing seasons would improve sensitivity to long-period and low-amplitude variability, particularly among Blazhko RR Lyrae stars with modulation periods of hundreds of days. Multi-colour photometry would allow improved characterization of the newly discovered SX Phoenicis members. Radial velocity follow-up of the unclassified variables V286 and V287 could confirm their binary nature, or additional photometric coverage showing long-term amplitude modulation could distinguish the hypothesized chromospheric activity as the source of variability from more typical pulsation-induced variability. Alternatively, simple spectroscopy could detect Ca II H and K emission as a signature of chromospheric activity.

The demonstration of the detectability of new SXP stars in a well-studied cluster naturally motivates the study of additional clusters. Selection criteria for future surveys could include the incidence of blue straggler stars, for which *Gaia* data would be invaluable, and a lack of known SXP stars. Given the characterization of TWIST’s photometric precision, it would be simple to simulate the observing time needed to identify these stars, and time could be allocated to each cluster on this basis.

For the *PLATO* simulations, the most immediate extension would be to incorporate improvements to the mission pipelines as they become available, enabling progressively more realistic predictions. Following the final field selection, applying the methodology to the true Northern field (if selected) would provide yet further realism. Since detectability of TTVs depends not only on the individual transit SNRs, but also on the TTV amplitudes, future work could use the injection of realistic TTVs and explore the use of photodynamical modelling and machine-learning approaches for extracting transit timing information from low signal-to-noise observations. Although such an approach requires system-specific modelling, the present work supports the identification of the most promising systems for in-depth analysis. Ultimately, once *PLATO* begins operations, the predictions presented here can be tested directly against flight data.

Bibliography

- Adams F. C., Laughlin G., 1997, *Reviews of Modern Physics*, 69, 337
- Agol E., Fabrycky D. C., 2025, in , *Handbook of Exoplanets*. Springer, Cham, pp 1–40, doi:10.1007/978-3-319-30648-3_7-2, https://link.springer.com/rwe/10.1007/978-3-319-30648-3_7-2
- Agol E., Steffen J., Sari R., Clarkson W., 2005, *Monthly Notices of the Royal Astronomical Society*, 359, 567
- Aigrain S., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 3211
- Airey R. J., et al., 2025, *Advances in Space Research*, 75, 5757
- Alarcon M. R., Licandro J., Serra-Ricart M., Joven E., Gaitan V., Sousa R. d., 2023, *Publications of the Astronomical Society of the Pacific*, 135, 055001
- Albrow M. D., et al., 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 2099
- Alcock C., et al., 2000, *The Astrophysical Journal*, 542, 257
- Alexander R., Pascucci I., Andrews S., Armitage P., Cieza L., 2013, *The Dispersal of Protoplanetary Disks*, doi:10.2458/azu'uapress'9780816531240-ch021, <https://arxiv.org/abs/1311.1819v1>
- Anderson J., Cool A. M., King I. R., 2003, *The Astrophysical Journal*, 597, L137
- André P., et al., 2010, *Astronomy & Astrophysics*, 518, L102
- Apergis I., et al., 2025, *RAS Techniques and Instruments*, 4, rzaf049
- Apergis I., et al., 2026, *RAS Techniques and Instruments*, p. rzag022
- Arakelyan N. R., Pilipenko S. V., Libeskind N. I., 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 918
- Arzoumanian D., et al., 2011, *Astronomy & Astrophysics*, 529, L6
- Astropy Collaboration et al., 2013, *A&A*, 558, A33
- Astropy Collaboration et al., 2018, *AJ*, 156, 123
- Astropy Collaboration et al., 2022, *ApJ*, 935, 167

- Balona L. A., Nemec J. M., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 2413
- Baranne A., et al., 1996, *Astronomy and Astrophysics Supplement Series*, 119, 373
- Barbary K., 2016, *Journal of Open Source Software*, 1, 58
- Barclay T., Pepper J., Quintana E. V., 2018, *The Astrophysical Journal Supplement Series*, 239, 2
- Bastian N., Lardo C., 2018, *Annual Review of Astronomy and Astrophysics*, 56, 83
- Batista V., Beaulieu J.-P., Bennett D. P., Gould A., Marquette J.-B., Fukui A., Bhattacharya A., 2015, *The Astrophysical Journal*, 808, 170
- Baumgardt H., Hilker M., 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 1520
- Becker A., 2015, HOTPANTS: High Order Transform of PSF ANd Template Subtraction, *Astrophysics Source Code Library*, record ascl:1504.004
- Benkő J. M., Bakos G. A., Nuspl J., 2006, *Monthly Notices of the Royal Astronomical Society*, 372, 1657
- Bertin E., Arnouts S., 1996, *Astronomy and Astrophysics Supplement Series*, 117, 393
- Bhardwaj A., Rejkuba M., Grijs R. d., Herczeg G. J., Singh H. P., Kanbur S., Ngeow C.-C., 2020, *The Astronomical Journal*, 160, 220
- Birnstiel T., Dullemond C. P., Brauer F., 2009, *Astronomy & Astrophysics*, 503, L5
- Bland-Hawthorn J., Gerhard O., 2016, *Annual Review of Astronomy and Astrophysics*, 54, 529
- Blažko S., 1907, *Astronomische Nachrichten*, 175, 325
- Bono G., Caputo F., Santolamazza P., 1997, *Astronomy and Astrophysics*, 317, 171
- Bono G., et al., 2020, *Astronomy & Astrophysics*, 644, A96
- Bonomo A. S., et al., 2023, *Astronomy & Astrophysics*, 677, A33
- Borgne J., et al., 2014, *Monthly Notices of the Royal Astronomical Society*, 441
- Borucki W. J., 2016, *Reports on Progress in Physics*, 79, 036901
- Borucki W. J., et al., 2010, *Science*, 327, 977
- Bowler B. P., 2016, *Publications of the Astronomical Society of the Pacific*, 128, 102001
- Boyle W. S., Smith G. E., 1970, *Bell Labs Technical Journal*, 49, 587
- Bradley L., et al., 2023, *astropy/photutils: 1.9.0*, doi:10.5281/zenodo.8248020
- Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *Monthly Notices of the Royal Astronomical Society*, 427, 127
- Brodie J. P., Strader J., 2006, *Annual Review of Astronomy and Astrophysics*, 44, 193

- Buchler J. R., Kolláth Z., 2011, *The Astrophysical Journal*, 731, 24
- Burke C. J., Mullally F., Thompson S. E., Coughlin J. L., Rowe J. F., 2019, *The Astronomical Journal*, 157, 143
- Busso G., et al., 2022, Gaia DR3 documentation Chapter 5: Photometric data, Gaia DR3 documentation, European Space Agency; Gaia Data Processing and Analysis Consortium., <https://ui.adsabs.harvard.edu/abs/2022gdr3.reptE...5B>
- Cabrera J., et al., 2026, Assessment of PLATO Science Performance, doi:10.48550/arXiv.2604.04818, <http://arxiv.org/abs/2604.04818>
- Canocchi G., et al., 2023, *Astronomy & Astrophysics*, 672, A144
- Cardelli J. A., Clayton G. C., Mathis J. S., 1989, *The Astrophysical Journal*, 345, 245
- Carretta E., et al., 2009, *Astronomy & Astrophysics*, 505, 117
- Carter J. A., Yee J. C., Eastman J., Gaudi B. S., Winn J. N., 2008, *The Astrophysical Journal*, 689, 499
- Carter J. A., et al., 2011, *Science*, 331, 562
- Casagrande L., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 2684
- Castro-González A., et al., 2024, *A&A*, 691, A233
- Catelan M., 2004, *The Astrophysical Journal*, 600, 409
- Catelan M., 2009, *Astrophysics and Space Science*, 320, 261
- Catelan M., Pritzl B. J., Smith H. A., 2004, *The Astrophysical Journal Supplement Series*, 154, 633
- Chaplin W. J., Miglio A., 2013, *Annual Review of Astronomy and Astrophysics*, 51, 353
- Chauvin G., 2024, *Comptes Rendus. Physique*, 24, 129
- Chote P., et al., 2025, *Advances in Space Research*, 76, 764
- Christensen-Dalsgaard J., 2012, Kepler Asteroseismology of Red-giant Stars, doi:10.48550/arXiv.1110.5012, <https://ui.adsabs.harvard.edu/abs/2012ASPC...462..503C>
- Christiansen J. L., et al., 2025, The NASA Exoplanet Archive and Exoplanet Follow-up Observing Program: Data, Tools, and Usage, doi:10.48550/arXiv.2506.03299, <https://ui.adsabs.harvard.edu/abs/2025arXiv250603299C>
- Claringbold A. B., Wheatley P. J., Kirk J., Ahner E.-M., Skillen I., Brogi M., King G. W., McCormac J., 2026, *Monthly Notices of the Royal Astronomical Society*, 547, stag384
- Clement C. M., et al., 2001, *The Astronomical Journal*, 122, 2587

- Clementini G., et al., 2017, *Astronomy & Astrophysics*, 605, A79
- Cohen R. E., Sarajedini A., 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 342
- Corsaro E., Fröhlich H.-E., Bonanno A., Huber D., Bedding T. R., Benomar O., De Ridder J., Stello D., 2013, *Monthly Notices of the Royal Astronomical Society*, 430, 2313
- Corwin T. M., Carney B. W., 2001, *The Astronomical Journal*, 122, 3183
- Craig M. W., et al., 2015, *Astrophysics Source Code Library*, p. ascl:1510.007
- Creevey O. L., et al., 2023, *Astronomy & Astrophysics*, 674, A26
- Cunningham T. J., Hancock B., Sun C., Yang G., Ortiz M., Wrigley C., Seshadri S., Pain B., 2003, in *IEEE Workshop on Charge-Coupled Devices & Advanced Image Sensors*, Elmau, Germany.
- Czesla S., Schröter S., Schneider C. P., Huber K. F., Pfeifer F., Andreasen D. T., Zechmeister M., 2019, *PyA: Python astronomy-related packages* (ascl:1906.010)
- Desidera S., et al., 2021, *Astronomy & Astrophysics*, 651, A70
- Diouf M., Lin Z., Krishna K., Jr K. C. T., 2022, in *Multiphoton Microscopy in the Biomedical Sciences XXII*. SPIE, pp 88–94, doi:10.1117/12.2610510, <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11965/119650D/Multiphoton-imaging-using-a-quantitative-CMOS-camera/10.1117/12.2610510.full>
- Doyle L. R., et al., 2011, *Science*, 333, 1602
- Dubber S. C., et al., 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 5103
- Durisen R. H., Boss A. P., Mayer L., Nelson A. F., Quinn T., Rice W. K. M., 2007, *Gravitational Instabilities in Gaseous Protoplanetary Disks and Implications for Giant Planet Formation*, doi:10.48550/arXiv.astro-ph/0603179, <https://ui.adsabs.harvard.edu/abs/2007prpl.conf..607D>
- Eschen Y. N. E., Bayliss D., Wilson T. G., Kunimoto M., Pelisoli I., Rodel T., 2024, *Monthly Notices of the Royal Astronomical Society*, 535, 1778
- Espinosa S., Vicuña M. L., Mendez R. A., Silva J. F., Orchard M., 2025, *Astronomy & Astrophysics*, 702, A91
- Evans J., 1998, *The history and practice of ancient astronomy*. Oxford University Press, New York
- Evans I. N., et al., 2024, *The Astrophysical Journal Supplement Series*, 274, 22
- Eyer L., et al., 2019, *Astronomy & Astrophysics*, 623, A110
- Fabrycky D., Tremaine S., 2007, *The Astrophysical Journal*, 669, 1298
- Fabrycky D. C., et al., 2014, *The Astrophysical Journal*, 790, 146

- Feast M. W., 1996, *Monthly Notices of the Royal Astronomical Society*, 278, 11
- Ferraro F. R., Pecci F. F., Cacciari C., Corsi C., Buonanno R., Fahlman G. G., Richer H. B., 1993, *The Astronomical Journal*, 106, 2324
- Ferraro F. R., et al., 1997, *Astronomy and Astrophysics*, 324, 915
- Figueira P., et al., 2025, *Astronomy & Astrophysics*, 700, A174
- Fischer D. A., Valenti J., 2005, *The Astrophysical Journal*, 622, 1102
- Forbes D. A., Bridges T., 2010, *Monthly Notices of the Royal Astronomical Society*, 404, 1203
- Forbes D. A., et al., 2018, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474, 20170616
- Ford E. B., Rasio F. A., Yu K., 2003, *Dynamical Instabilities in Extrasolar Planetary Systems*, doi:10.48550/arXiv.astro-ph/0210275, <https://ui.adsabs.harvard.edu/abs/2003ASPC..294..181F>
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Fortier A., et al., 2024, *Astronomy & Astrophysics*, 687, A302
- Fossum E., 1997, *IEEE Transactions on Electron Devices*, 44, 1689
- Fossum E. R., Hondongwa D. B., 2014, *IEEE Journal of the Electron Devices Society*, 2, 33
- Gaia Collaboration 2022, *VizieR Online Data Catalog*, 1358, I/358
- Gaia Collaboration et al., 2023, *A&A*, 674, A1
- Garro E. R., Minniti D., Fernández-Trincado J. G., 2024, *Astronomy & Astrophysics*, 687, A214
- Gaudi B. S., 2012, *Annual Review of Astronomy and Astrophysics*, 50, 411
- Ge J., et al., 2022, *ET White Paper: To Find the First Earth 2.0*, doi:10.48550/arXiv.2206.06693
- Geller A. M., et al., 2017, *The Astrophysical Journal*, 840, 66
- Gieles M., Zwart S. F. P., Baumgardt H., Athanassoula E., Lamers H. J. G. L. M., Sipior M., Leenaarts J., 2006, *Monthly Notices of the Royal Astronomical Society*, 371, 793
- Gill A. S., et al., 2022, in *X-Ray, Optical, and Infrared Detectors for Astronomy X. SPIE*, pp 374–388, doi:10.1117/12.2627564
- Gilliland R. L., et al., 2010, *Publications of the Astronomical Society of the Pacific*, 122, 131
- Gnedin O. Y., Ostriker J. P., 1997, *The Astrophysical Journal*, 474, 223
- Gould A., 2000, *The Astrophysical Journal*, 542, 785

- Gratton R. G., Carretta E., Bragaglia A., 2012, *The Astronomy and Astrophysics Review*, 20, 50
- Groh J. H., Meynet G., Ekström S., Georgy C., 2014, *Astronomy & Astrophysics*, 564, A30
- Guidry J. A., et al., 2021, *The Astrophysical Journal*, 912, 125
- Guyon O., 2018, *Annual Review of Astronomy and Astrophysics*, 56, 315
- Hadden S., Lithwick Y., 2017, *The Astronomical Journal*, 154, 5
- Haisch Jr. K. E., Lada E. A., Lada C. J., 2001, *The Astrophysical Journal*, 553, L153
- Hall D. S., 1976, *International Astronomical Union Colloquium*, 29, 287
- Hardy L. K., Butterley T., Dhillon V. S., Littlefair S. P., Wilson R. W., 2015, *Monthly Notices of the Royal Astronomical Society*, 454, 4316
- Harris W. E., 2010, *A New Catalog of Globular Clusters in the Milky Way*, doi:10.48550/arXiv.1012.3224, <http://arxiv.org/abs/1012.3224>
- Hartman J. D., Kaluzny J., Szentgyorgyi A., Stanek K. Z., 2005, *The Astronomical Journal*, 129, 1596
- Heggie D., Hut P., 2003, *The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics*. Cambridge University Press, Cambridge, doi:10.1017/CBO9781139164535, <https://www.cambridge.org/core/books/gravitational-millionbody-problem/AFA2C2F4821B5FEF991A22470E889C22>
- Helled R., et al., 2014, *Giant Planet Formation, Evolution, and Internal Structure*, doi:10.2458/azu'uapress'9780816531240-ch028, <https://ui.adsabs.harvard.edu/abs/2014prpl.conf..643H>
- Herschel W., 1789, *Philosophical Transactions of the Royal Society of London*, 79, 212
- Hines T., 2016, *Package containing the tools necessary for radial basis function (RBF) applications*, <https://www.github.com/treverhines/RBF>
- Hipke M., David T. J., Mulders G. D., Heller R., 2019, *The Astronomical Journal*, 158, 143
- Holczer T., et al., 2016, *The Astrophysical Journal Supplement Series*, 225, 9
- Holman M. J., Murray N. W., 2005, *Science*, 307, 1288
- Howell S. B., 2000, *Handbook of CCD Astronomy*. Cambridge University Press, <https://ui.adsabs.harvard.edu/abs/2000hccd.book.....H>
- Jackson B., Greenberg R., Barnes R., 2008, *The Astrophysical Journal*, 678, 1396
- Jähne B., 2010, *Optik & Photonik*, 5, 53
- Janesick J. R., 2001, *Scientific charge-coupled devices*. SPIE Optical Engineering Press, <https://ui.adsabs.harvard.edu/abs/2001sccd.book.....J>

Janesick J., Klaasen K., Elliott T., 1985, in Solid state imaging arrays, pp 7–19, doi:10.1117/12.950297

Jannsen N., et al., 2024, *Astronomy & Astrophysics*, 681, A18

Jenkins J. M., et al., 2010, *The Astrophysical Journal Letters*, 713, L87

Jontof-Hutter D., Wolfgang A., Ford E. B., Lissauer J. J., Fabrycky D. C., Rowe J. F., 2021, *The Astronomical Journal*, 161, 246

Jordán A., et al., 2007, *The Astrophysical Journal Supplement Series*, 171, 101

Jurcsik J., 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 80

Jurcsik J., et al., 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1006

Jurcsik J., Szeidl B., Clement C., Hurta Z., Lovas M., 2011, *Monthly Notices of the Royal Astronomical Society*, 411, 1763

Jurcsik J., et al., 2012, *Monthly Notices of the Royal Astronomical Society*, 419, 2173

Jurcsik J., et al., 2015, *The Astrophysical Journal Supplement Series*, 219, 25

Jurcsik J., et al., 2017, *Monthly Notices of the Royal Astronomical Society*, 468, 1317

Jurcsik J., Hajdu G., Juhász A., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 2468

Kallinger T., et al., 2014, *Astronomy & Astrophysics*, 570, A41

Kaluzny J., Hilditch R. W., Clement C., Rucinski S. M., 1998, *Monthly Notices of the Royal Astronomical Society*, 296, 347

Kane M., Ragozzine D., Flowers X., Holczer T., Mazeh T., Relles H. M., 2019, *The Astronomical Journal*, 157, 171

King D., 1985, RGO/La Palma technical note no 31

Kipping D. M., 2009, *Monthly Notices of the Royal Astronomical Society*, 392, 181

Kipping D., 2023, *Monthly Notices of the Royal Astronomical Society*, 523, 1182

Kleinman S. J., et al., 1994, *The Astrophysical Journal*, 436, 875

Kley W., Nelson R. P., 2012, *Annual Review of Astronomy and Astrophysics*, 50, 211

Koch D. G., et al., 2010, *The Astrophysical Journal Letters*, 713, L79

Kreidberg L., 2015, *Astrophysics Source Code Library*, p. ascl:1510.002

Kruijssen J. M. D., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 3008

Kruijssen J. M. D., 2015, *Monthly Notices of the Royal Astronomical Society*, 454, 1658

Kudashkina L. S., 2019, *Astrophysics*, 62, 556

- Kuehn C. A., Smith H. A., Catelan M., Pritzl B. J., De Lee N., Borissova J., 2012, *The Astronomical Journal*, 144, 186
- Kumar N., Bhardwaj A., Singh H. P., Rejkuba M., Marconi M., Prugniel P., 2024, *Monthly Notices of the Royal Astronomical Society*, 531, 2976
- Kıroğlu F., Weatherford N. C., Kremer K., Ye C. S., Fragione G., Rasio F. A., 2022, *The Astrophysical Journal*, 928, 181
- Lada C. J., Lada E. A., 2003, *Annual Review of Astronomy and Astrophysics*, 41, 57
- Lamers H. J. G. L. M., Levesque E. M., 2017, in , *Understanding Stellar Evolution*. IOP Publishing, doi:10.1088/978-0-7503-1278-3ch21, <https://iopscience.iop.org/book/mono/978-0-7503-1278-3/chapter/bk978-0-7503-1278-3ch21.pdf>
- Lang D., Hogg D. W., Mierle K., Blanton M., Roweis S., 2010, *The Astronomical Journal*, 139, 1782
- Langford Z., Agol E., 2025, *Monthly Notices of the Royal Astronomical Society*, 540, 106
- Law N. M., et al., 2022, *Publications of the Astronomical Society of the Pacific*, 134, 035003
- Leavitt H. S., Pickering E. C., 1912, *Harvard College Observatory Circular*, 173, 1
- Lecar M., Franklin F. A., Holman M. J., Murray N. W., 2001, *Annual Review of Astronomy and Astrophysics*, 39, 581
- Lecavelier des Etangs A., Lissauer J. J., 2022, *New Astronomy Reviews*, 94, 101641
- Lee J.-W., Sneden C., 2021, *The Astrophysical Journal*, 909, 167
- Lee Y.-W., Demarque P., Zinn R., 1990, *The Astrophysical Journal*, 350, 155
- Leleu A., Chatel G., Udry S., Alibert Y., Delisle J.-B., Mardling R., 2021, *Astronomy & Astrophysics*, 655, A66
- Lesser M., 2015, *Publications of the Astronomical Society of the Pacific*, 127, 1097
- Libralato M., Bedin L. R., Nardiello D., Piotto G., 2016a, *Monthly Notices of the Royal Astronomical Society*, 456, 1137
- Libralato M., et al., 2016b, *Monthly Notices of the Royal Astronomical Society*, 463, 1780
- Libralato M., et al., 2022, *The Astrophysical Journal*, 934, 150
- Licandro J., Tonry J., Serra-Ricart M., Alarcon M. R., Denneau L., 2022, *European Planetary Science Congress*
- Lin A. S. J., et al., 2023, *The Astronomical Journal*, 166, 90
- Lithwick Y., Xie J., Wu Y., 2012, *The Astrophysical Journal*, 761, 122
- Lomb N. R., 1976, *Astrophysics and Space Science*, 39, 447

- Longmore A. J., Fernley J. A., Jameson R. F., 1986, *Monthly Notices of the Royal Astronomical Society*, 220, 279
- Longmore A. J., Dixon R., Skillen I., Jameson R. F., Fernley J. A., 1990, *Monthly Notices of the Royal Astronomical Society*, 247, 684
- Louden T., Kreidberg L., 2018, *Monthly Notices of the Royal Astronomical Society*, 477, 2613
- Lovis C., Fischer D. A., 2010, in Seager S., ed., *Exoplanets*. University of Arizona Press, Tucson, pp 27–53
- Lyman J. D., 2021, *Astrophysics Source Code Library*, p. ascl:2103.003
- Lyman J. D., et al., 2026, The Gravitational-wave Optical Transient Observer (GOTO) data pipeline and workflow for transient discovery (arXiv:2603.02330), <https://arxiv.org/abs/2603.02330>
- Maeyama Y., 2002, *Acta Historica Astronomiae*, 16, 113
- Mahalanobis P. C., 1936, *Proceedings of the National Institute of Sciences of India*, 2, 49
- Malbet F., et al., 2022, in *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. SPIE, pp 574–585, doi:10.1117/12.2629927
- Mamajek E. E., et al., 2015, IAU 2015 Resolution B2 on Recommended Zero Points for the Absolute and Apparent Bolometric Magnitude Scales, doi:10.48550/arXiv.1510.06262, <http://arxiv.org/abs/1510.06262>
- Mandel K., Agol E., 2002, *The Astrophysical Journal*, 580, L171
- Marchiori V., et al., 2019, *Astronomy & Astrophysics*, 627, A71
- Marois C., Macintosh B., Barman T., Zuckerman B., Song I., Patience J., Lafrenière D., Doyon R., 2008, *Science*, 322, 1348
- Marois C., Zuckerman B., Konopacky Q. M., Macintosh B., Barman T., 2010, *Nature*, 468, 1080
- Marín-Franch A., et al., 2009, *The Astrophysical Journal*, 694, 1498
- Massari D., Lapenna E., Bragaglia A., Dalessandro E., Contreras Ramos R., Amigo P., 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 4162
- Massari D., Koppelman H. H., Helmi A., 2019, *Astronomy & Astrophysics*, 630, L4
- Massari D., Marasco A., Beltramo-Martin O., Milli J., Fiorentino G., Tolstoy E., Kerber F., 2020, *Astronomy & Astrophysics*, 634, L5
- May B. H., 2008, A survey of radial velocities in the zodiacal dust cloud, doi:<https://doi.org/10.25560/1333>, <http://hdl.handle.net/10044/1/1333>
- Mayor M., Queloz D., 1995, *Nature*, 378, 355

- McCormac J., Pollacco D., Skillen I., Faedi F., Todd I., Watson C. A., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 548
- McCormac J., Skillen I., Pollacco D., Faedi F., Ramsay G., Dhillon V. S., Todd I., Gonzalez A., 2014, *Monthly Notices of the Royal Astronomical Society*, 438, 3383
- McKee C. F., Ostriker E. C., 2007, *Annual Review of Astronomy and Astrophysics*, 45, 565
- McNamara D. H., 2011, *The Astronomical Journal*, 142, 110
- Messier C., 1781, *Catalogue des Nébuleuses et des Amas d'Étoiles (Catalog of Nebulae and Star Clusters)*, *Connaissance des Temps ou des Mouvements Célestes*, for 1784, p. 227-267
- Meylan G., Sarajedini A., Jablonka P., Djorgovski S. G., Bridges T., Rich R. M., 2001, *The Astronomical Journal*, 122, 830
- Miret-Roig N., 2023, *Astrophysics and Space Science*, 368, 17
- Mitchell M. A., et al., 2026a, *Monthly Notices of the Royal Astronomical Society*, 545, staf2100
- Mitchell M. A., et al., 2026b, *Monthly Notices of the Royal Astronomical Society*, 550, stag1174
- Murphy S. J., Hey D., Van Reeth T., Bedding T. R., 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 2380
- Mustill A. J., Villaver E., 2012, *The Astrophysical Journal*, 761, 121
- Nascimbeni V., et al., 2022, *Astronomy & Astrophysics*, 658, A31
- Nemec J., Mateo M., 1990, in *SX Phoenicis Stars*. pp 64–84, <https://ui.adsabs.harvard.edu/abs/1990ASPC...11...64N>
- Netzel H., Smolec R., Soszyński I., Udalski A., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 1229
- Ngeow C.-C., Bhardwaj A., Graham M. J., Groom S. L., Masci F. J., Riddle R., 2022, *The Astronomical Journal*, 164, 191
- Olech A., Dziembowski W. A., Pamyatnykh A. A., Kaluzny J., Pych W., Schwarzenberg-Czerny A., Thompson I. B., 2005, *Monthly Notices of the Royal Astronomical Society*, 363, 40
- Oosterhoff P. T., 1939, *The Observatory*, 62, 104
- Oosterhoff P. T., 1944, *Bulletin of the Astronomical Institutes of the Netherlands*, 10, 55
- Orosz J. A., et al., 2012, *Science*, 337, 1511
- Osborn J., Föhring D., Dhillon V. S., Wilson R. W., 2015, *Monthly Notices of the Royal Astronomical Society*, 452, 1707
- Paczynski B., 1986, *Astrophysical Journal*, Part 1, 304

- Paczynski B., 1994, *Acta Astronomica*, 44, 235
- Paegert M., Stassun K. G., Collins K. A., Pepper J., Torres G., Jenkins J., Twicken J. D., Latham D. W., 2021, TESS Input Catalog versions 8.1 and 8.2: Phantoms in the 8.0 Catalog and How to Handle Them, doi:10.48550/arXiv.2108.04778, <https://ui.adsabs.harvard.edu/abs/2021arXiv210804778P>
- Parviainen H., Aigrain S., 2015, *Monthly Notices of the Royal Astronomical Society*, 453, 3821
- Pecaut M. J., Mamajek E. E., 2013, *The Astrophysical Journal Supplement Series*, 208, 9
- Penny M. T., Scott Gaudi B., Kerins E., Rattenbury N. J., Mao S., Robin A. C., Calchi Novati S., 2019, *The Astrophysical Journal Supplement Series*, 241, 3
- Pepe F. A., et al., 2010, in *Ground-based and Airborne Instrumentation for Astronomy III*. SPIE, pp 209–217, doi:10.1117/12.857122, <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/7735/77350F/ESPRESSO--the-Echelle-spectrograph-for-rocky-exoplanets-and-stable/10.1117/12.857122.full>
- Petit A. C., Laskar J., Boué G., 2018, *Astronomy & Astrophysics*, 617, A93
- Pfeffer J., Kruijssen J. M. D., Crain R. A., Bastian N., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 4309
- Pickering E. C., 1890, *Astronomische Nachrichten*, 123, 207
- Pickering E. C., Bailey S. I., 1895, *The Astrophysical Journal*, 2, 321
- Pietrukowicz P., Kaluzny J., Thompson I. B., Jaroszynski M., Schwarzenberg-Czerny A., Krzeminski W., Pych W., 2005, Cluster AgeS Experiment (CASE): Dwarf Novae and a Probable Microlensing Event in the Globular Cluster M22, doi:10.48550/arXiv.astro-ph/0509797, <https://ui.adsabs.harvard.edu/abs/2005AcA....55..261P>
- Pietrukowicz P., Minniti D., Jetzer P., Alonso-García J., Udalski A., 2011, *The Astrophysical Journal Letters*, 744, L18
- Piskunov A. E., Kharchenko N. V., Röser S., Schilbach E., Scholz R.-D., 2006, *Astronomy & Astrophysics*, 445, 545
- Pollacco D. L., et al., 2006, *Publications of the Astronomical Society of the Pacific*, 118, 1407
- Pollack J. B., Hubickyj O., Bodenheimer P., Lissauer J. J., Podolak M., Greenzweig Y., 1996, *Icarus*, 124, 62
- Price E. M., Rogers L. A., 2014, *The Astrophysical Journal*, 794, 92
- Pritzl B., Smith H. A., Catelan M., Sweigart A. V., 2000, *The Astrophysical Journal*, 530, L41
- Prudil Z., Arellano Ferro A., 2024, *Monthly Notices of the Royal Astronomical Society*, 534, 3654

- Prusti T., et al., 2016, *Astronomy & Astrophysics*, 595, A1
- Pych W., Kaluzny J., Krzeminski W., Schwarzenberg-Czerny A., Thompson I. B., 2001, *Astronomy & Astrophysics*, 367, 148
- Ragozzine D., Holman M. J., 2010, The Value of Systems with Multiple Transiting Planets, doi:10.48550/arXiv.1006.3727, <http://arxiv.org/abs/1006.3727>
- Ramsay S., et al., 2021, *The Messenger*, 182, 3
- Rathour R. S., Smolec R., Hajdu G., Ziólkowska O., Hocdé V., Soszyński I., Udalski A., Karczmarek P., 2025, *Astronomy & Astrophysics*, 695, A114
- Rauer H., et al., 2025, *Experimental Astronomy*, 59, 26
- Richards J. W., Starr D. L., Miller A. A., Bloom J. S., Butler N. R., Brink H., Crellin-Quick A., 2012, *The Astrophysical Journal Supplement Series*, 203, 32
- Ricker G. R., et al., 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Rieke F., Baylor D. A., 1998, *Reviews of Modern Physics*, 70, 1027
- Riello M., et al., 2021, *Astronomy & Astrophysics*, 649, A3
- Rimoldini L., et al., 2019, *Astronomy & Astrophysics*, 625, A97
- Rimoldini L., et al., 2023, *Astronomy & Astrophysics*, 674, A14
- Ripepi V., et al., 2024, *Astronomy & Astrophysics*, 682, A1
- Rodríguez E., López-González M. J., 2000, *Astronomy and Astrophysics*, 359, 597
- Rozyczka M., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 517, 2485
- Sahu K. C., Casertano S., Livio M., Gilliland R. L., Panagia N., Albrow M. D., Potter M., 2001, *Nature*, 411, 1022
- Sahu K. C., Anderson J., King I. R., 2002, *The Astrophysical Journal*, 565, L21
- Samadi R., et al., 2019, *Astronomy & Astrophysics*, 624, A117
- Samus' N. N., Kazarovets E. V., Durlevich O. V., Kireeva N. N., Pastukhova E. N., 2017, *Astronomy Reports*, 61, 80
- Sandage A. R., 1953, *Astronomical Journal*, Vol. 58, p. 61-75 (1953), 58, 61
- Sarajedini A., 2025, *The Astronomical Journal*, 170, 189
- Sawyer Hogg H., 1973, *Publications of the David Dunlap Observatory*, 3, 6
- Scargle J. D., 1982, *The Astrophysical Journal*, 263, 835
- Schröder K.-P., Cannon Smith R., 2008, *Monthly Notices of the Royal Astronomical Society*, 386, 155

- Schwarzenberg-Czerny A., 1997, *The Astrophysical Journal*, 489, 941
- Sebastian A., 2001, *A Dictionary of the History of Science*. Taylor & Francis
- Shapley H., 1918a, *Publications of the Astronomical Society of the Pacific*, 30, 42
- Shapley H., 1918b, *The Astrophysical Journal*, 48, 154
- Shields A. L., Ballard S., Johnson J. A., 2016, *Physics Reports*, 663, 1
- Shupe D. L., Moshir M., Li J., Makovoz D., Narron R., Hook R. N., 2005, in Shopbell P., Britton M., Ebert R., eds, *Astronomical Society of the Pacific Conference Series Vol. 347, Astronomical Data Analysis Software and Systems XIV*. p. 491, <https://ui.adsabs.harvard.edu/abs/2005ASPC...347..491S>
- Siegel M. H., Porterfield B. L., Balzer B. G., Hagen L. M. Z., 2015, *The Astronomical Journal*, 150, 129
- Skarka M., Prudil Z., Jurcsik J., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 1237
- Smith H. A., 2004, *RR Lyrae Stars*. Cambridge University Press
- Smolec R., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 447, 3756
- Smolec R., Moskalik P., Kałużny J., Pych W., Różyczka M., Thompson I. B., 2017, *Monthly Notices of the Royal Astronomical Society*, p. stx088
- Soszynski I., et al., 2014, *Over 38000 RR Lyrae Stars in the OGLE Galactic Bulge Fields*, doi:10.48550/arXiv.1410.1542, <http://arxiv.org/abs/1410.1542>
- Soszyński I., et al., 2015, *The OGLE Collection of Variable Stars. Anomalous Cepheids in the Magellanic Clouds*, doi:10.48550/arXiv.1508.00907, <https://ui.adsabs.harvard.edu/abs/2015AcA....65..233S>
- Soszyński I., et al., 2016, *The OGLE Collection of Variable Stars. Over 45 000 RR Lyrae Stars in the Magellanic System*, doi:10.48550/arXiv.1606.02727, <http://arxiv.org/abs/1606.02727>
- Sozzetti A., 2024, *Comptes Rendus. Physique*, 24, 45
- Spitzer Jr. L., 1969, *The Astrophysical Journal*, 158, L139
- Stacey W. S., Heinke C. O., Elsner R. F., Edmonds P. D., Weisskopf M. C., Grindlay J. E., 2011, *The Astrophysical Journal*, 732, 46
- Stellingwerf R. F., 1978, *The Astrophysical Journal*, 224, 953
- Stetson P. B., 1987, *Publications of the Astronomical Society of the Pacific*, 99, 191
- Sullivan P. W., et al., 2015, *The Astrophysical Journal*, 809, 77
- Szabó R., 2013, *Proceedings of the International Astronomical Union*, 9, 241

- Szabó R., et al., 2010, *Monthly Notices of the Royal Astronomical Society*, 409, 1244
- Szabó R., et al., 2014, *Astronomy & Astrophysics*, 570, A100
- Sódor A., et al., 2011, *Monthly Notices of the Royal Astronomical Society*, 411, 1585
- Teranishi N., Kohono A., Ishihara Y., Oda E., Arai K., 1982, in 1982 International Electron Devices Meeting. pp 324–327, doi:10.1109/IEDM.1982.190285, <https://ieeexplore.ieee.org/abstract/document/1482819>
- Thompson S. J., et al., 2016, in *Ground-based and Airborne Instrumentation for Astronomy VI*. SPIE, pp 1949–1961, doi:10.1117/12.2232111, <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9908/99086F/HARPS3-for-a-roboticized-Isaac-Newton-Telescope/10.1117/12.2232111.full>
- Thompson S. E., et al., 2018, *The Astrophysical Journal Supplement Series*, 235, 38
- Thompson I. B., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 492, 4254
- Torres G., et al., 2017, *The Astronomical Journal*, 154, 264
- Trabucchi M., Mowlavi N., 2022, *Astronomy & Astrophysics*, 658, L1
- Trabucchi M., Mowlavi N., Lebzelter T., 2021, *Astronomy & Astrophysics*, 656, A66
- Trager S. C., King I. R., Djorgovski S., 1995, *The Astronomical Journal*, 109, 218
- Trombino D., 1980, *Journal of the British Astronomical Association*, 90, 565
- Valberg A., 2005, *Light Vision Color*. John Wiley & Sons
- Valentinuzzi M. E., 2019, *IEEE Pulse*, 10, 23
- Van Cleve J. E., Caldwell D. A., 2016, Technical report, *Kepler Instrument Handbook*, <https://ui.adsabs.harvard.edu/abs/2016ksci.rept....1V>. NASA Ames Research Center, <https://ui.adsabs.harvard.edu/abs/2016ksci.rept....1V>
- VandenBerg D. A., Brogaard K., Leaman R., Casagrande L., 2013, *The Astrophysical Journal*, 775, 134
- VandenBerg D. A., Denissenkov P. A., Catelan M., 2016, *The Astrophysical Journal*, 827, 2
- VanderPlas J. T., 2018, *The Astrophysical Journal Supplement Series*, 236, 16
- Vanderbosch Z. P., et al., 2021, *The Astrophysical Journal*, 917, 41
- Vasiliev E., Baumgardt H., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 5978
- Veras D., 2016, *Royal Society Open Science*, 3, 150571
- Vigan A., et al., 2021, *Astronomy & Astrophysics*, 651, A72
- Virtanen P., et al., 2020, *Nature Methods*, 17, 261

- Wallace A. L., Casey A. R., Brown A. G. A., Castro-Ginard A., 2025, *Monthly Notices of the Royal Astronomical Society*, 536, 2485
- Wallerstein G., 2002, *Publications of the Astronomical Society of the Pacific*, 114, 689
- Wenger M., et al., 2000, *Astronomy and Astrophysics Supplement Series*, 143, 9
- Wheatley P. J., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 4476
- Willfried Baatz 1997, *Photography. Barron's*, <http://archive.org/details/photography00baatz>
- Winn J. N., 2010, in Seager S., ed., *Exoplanets*. University of Arizona Press, Tucson, pp 55–77
- Winn J. N., et al., 2007, *The Astronomical Journal*, 134, 1707
- Wolszczan A., 1994, *Science*, 264, 538
- Wolszczan A., Frail D. A., 1992, *Nature*, 355, 145
- Wurster J., Li Z.-Y., 2018, *Frontiers in Astronomy and Space Sciences*, 5
- Wyrzykowski L., et al., 2011, *Monthly Notices of the Royal Astronomical Society*, 416, 2949
- Yanoff M., Duker J. S., 2009, *Ophthalmology. Elsevier Health Sciences*
- Young A. T., 1967, *The Astronomical Journal*, 72, 747
- Zejda M., Paunzen E., Baumann B., Mikulášek Z., Liška J., 2012, *Astronomy & Astrophysics*, 548, A97
- Zhao Y., Heinke C. O., Cohn H. N., Lugger P. M., Cool A. M., 2019, *Monthly Notices of the Royal Astronomical Society*, 483, 4560
- Zwart S. F. P., McMillan S. L. W., Gieles M., 2010, *Annual Review of Astronomy and Astrophysics*, 48, 431
- van Albada T. S., Baker N., 1973, *The Astrophysical Journal*, 185, 477