WARWICK
THE UNIVERSITY OF WARWICK

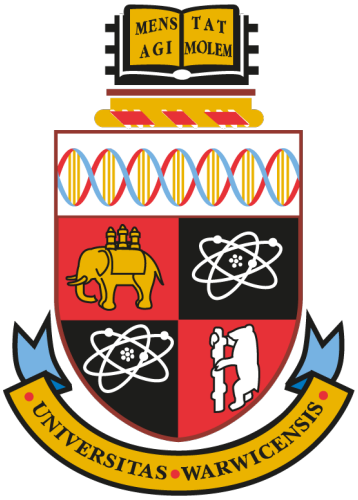

Characterisation of Resident Space Objects Using High-Precision Photometry

by

Robert James Scott Airey

A thesis submitted for the degree of

Doctor of Philosophy

UNIVERSITY OF WARWICK

Department of Physics

May 2026

Contents

Contents iii

List of figures vii

List of tables xi

Acknowledgements xiii

Declaration xv

Abstract xvii

List of abbreviations xix

1 Introduction 1

1.1 Orbital Regimes . . . 1

1.1.1 The GEO-belt . . . 7

1.2 Resident Space Objects . . . 7

1.2.1 Classification and Identification . . . 8

1.2.2 Challenges and Trends . . . 12

1.2.2.1 Space Debris: The Current Picture . . . 12

1.2.2.2 Challenges at LEO & GEO . . . 14

1.2.2.3 Mitigation Strategies . . . 15

1.3 Thesis outline . . . 18

2 Methods 19

2.1 Technical Background . . . 19

2.1.1 Reference Frames, Geometries and Time . . . 19

2.1.1.1 Fundamental Concepts . . . 19

2.1.1.2 Common Reference Frames and Coordinates . . . 21

2.1.1.3 Useful Geometry . . . 24

2.1.1.4 Time Systems in Space Applications . . . 26

2.1.2 Optical Brightness Models . . . 28

2.1.2.1 Fundamentals . . . 29

2.1.2.2 Bidirectional Reflectance Distribution Function . . . 32

2.2 Data Sources . . . 34

2.2.1 Instrumental and Observational Capabilities Required for RSO Characterisation . . . 34
2.2.2 SuperWASP . . . 36
2.2.3 STING . . . 38
2.2.4 CLASP . . . 40
2.3 Observational Constraints and Strategies . . . 41
2.4 Data Reduction . . . 47
2.4.1 Instrumentation . . . 47
2.4.2 Standard astronomical calibrations . . . 48
2.4.3 Astrometric and photometric calibrations . . . 50
2.4.4 Target Detection and Extraction . . . 51
2.4.5 Aperture photometry and light curve extraction . . . 52
2.5 Simulating RSO Light Curves . . . 54
2.5.1 Orbital and observing geometry . . . 54
2.5.2 Simplified Spacecraft Geometry . . . 55
2.5.3 Surface-scattering model . . . 57
2.5.4 Facet flux summation . . . 57
2.5.5 Attitude models . . . 57
2.5.5.1 Uniform spin . . . 58
2.5.5.2 Torque-free tumbling . . . 59
2.5.6 Illustrative light curve comparison . . . 60
2.6 Light Curve Analysis Techniques . . . 62
2.6.1 Period Determination Methods . . . 62
2.6.2 Relating Light Curve Features to Physical Features . . . 65
2.7 Markov Chain Monte Carlo . . . 67
2.7.1 Assumptions and limitations . . . 68

3 Characterisation of the STING Instrument 71
3.1 Motivations . . . 71
3.2 Photometric Performance . . . 71
3.3 Horizon Limits . . . 75
3.4 GPS Ephemeris Comparison . . . 79
3.4.1 GPS Ephemeris Data . . . 80
3.4.2 Observing Strategy . . . 81
3.4.3 Position Extraction and Comparison . . . 83
3.5 Summary . . . 87

4 A Comprehensive Survey of the GEO-belt with STING 89
4.1 Motivations . . . 89
4.2 Observational Strategy . . . 90
4.3 Overview of Survey Demographics . . . 90
4.4 Results and Analysis . . . 97

4.4.1 Induced Colour Error . . . 97
4.4.2 Colour Contrasts Between Bus Configurations . . . 98
4.4.3 Colour-Colour Diagram . . . 101
4.4.4 Light Curve Features . . . 102
4.4.4.1 LDPE-3A Manoeuvres . . . 102
4.4.4.2 Fitting Manoeuvres . . . 103
4.4.4.3 Blue Colour Features . . . 113
4.4.4.4 Triple Peak Features . . . 114
4.4.4.5 Short Timescale Glints . . . 116
4.4.5 Average Colour Measurements . . . 118
4.4.6 Colour Features in Relation with Physical Parameters . . . 120
4.4.7 Solar Panel Offsets . . . 123
4.4.7.1 Solar Panel Glint Colours . . . 125
4.5 Discussion . . . 125
4.6 Summary . . . 127

5 Mapping the Spin-state of the Ajisai Satellite 129
5.1 Motivations . . . 129
5.2 Ajisai Background . . . 130
5.3 Methodology . . . 132
5.3.1 Reference Frame Transformation and Vector Calculations . . . 133
5.3.2 Interpreting Brightness Maps . . . 136
5.3.3 MCMC Fitting . . . 136
5.4 SuperWASP Observations of Ajisai (2019) . . . 138
5.4.1 Instrument and Observations . . . 139
5.4.2 Results . . . 139
5.5 STING Observation of Ajisai (2025) . . . 144
5.5.1 Instrument and Observations . . . 144
5.5.2 Observation Planning . . . 144
5.5.3 Detection Pipeline . . . 146
5.5.4 Results . . . 147
5.6 Summary . . . 150

6 Three Techniques for Untangling Gorizont-6's Spin-state 153
6.1 Motivations . . . 153
6.2 Gorizont-6 Background . . . 155
6.3 Polarimetry . . . 156
6.4 Observational Strategy . . . 157
6.5 Results . . . 159
6.5.1 Period Analysis . . . 159
6.5.2 Phase Folded Light Curves . . . 164
6.5.2.1 Polarisation behaviour . . . 166

6.5.3 Spin Axis Constraint 167
6.6 Discussion 169
6.7 Summary 171

7 Summary **173**
7.1 Meeting the Required Capabilities 174
7.2 Future Outlook 175

Bibliography **181**

A BRDF Models **203**

B Typical Frame Transforms **207**

C Ajisai Empirical Model Calculations **211**

D Haversine Equations **213**

E Ajisai Individual Corner Plots **215**

List of figures

1.1 Illustration of the Keplerian orbital elements. 2
1.2 A diagram representing the most populous orbital regimes with respect to the Earth . 4
1.3 A diagram representing the locations of the facilities used within the Space Surveillance Network (SSN) 8
1.4 Example TLE of a satellite 10
1.5 Fractional representation of fragmentation history 11
1.6 Fractional representation of the object classes scraped from the DISCOSweb database 14
1.7 Distribution of estimated sub-satellite longitudes 16

2.1 Representation of equatorial and horizontal coordinate systems 23
2.2 ECEF coordinates (x, y, z) shown in relation to a geodetic latitude and longitude position 24
2.3 An illustration of the construction of the phase angle bisector vector 25
2.4 Relationship between the sidereal and mean solar day 28
2.5 Demonstration of the difference between a specular reflection versus a diffuse reflection 29
2.6 Parameter sensitivities in satellite visual brightness modelling using simple geometric shapes 31
2.7 Illustration of the BRDF scattering geometry 32
2.8 SuperWASP at the Roque de los Muchachos Observatory in August 2019 37
2.9 A photograph of the STING instrument, at the Roque de los Muchachos observatory 39
2.10 CLASP at the Roque de los Muchachos observatory 39
2.11 Total throughput against wavelength for the CLASP and STING cameras and filters . 40
2.12 Earth shadow geometry and apparent-disk configurations for THOR 6 during eclipse ingress 41
2.13 Projected separation and illumination threshold for THOR 6 and Ajisai across an equinox night 42
2.14 Peak apparent angular rates of catalogued satellites as seen from the ORM site 45
2.15 Comparison of pixel-level readout architectures in CCD and CMOS image sensors . 47
2.16 Illustrative example of a multi-colour calibrated STING light curve 52
2.17 Simulated light curves for uniform spin and torque-free tumbling of a box-wing satellite 61
2.18 Periodograms following the application of the Lomb-Scargle, PDM and FFT methods 64
2.19 Examples of differing light curve features of satellites in different orbital regimes . . 66

3.1 Fractional RMS of STING light curves against calibrated magnitude for the four cameras, compared with the noise model 74

3.2 Example Alt–Az observing sequence for the 9th February 2026 horizon survey . . . 75
3.3 Alt–Az zero-point maps for the four cameras, used to identify horizon limits 76
3.4 Combined obstruction and horizon-limit visualisation 78
3.5 Predicted and measured morphology of a GPS satellite (NAVSTAR 81) streak observed in three cameras 83
3.6 Hexbin density maps of the mid-exposure SP3 ephemeris residuals in local along-track and cross-track coordinates 86

4.1 The observational fields (A-Q) used in this survey of the GEO belt 91
4.2 Demographics of observed objects upon conclusion of the survey 92
4.3 An example four colour light curve of INTELSAT 901 (+ MEV-1) on the 25th April 2023 94
4.4 Illustrative example of G_{RGB} band light curves for SSL-1300 class satellites 95
4.5 Illustrative example of G_{RGB} band light curves for Eurostar-3000 class satellites . . . 96
4.6 Binned RMS values calculated for the sequential and simultaneous measurements of $B_{RGB} - i$ 97
4.7 Colour light curve maps for SSL-1300, Eurostar-3000 and BSS-702 satellites 99
4.8 Colour light curve maps for STAR-2, Spacebus-4000, A2100 and other satellites . . 100
4.9 Colour-colour diagrams showing the average colour measurements of individual satellites 102
4.10 Binned (bin widths of 7.5 minutes) four colour phase light curves for satellite LDPE-3A (NORAD: 55264) observed on the 15th April 2023 104
4.11 Top-down view of the simple 3D model of the LDPE-3A satellite 108
4.12 A demonstration of fitting the simulated light curve to the observed ‘dips’ caused by LDPE-3A manoeuvring 111
4.13 Reconstructed attitude angles for two LDPE-3A manoeuvres derived from the fitted model 112
4.14 Colour light curve maps for Spacebus-4000 and SSL-1300 satellites showing the blue features 113
4.15 Triple peak features found in the light curves of 6 example satellites in the survey sample 115
4.16 Artistic impression of the BADR-4 satellite 116
4.17 Portion of the BADR-6 light curve in the B_{RGB} and i bands 116
4.18 Examples of the glinting regions seen within light curves of satellites in the dataset . 117
4.19 Distribution of the duration of glinting regions in minutes for all satellites that had glinting regions 119
4.20 Photometric distributions of glinting regions within all four colour bands and the three colour index spaces 119
4.21 Time spent on orbit and dry mass for satellites with a variety of different bus configurations as a function of the average colour 121
4.22 Pearson correlation matrix between physical parameters and colour change measurements 121
4.23 The measured $B_{RGB} - i$ values of the near-symmetrical colour changes as a function of different physical properties 122
4.24 Three examples of the colour response across glints caused from specular reflections off of the solar panels 124

4.25 The distribution of solar panel offsets for 54 satellites in the sample 124
4.26 An artistic impression of the STAR-2 bus configuration design 127

5.1 The Ajisai satellite with mirror panels and CCRs covering its spherical surface 131
5.2 Propagation of the empirical model for the orientation and period of Ajisai's spin axis 132
5.3 Diagram illustrating the defined angles, α and δ of the spin-axis vector, $\vec{S}$, with respect to the axes of the ICRF 134
5.4 Illustration of the rotation sequence to transform a vector $\vec{v}$ from the ICRF to body frame 135
5.5 Examples of brightness maps from the 2019-08-01 observation with three different parameter sets compared to the Ajisai mirror positions 136
5.6 Light curves of the 2019 SuperWASP observations of the Ajisai satellite 138
5.7 The target motion and resulting extraction cadence plotted for the four Ajisai observations. 139
5.8 Corner plots from the MCMC runs on the data from each of the four observation nights of Ajisai in 2019 140
5.9 Brightness maps generated using the MCMC-derived parameter sets 141
5.10 Residual measurements of spin-state parameters (estimate − empirical) for each SuperWASP observation of Ajisai. 142
5.11 Illustration of the detection pipeline for Ajisai in STING science frames 145
5.12 Corner plot from the MCMC run on the data from the night of 2025-08-29 147
5.13 Brightness map generated using the MCMC-derived parameter sets for the 29th August 2025 STING observation of Ajisai 148

6.1 Schematic illustration of the YORP effect for the satellite **Gorizont-6** 153
6.2 An example illustration of one of the satellites belonging to the Gorizont (KAUR-3) family 156
6.3 A collection of light curves of different Gorizont satellites 157
6.4 **Top:** Cartographic map illustrating the site locations for the NGTS (Paranal), STING (La Palma) and LT (La Palma) instruments. **Bottom Left:** A photograph of the Marana sCMOS camera, highlighted within the blue rectangle, mounted on one of the NGTS telescopes. This photograph is reproduced from Apergis et al. (2026) with permission. **Bottom Right:** A photograph of the ORM on La Palma, where STING and the LT are located. 159
6.5 PDM periodograms and phase-folded light curves for 2023–12–12 at $P'_G = 481.0$ s and its first three harmonics 163
6.6 Evolution of the PDM-derived rotation period of Gorizont-6 as a function of time . . 164
6.7 The six observations of Gorizont-6 with the STING telescope and polarimetry measurements (linear polarisation and polarisation angle) from the LT 165
6.8 Simultaneous STING (multi-colour) and NGTS light curves of Gorizont-6 acquired on the 12th December 2023 168
6.9 A diagram explaining the spin-pole estimation from matched glint pairs on the unit sphere. 169

6.10 Total cost surface as a function of trial spin-axes constructed from α,δ of the spin-pole in the ICRS 170

7.1 Demographics of unique targets observed with STING 178

E.1 Corner plot from the MCMC run on the data from the night of 2019-08-01 215
E.2 Corner plot from the MCMC run on the data from the night of 2019-08-02 216
E.3 Corner plot from the MCMC run on the data from the night of 2019-08-07 217
E.4 Corner plot from the MCMC run on the data from the night of 2019-08-16 218

List of tables

1.1 Description of the orbital regimes that many RSOs occupy . . . 3
1.2 Categories and definitions of the main satellite breakup causes as defined by ESA . . 12

2.1 Representative phase functions for idealised shapes using specular and diffuse scattering assumptions . . . 30
2.2 Summary table of the specifications of the STING telescope . . . 38
2.3 Summary table of the specifications of the CLASP telescope . . . 41
2.4 Model Parameters used in the illustrative Cook-Torrance BRDF light curve simulation 56
2.5 Recovered period estimates following the application of the Lomb-Scargle, PDM and FFT methods . . . 64

3.1 Camera-wise along-track and cross-track residual offsets after iterative 3σ clipping . 86

4.1 Summary of objects observed in each of the 17 survey fields . . . 93
4.2 Count of bus configurations for the unique active GEO sample of satellites . . . 94
4.3 Standard deviation measurements for both sequential and simultaneous colour measurements ($B_{RGB} - i$) of short timescale glinting regions . . . 98
4.4 Summary of fitted manoeuvre parameters for the two LDPE-3A events . . . 109
4.5 Interpretation of the fitted manoeuvre parameters . . . 110
4.6 Measurements and standard errors for the near-symmetrical colour changes . . . 113
4.7 Time measurements of the glinting regions for satellites . . . 120
4.8 Colour responses of the measured solar panel offsets in the sample by the type of offset 125

5.1 Table of Ajisai priors corresponding to the observation on the night of 2019-08-07 . 137
5.2 Summary of the 2019 SuperWASP observations of Ajisai . . . 139
5.3 Best estimates of parameters and their uncertainties derived from post burn-in MCMC samples for the 2019 SuperWASP observations . . . 141
5.4 Best estimates of parameters, their uncertainties, and comparison with empirical values with their uncertainties stated in Shakun et al. (2026). . . . 150

6.1 Overview of the Gorizont-6 observations . . . 160
6.2 PDM-derived periods and bootstrap mean periods with 1σ uncertainties for all STING observation epochs and bandpasses (and $B_{RGB} - i$ colour) . . . 162

Acknowledgements

> It is easy to acknowledge, but almost impossible to realise for long, that we are mirrors whose brightness, if we are bright, is wholly derived from the sun that shines upon us.
>
> — C.S. Lewis, *The Four Loves*

First and foremost, I am sincerely grateful to everyone in the Physics Department for their support throughout my PhD. In particular, I thank all my fellow PhD students for the camaraderie, lunchtime discussions, and 5-a-side games that made the stresses of this journey far more enjoyable. Special thanks to my office-mates Jamie and Alastair, and to Sergio, Isobel, Pratishtha for your friendship and the many heart-to-heart conversations that helped me stay balanced and positive in the final stages.

To my hometown friends from Oldham, Connah and Drew: thank you for over a decade of friendship and unwavering support since our school days. Through shared humour and our more recent trips and meetups in Scotland, you've given me laughter, perspective, and a sense of normality I've cherished.

To my close friends Kyeman, Paul, and James: thank you for the social balance, for the many hours talking and gaming, playing pool, and the Jug o' Tea. I'm grateful we've stayed in touch over the years.

I also want to thank my dear friends from Mexico, whom I was fortunate to meet online and who have become an important part of my life. Especially Rey, Rafa, Quiri, Lexy, Jonathan, and everyone in *Los Tejuinos de fornais*. Thank you for your love, encouragement, humour, and for sharing your culture so generously (Spanish lessons included). Gracias por todo.

I am grateful to the legendary 442 Charter Avenue for being a safe haven during this final year. Thank you to Ioannis for your honesty, encouragement, and love of football; to James B for your thoughtful perspective and SDA solidarity; to Morgan for your humour, decisiveness, and authenticity; and to Marcelo for your warmth, generosity, and quiet understanding. You each made me feel welcome, even while carrying your own pressures.

To the Space Domain Awareness group - thank you for your advice, patience, and encouragement whenever I got stuck! And equally a special shoutout to Klaas Wiersema for being supportive and encouraging. And to my supervisor, Paul Chote: thank you for trusting me with such a broad project (and a robotic telescope), and for your expertise and patience. Your guidance helped me

grow not only as a scientist, but as a person, with a deeper appreciation of what rigorous research demands.

To my parents: thank you for your sacrifices, your unwavering love, and for teaching me to treat everyone with kindness and respect. Thank you also for your genuine curiosity about my work and about science more broadly, something that has fuelled my own curiosity since a visit to Jodrell Bank over twenty years ago.

To my siblings, Peter, Andrew, and Fiona: thank you for your love, and for the comfort of our shared roller-coaster of supporting Aston Villa. To my niece, Ada, I hope I can be even a small inspiration to you, as you are to me through your compassion and energy. To my young niece, Scarlett: welcome to the world—you have already brought so much joy, and I hope this work can one day stand as a small example of curiosity and perseverance. Though they can't read this; to all the animals who have been part of my life over this PhD; The awesome foursome (+ Whisper), Lightning, Missy, Sox, Ted, Nella and Bella, your unconditional love has always made being home happier.

To my late sister, Charlotte: though your time was far too brief, your spirit endures and continues to shape who I am. I do everything with the hope of making you proud, and with the feeling that I will always have a guardian angel. I dedicate this thesis, and my PhD journey, to your memory.

I dedicate this thesis to the memory of my sister, Charlotte

Declaration

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The work presented in this thesis, including data generation and analysis, was carried out by myself unless otherwise explicitly stated, under the supervision of Dr Paul Chote between October 2022 and May 2026. Parts of this thesis have been published by myself. Material presented in Chapter 3 and Chapter 4 is published in Airey et al. (2025a). Material presented in Chapter 5 is published in Airey et al. (2025b). Material presented in Chapter 6 is published in Airey et al. (2026). For §1.1, information from ESA DISCOS (Database and Information System Characterising Objects in Space) was used as a single-source reference for launch information, object registration details, launch vehicle descriptions, and spacecraft information for all trackable, unclassified objects. I acknowledge ESA's efforts in maintaining and operating this database and its associated APIs.

The observation scheduling procedures and data reduction pipelines used to extract the light curves analysed in this thesis were written and developed by Paul Chote. The polarimetry analysis described in Chapter 6 formed part of a larger proposal for observing time on the Liverpool Telescope (LT), drafted by Klaas Wiersema (University of Hertfordshire), whom also analysed and reduced the polarimetry data of said chapter.

Elements of the work presented in this thesis have also formed the basis of several research talks and poster presentations, including "Constraining the Attitudes of RSOs from Photometric Measurements", presented as a conference talk at the Operating in the Future Electromagnetic Environment Symposium in Newport, United Kingdom (Nov 2025), and "A comprehensive survey of the GEO-belt using simultaneous four-colour observations with STING", presented as a poster at the National Astronomy Meeting (NAM) at Durham University, United Kingdom (Jul 2025).

Abstract

The rapid expansion of human activity in space has led to sustained growth in the population of active satellites, defunct spacecraft, rocket bodies and debris throughout the near-Earth environment. This growth presents distinct challenges in both Low Earth Orbit (LEO) and Geostationary Earth Orbit (GEO), where the ability to characterise resident space objects (RSOs) is central to space domain awareness and the long-term sustainability of orbital operations. While radar and resolved imaging can be effective for some LEO applications, objects in GEO are too distant for resolved ground-based imaging and are less amenable to radar detection. High-precision optical photometry therefore provides one of the most practical routes for studying their behaviour, geometry and attitude.

This thesis investigates how high-precision photometry can be used to characterise RSOs across both GEO and LEO regimes. The SuperWASP-N facility was upgraded to STING, a dedicated four-colour wide-field instrument for space domain awareness, and its performance was characterised in terms of photometric precision, astrometric accuracy and practical visibility constraints. STING was then used to carry out the first large systematic colour survey of the GEO belt, producing full-night simultaneous multi-colour light curves over a five-week campaign. These observations show that simultaneous colour measurements reduce errors introduced by sequential filter cycling during rapid brightness changes, and that colour evolution, glints and solar-panel offsets provide useful discriminants between satellite classes and, in some cases, between individual objects. The survey further shows that some bus families exhibit distinctive colour behaviour, with tentative evidence for colour evolution within parts of the Eurostar-3000 population.

At LEO, the thesis presents a high-cadence photometric study of the geodetic satellite Ajisai. Using four SuperWASP observations from August 2019, TLE-derived geometry and an MCMC-based fitting framework were used to transform observed glints into body-fixed brightness maps, allowing the spin-state and surface reflectivity pattern to be recovered without requiring a detailed BRDF model or complete a priori shape knowledge. The resulting solutions are consistent with empirical predictions and demonstrate that rapid optical photometry can constrain the attitude of fast-spinning objects using modest ground-based instrumentation. This was expanded upon with observations of the satellite with the STING instrument.

The thesis also examines the inactive GEO satellite Gorizont-6 through coordinated multi-colour photometric and polarimetric observations. Photometry across six epochs revealed spin-down over nine months consistent with YORP-driven rotational evolution, while simultaneous polarimetry broke period ambiguities and linked light curve features to specific spacecraft surfaces. Taken together, the results demonstrate that high-precision photometry, especially when combined with simultaneous colour and targeted polarimetry, provides a powerful, scalable nd physically informative route to the characterisation of non-cooperative space objects.

List of abbreviations

AAVSO	American Association of Variable Star Observers
ADR	Active Debris Removal
APASS	AAVSO Photometric All-Sky Survey
BIPM	International Bureau of Weights and Measures - '*Bureau International des Poids et Mesures*'
BPN	Bias Precession Nutation
BRDF	Bidirectional Reflectance Distribution Function
CCD	Charged Coupled Device
CCR	Corner Cube Reflector
CMOS	Complimentary Metal Oxide Semi-conductor
COM	Centre of Mass
CRS	Celestial Reference Sphere
DISCOS	Database and Information System Characterising Objects in Space
ECEF	Earth Centred Earth Fixed
ECI	Earth Centred Inertial
ESA	European Space Agency
FFT	Fast Fourier Transform
FOV	Field Of View
FWHM	Full Width Half Maximum
GCRS	Geocentric Celestial Reference Sphere
GEO	Geostationary Earth Orbit
GEODSS	Ground-based Electro-Optical Deep Space Surveillance

GNSS	Global Navigation Satellite System
GP	General Perturbations
GPS	Global Positioning System
GPST	Global Positioning System Time
GSO	Geosynchronous Earth Orbit
HA	Hour Angle
HCS	Horizontal Coordinate System
IADC	Inter-Agency Space Debris Coordination Committee
IAU	International Astronomical Union
ICRF	International Celestial Reference Frame
ICRS	International Celestial Reference System
IERS	International Earth Rotation and Reference Systems Service
IRP	International Reference Pole
ITRF	International Terrestial Reference Frame
JD	Julian Day
LEO	Low Earth Orbit
LST	Local Sidereal Time
LT	Liverpool Telescope
LVLH	Local Vertical Local Horizontal
MCMC	Markov Chain Monte Carlo
MEO	Medium Earth Orbit
MJD	Modified Julian Day
MOPTOP	Multi-colour OPTimised Optical Polarimeter
NASDA	National Space Development Agency of Japan
NGA	National Geospatial-Intelligence Agency
NIR	Near Infrared
NORAD	North American Aerospace Defense Command
NPA	Non Principal Axis

NSO	Navigation Satellite Orbit
ORM	Roque de los Muchachos Observatory -'*Observatorio del Roque de los Muchachos* '
OSC	Orbital Science Corporation
PAB	Phase Angle Bisector
PDM	Phase Dispersion Minimisation
QE	Quantum Efficiency
RA	Right Ascension
RAAN	Right Ascension of Ascending Node
RASA	Rowe Ackermann Schmidt Astrograph
RF	Radio Frequency
RMS	Root Mean Square
RPM	Revolutions Per Minute
RPO	Rendezvous Proximity Operation
RSO	Resident Space Object
sCMOS	Scientific Complimentary Metal Oxide Semi-conductor
SDA	Space Domain Awareness
SED	Spectral Energy Density
SGP	Simplified General Perturbations
SLR	Satellite Laser Ranging
SRP	Solar Radiation Pressure
SSA	Space Situational Awareness
SSN	Space Surveillance Network
SSO	Sun Synchronous Orbit
SuperWASP	Super Wide Angle Search for Planets
SWIR	Short Wave Infrared
TAI	International Atomic Time - '*Temps Atomique International*'
TCS	Telescope Control System
TETE	True Equator and True Equinox of date

TLE	Two Line Element
TT	Terrestial Time
USNO	United States Naval Observatory
UT1	Universal Time
UTC	Coordinated Universal Time
VLBI	Very Long Baseline Interferometry
WASP	Wide Angle Search for Planets
WCS	World Coordinate System
WGS	World Geodetic System
YORP	Yarkovsky–O'Keefe–Radzievskii–Paddack

1 Introduction

"Change will preserve us. It is the lifeblood of the isles. It will move mountains! It will mount movements!"

Sheogorath, The Elder Scrolls IV: Oblivion, Bethesda Studios

The first artificial satellite, Sputnik-1, was launched in October 1957 by the Soviet Union. This was the beginning of the Space Race, which in turn was the first era of the 'space age'. The launch of Sputnik-1 and many others in this early phase of the 'space age' demonstrated the feasibility of placing objects into orbit around the Earth. This era resulted in much innovation in science and technology, particularly rocket engineering and material science. Artificial satellites have become intertwined with almost every aspect of modern society, from large-scale communication and navigation networks to weather forecasting and Earth monitoring. Artificial satellites have also afforded us the opportunity to perform high-level astronomy through taking observations with various space-based observatories.

However, the orbital regimes around Earth that are vital for these networks have become increasingly congested, primarily due to the proliferation of satellite constellations and the growing accumulation of space debris, which together increase the risk of collisions. These issues are vitally important to consider and to mitigate, and this is the overarching goal of SDA (space domain awareness). Greater characterisation of resident space objects (RSOs) in the near-Earth environment is therefore required to support the sustainable use of space, as improved knowledge of these objects can aid their identification, improve predictions of their behaviour, and support collision avoidance and debris mitigation efforts.

This thesis explores how high-precision optical photometry can be used toward improving the characterisation effort for various types of RSO and how this can contribute to improved SDA.

1.1 Orbital Regimes

Any satellite in the near Earth environment will be influenced primarily by the gravitational force exerted by the central mass of the Earth, such that with a given tangential velocity, the satellite will be expected to follow a curved trajectory around the Earth. Given that the mass of a satellite is negligible in comparison to Earth mass and assuming the Earth to be spherically symmetric, the orbital motion of a satellite can be described by Newton's law for an unperturbed two body problem

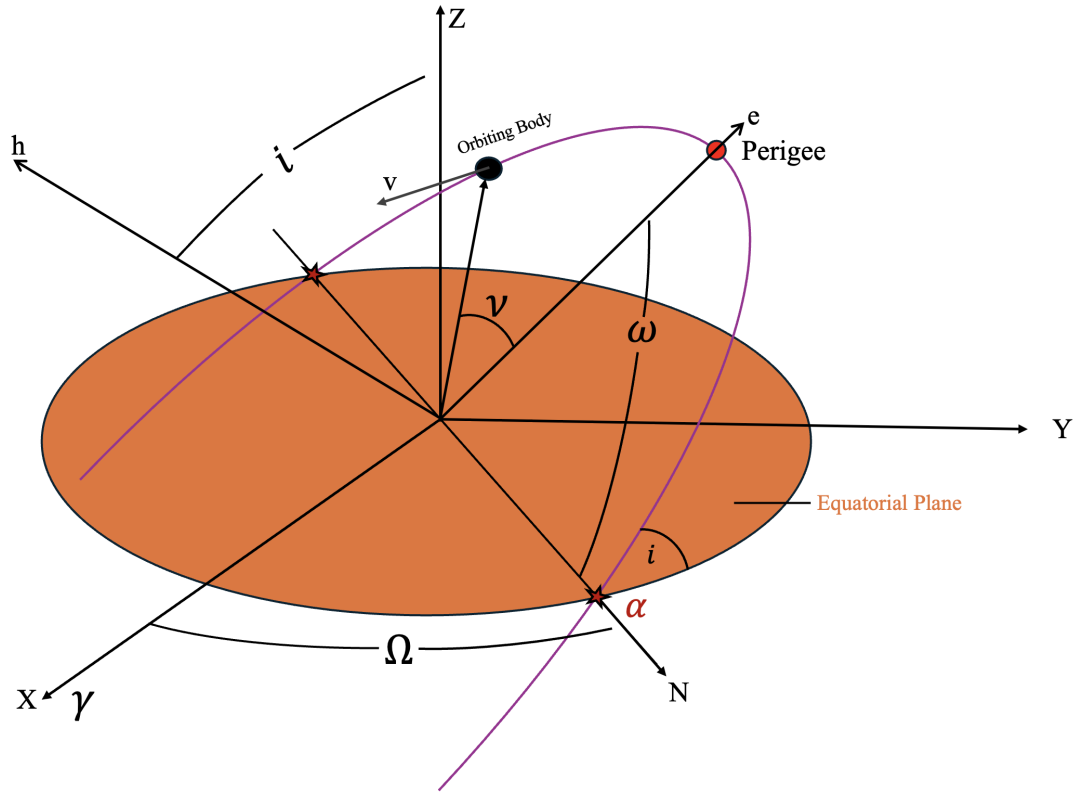

FIGURE 1.1: Illustration of the Keplerian orbital elements within a geocentric equatorial reference frame. The equatorial plane (in orange) is the plane of reference bound by XY and the vernal equinox (γ) serves as the X direction in this frame. An orbiting body is shown traversing its orbit (purple) with a given velocity, v. The eccentricity and specific angular momentum vectors, e and h respectively are shown alongside the four angular elements: RAAN (Ω), inclination (i), argument of perigee (ω) and true anomaly (ν). All angles are measured counter clockwise. The ascending node is marked by the red star symbol and denoted by α.

(Montenbruck and Eberhard, 2000),

$$\ddot{\mathbf{r}} = -\frac{GM_{\oplus}}{r^3}\vec{\mathbf{r}}, \tag{1.1}$$

where $\ddot{\mathbf{r}}$ is the vector describing the acceleration of the body, $GM_{\oplus} \approx 3.986 \times 10^{14} m^3 s^{-2}$, the fraction $-\vec{\mathbf{r}}/r$ describes the unit vector pointing from the satellite to the Earth's centre and r is the scalar distance between the two bodies.

Since the gravitational force experienced by the satellite points towards the Earth centre, this consequently means that the orbital motion is confined to a fixed plane since there is no resultant acceleration perpendicular to the orbital plane. These dynamics in turn give rise to Kepler's laws and the confined orbital planes when intersecting a cone result in distinct trajectories such as a circle, ellipse or parabola.

The orbit of a satellite can be typically described by six Keplerian orbital elements (e, a, i, Ω, ω, ν), illustrated in Figure 1.1. The eccentricity (e, 0 for a perfectly circular orbit) describes the deviation of an elliptical orbit away from a perfect circle, with values < 1 governing an elliptical trajectory. The semi-major axis (a) is calculated as half of the distance between apogee and perigee, the furthest and nearest approach from the primary body respectively. Both these elements govern the shape and extent of the orbit.

The right ascension of ascending node (RAAN, Ω) describes the angle between the ascending

node (where orbital plane intersects reference plane) and the given reference frame's reference direction, Υ. The inclination (i) describes the vertical tilt of the orbital plane relative to the reference plane, measured at the ascending node. The argument of perigee (ω) describes the orientation of the orbit within the orbital plane and is measured as the angle between the ascending node and orbital perigee.

The RAAN, inclination and argument of perigee are elements which describe the orientation of the plane of the orbit, and the orientation of the orbit within the defined plane. These three elements combined can be described as Euler angles which define the orientation of the orbit relative to a reference coordinate system (Capderou, 2005).

The final element in this set is the true anomaly, ν, which describes the angular displacement of the orbiting body at a particular epoch time, measured relative to the orbital perigee.

The near-Earth environment can be sub-divided into various orbital regimes which have become further populated with spacecraft over time. These regimes possess distinct characteristics that operators have exploited to satisfy different mission requirements. Such regimes are typically defined in terms of the inclination and the perigee and apogee (h_p and h_a), the orbital heights of the nearest and furthest points of an orbit above the Earth's surface respectively. Figure 1.2 provides a scaled representation of the key orbital regimes relative to Earth, while Table 1.1 summarises the orbital parameters used to define their respective boundaries.

Orbital Regime	Description	Defining Limits
LEO	Low Earth Orbit	$h_p,\ h_a \in [0, 2000]$ km
NSO	Navigation Satellites Orbit	$i \in [50, 70]^\circ$ $h_p,\ h_a \in [18100, 24300]$ km
MEO	Medium Earth Orbit	$h_p,\ h_a \in [2000, 31570]$ km
GSO	Geosynchronous Earth Orbit	$i \in [0, 25]^\circ$ $h_p,\ h_a \in [35586, 35986]$ km
GEO	Geostationary Earth Orbit	$i = 0^\circ$ $h_p = h_a = 35786$ km

Table 1.1: Description of the orbital regimes that many RSOs occupy, adapted from (ESA Space Debris Office, 2026), where h_p and h_a are the perigee and apogee heights respectively, and i is the orbital inclination. Heights are given in km and angles in degrees.

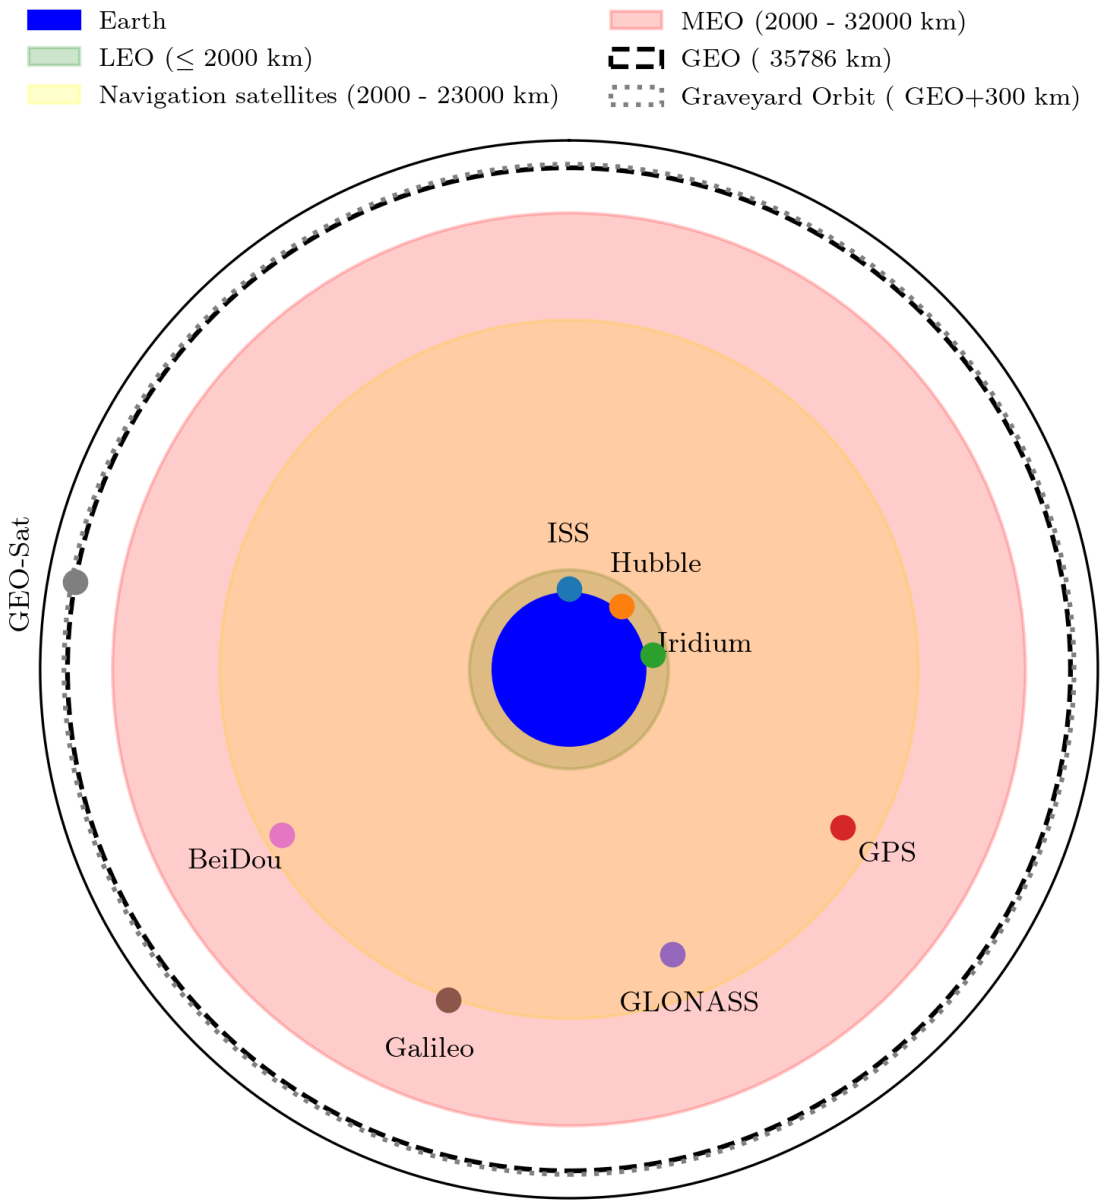

Figure 1.2: A diagram representing the most populous orbital regimes with respect to the Earth (blue). Example satellites are annotated with respect to the orbital regime they operate within.

Low Earth orbit (LEO) is the orbital regime closest to Earth and the most densely populated of all regimes. Typically this orbit is defined by orbital heights (measured from the Earth's surface) between approximately 200 and 2000 km with orbital eccentricities and periods of less than 0.25 and 128 minutes respectively (Emery and Camps, 2017). The usefulness of this regime is primarily derived from the comparatively low Δv required to place a satellite into this orbit. Δv represents the total change in velocity that a launch vehicle or spacecraft must provide to reach and maintain a particular orbit and is therefore a useful measure of the propulsive cost of a mission. Launch, orbit raising, plane changes, station-keeping and de-orbiting each carry a Δv price as part of a total Δv budget. Lower orbital heights generally require less Δv to reach than higher-energy orbital regimes, reducing the required propellant mass and allowing a greater proportion of the launch vehicle's capacity to be allocated to payload. Reaching LEO from the surface requires a Δv of $\sim$9.4 $\mathrm{km/s}$. Comparatively, to reach geostationary Earth orbit (GEO), a further $\sim$3.9 $\mathrm{km/s}$ is required (Orbital Radar, 2026).

Additionally, LEO is a useful regime due to the smaller communications latency between ground and satellite, making this regime particularly useful for communication purposes. However, due to the nature of a lower orbital height i.e., the Earth having a larger angular extent, there is a narrower coverage from a fixed narrow beam size, therefore constellations of satellites are needed to provide continuous coverage for a particular region of the Earth.

LEO can be sub-divided further into other useful orbits such as polar orbits and Sun-synchronous orbits (SSO), both of which are highly inclined with inclinations around approximately 98° and with orbital heights around 780 km resulting in retrograde orbits with orbital periods of around 90-100 minutes. These orbits in particular are very useful for imaging/monitoring of the Earth's surface and atmosphere due to the satellite always being illuminated by the Sun and the smaller illumination angle from the Sun creating good lighting conditions for imaging. Polar Orbits are ones in which a satellite will pass over the poles ($\pm 10°$) as the Earth rotates underneath, thus enabling the entire Earth to be mapped over the course of a full rotation. SSOs are a type of near-polar orbit, where the satellite in this orbit will pass over the same spot at the same mean solar time every pass, which is to say that the illumination angle remains approximately the same every time the satellite passes over a given target. However, an SSO is not achievable without an induced precession of about 1 degree eastwards per day with respect to the celestial reference sphere (CRS) in order to account for the Earth's movement around the Sun (Rosengren, 1992). The precession is induced via tuning the inclination to the height of the orbit such that Earth's equatorial bulge, which perturbs inclined orbits, causes the orbital plane of the spacecraft to precess at the desired rate stated prior.

Satellites in LEO are subject to a combination of perturbative forces that affect their orbits. One of the most significant influences is Earth's non-spherical gravity resulting from the planet's aforementioned equatorial bulge and uneven mass distribution. The departure of the Earth's gravity field from spherical symmetry is described by an infinite spherical-harmonic series, whose latitude-dependent terms carry the zonal coefficients J_k. The J_2 term, which represents the Earth's oblateness, is more than two orders of magnitude larger than any other coefficient and is therefore the dominant non-spherical gravitational perturbation acting on Earth-orbiting satellites. This term explains secular changes like nodal regression and perigee rotation (Vess and Starin, 2003), which have the effects of moving the RAAN at a constant rate and precessing the argument of perigee

respectively.

In addition, atmospheric drag in the upper atmosphere exerts a drag force that reduces a satellite's orbital energy, leading to a gradual orbital decay; this effect is especially pronounced during the solar maximum when the atmosphere expands (Lewis et al., 2011). Moreover, solar radiation pressure (SRP) arising from the momentum transfer of photons impacting and reflecting off a satellite's surfaces imparts small yet cumulative changes in orbital elements, particularly for satellites with high area-to-mass ratios (Wang et al., 2023). Collectively, these forces require regular station-keeping manoeuvres to maintain the desired orbit and ensure mission success.

The next orbital regime outwards is medium Earth orbit (MEO) with satellites within this regime typically reaching orbital heights ranging from between 2000 to 31,570 km. The typical use of this orbit is reserved for navigational satellite constellations such as GPS (global positioning system), GLONASS (global navigation satellite system), Beidou and Galileo. These satellites, with orbital heights of around 20,000 km move along prograde semi-synchronous orbits, such that they pass over the same two points on the equator every day (Riebeek, 2009). MEO satellites, similar to those in LEO, are perturbed by SRP, which is the most dominant non-gravitational perturbing force (Bury et al., 2020). In addition to this, MEO satellites must be designed with shielding in mind as this orbital regime contains two Van Allen radiation belts. These belts are regions of highly energised charged particles with energies sufficient enough to penetrate most materials and damage internal electronics (Reeves et al., 2020). More recently satellites have been deployed in MEO for commercial communication purposes (SES S.A., 2026).

Stationary Earth Orbits were first discussed in the context of a useful orbit for a network of communication satellites (Hoyt, 1964). The concept of GEO was first coined by Arthur Clarke in the article 'Extra-Terrestrial Relays: Can Rocket Stations Give World-wide Radio Coverage?' (Clark, 1966) in which this orbit was discussed as a useful vantage point for communication satellites. Geostationary Earth orbit (GEO) is a circular orbit (e=0) at a height of approximately 35,786 km above the Earth's equator ($\imath = 0$). Satellites within this regime are at rest with respect to the Earth's rotation given their orbital period is equal to a sidereal day ($23\text{h}\,56\text{m}\,04\text{s}$) and thus appear stationary to an observer on the ground. In addition to the large footprint i.e., due to the Earth having a smaller angular extent, there is larger coverage from a fixed narrow beam size, this makes them ideal for telecommunications, meteorology, data relaying and military activities. However, this comes at the expense of a greater signal latency due to the greater line of sight distance than previously discussed regimes. Specifically, the minimum latency to GEO and LEO is $\sim$ 600 ms and $20 - 40$ ms respectively (Reliasat, 2024).

Objects in GEO are perturbed by the gravitational influence of the Moon, the Sun, and the oblateness of the Earth (J_2 effect (Sehnal, 1960)). Without station-keeping manoeuvres, these perturbations cause the inclination of GEO orbits to vary between approximately $\pm 15°$ over a period of $\sim$53 years (Anderson et al., 2015). Station-keeping manoeuvres are therefore required throughout the operational lifetime of a GEO satellite and are typically performed in the North-South and East-West directions. North-South manoeuvres counteract lunar-solar perturbations on the orbital inclination, while East-West manoeuvres control the satellite's longitude, which is affected by drift due to the Earth's oblateness and the non-zero eccentricity induced by SRP effects (Guelman, 2014). At the end of a GEO satellite's operational lifetime, direct de-orbiting is generally impractical

because it requires a delta-v of approximately 1500 m/s (Bolonkin, 2014). This is equivalent to many years of station-keeping fuel expenditure and is therefore not typically feasible for satellites nearing the end of their mission. Instead, the standard disposal strategy is to manoeuvre the satellite into a graveyard orbit, located approximately 300 km above the GEO region. This requires a much smaller delta-v of approximately 11 m/s (Bolonkin, 2014), comparable to only a few months of station-keeping manoeuvres.

Once placed in the graveyard orbit, GEO satellites are typically passivated, this involves removing or depleting remaining sources of internal energy, such as residual propellant and batteries, in order to reduce the risk of fragmentation or explosions. Such events have previously been caused by internally stored energy in defunct spacecraft (Bonnal, 2007), an example of this in particular is Intelsat 29E, which suffered a fuel leak and entered a tumbling state shortly after (Henry, 2019).

1.1.1 The GEO-belt

The GEO belt refers to the region of near-Earth space populated by objects at geosynchronous orbital heights, encompassing not only the geostationary satellites described previously but also the inclined and drifting objects that result from the perturbations and disposal practices discussed above. Unlike the lower regimes, atmospheric drag is negligible at these heights, so objects do not decay naturally and remain on orbit indefinitely. Over decades of continuous use, the GEO belt has therefore accumulated a complex population of operational satellites, defunct spacecraft, rocket bodies, and debris fragments (ESA Space Debris Office, 2026). This population is not distributed uniformly in longitude, since uncontrolled objects librate about the stable equilibrium points at $\approx 75°$E and $\approx 105°$W, produced by the longitude-dependent acceleration arising from the ellipticity of the Earth's equator (Soop, 1994). The resulting clustering is of particular interest for survey observations.

Optical observations have become a primary means of surveying and characterising the GEO environment, offering sensitivity to small objects and enabling photometric studies of their physical and dynamical properties. At geosynchronous ranges the returned radar signal scales as R^{-4} whereas reflected sunlight scales as R^{-2}, which makes high-precision optical photometry particularly well suited to this region. The observing geometry is, however, latitude dependent. Although geostationary objects lie close to the Earth's equatorial plane in a geocentric sense, their apparent topocentric positions depend on the observer's latitude. Due to parallax, the GEO belt appears south of the celestial equator for northern hemisphere observers; from La Palma it lies at approximately $-4°$ to $-5°$ declination, while from higher-latitude locations such as the UK it appears further south. Conversely, from southern hemisphere observing sites, the GEO belt appears north of the celestial equator.

1.2 Resident Space Objects

As discussed in §1.1, the orbital regimes surrounding the Earth are populated by a variety of different RSOs with various objectives, designs and properties. Over the decades, the number of RSOs has grown at a considerable rate, with particular focus on communication and military networks. This

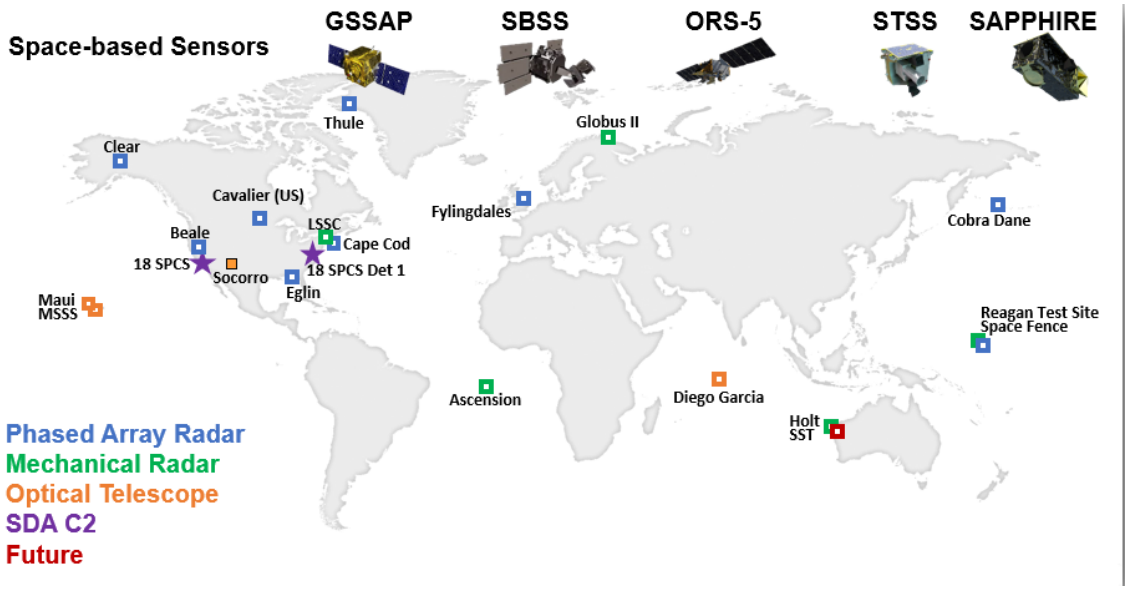

Figure 1.3: A diagram representing the locations of the facilities used within the Space Surveillance Network (SSN). It is a combination of ground-based optical sensors, radar arrays and space-based sensors taken from Force (2022).

growth is not merely a continuation of historical trends. The specific cost of access to orbit has fallen substantially (Jones, 2018), driven by partially reusable launch vehicles and high-cadence ride-share provision, while satellite manufacture has shifted from bespoke, long-lifetime platforms toward mass-produced spacecraft with comparatively short design lives, replenished on a continuing basis over the operational lifetime of a constellation (Sturza and Saura Carretero, 2021). Combined with commercial demand for global broadband coverage which, unlike GEO, requires hundreds to thousands of satellites to deliver continuous low-latency service, these factors have made large constellations economically viable.

1.2.1 Classification and Identification

The United States Space Surveillance Network (SSN) detects, tracks, catalogues and identifies artificial objects orbiting Earth. This includes active and defunct satellites, rocket bodies and fragmentation debris.

Figure 1.3 displays the entirety of the space surveillance network. The Space Surveillance Network is made up of electro-optical, passive radio frequency (RF) and radar sensors, which are used to maintain a consistent watch on the near-Earth environment.

It is worth clarifying that the SSN is not, however, the sole source of tracking data for the near-Earth environment. Within Europe, the EU Space Surveillance and Tracking (EU SST) Partnership networks national radar and optical assets across member states, complemented by commercial sensor provision, to deliver collision avoidance, re-entry and fragmentation analysis services (EU SST Partnership). Commercial providers also operate independent networks in their own right, notably LeoLabs (phased-array radar, primarily LEO) and ExoAnalytic Solutions (a globally distributed optical telescope network, primarily GEO). In the United States, responsibility for civil conjunction-assessment services has transitioned from the Department of Defense to the

Department of Commerce through the Traffic Coordination System for Space (TraCSS), which draws on commercially procured observations alongside government data (NOAA Office of Space Commerce). Tracking and characterisation are therefore increasingly distributed activities, drawing on government, allied and commercial sensors of widely varying aperture and design. This is directly relevant to the work presented within this thesis, since it establishes high-precision optical characterisation from non-governmental/commercial facilities.

Three of the sites denoted in Figure 1.3 form the Ground-based Electro-Optical Deep Space Surveillance (GEODSS), Socorro, Maui and Diego Garcia respectively. The goal of observing from these sites is to track artificial objects beyond GEO. Each site has three 40 inch telescopes with ~ two degree FOVs. As with any ground-based optical sensor, the variability in weather directly impacts upon the effectiveness of the system. Each GEODSS site tracks ~3,000 objects per night, the majority of the deep-space catalogue being inaccessible to radar at these orbital heights. Absolute catalogue figures quoted in this section reflect the state of the public catalogue at the time of writing and are subject to continuous revision as sensor capability and launch rates evolve.

Typically the smallest debris fragments reliably observable by the SSN are objects with diameters of 10 cm in the LEO regime (mostly by radar) (Matney, 2016) and between 50-100 cm in the geosynchronous domain (Manis et al., 2023). Each identified and tracked object has an associated 5 digit NORAD identification number. According to the latest boxscore[1], the SSN has tracked and catalogued over 63,000 objects, with ~ 30,000 of these remaining in orbit at the time of writing.

A general perturbations (GP) theory is used to maintain the resultant catalogues following the observations made by the SSN as it provides a general analytical solution to the satellite equations of motion. The Keplerian orbital elements and partial derivatives are described by series expansions based upon initial conditions set by the differential equations. The GP theory can be simplified by a series of assumptions such as; truncating the Earth's gravitational potential to a few zonal harmonics, modelling the atmosphere as a static and exponentially decaying spherical field and by only partially modelling the complexity of third body influences. These models are based on developments made by Lane and Cranford in the sixties (see Lane, 1965; Lane and Cranford, 1969) and onwards as they developed models used to predict satellite positions based on a limited set of data elements. This eventually led to SGP (Simplified General Perturbations) models, which have become the standard model used to propagate satellite orbits. Lane's models were widely adopted by the military and NASA and led to the creation of the TLE format, which is specific to the SGP models.

The TLE format gives the mean Keplerian orbital elements which describe a satellite orbit. Each TLE consists of two lines of 69 characters; the format and key elements are annotated in Figure 1.4. The typical accuracy of a TLE at epoch from the SGP models is larger than ~1 km in LEO, ~3 km in MEO (Ly et al., 2020). In GEO, Früh and Schildknecht (2012) quote on average, 25 km and 10 km residuals in along-track and cross-track directions respectively. The rate at which this positional accuracy decays depends on the largest perturbation the SGP theory fails to capture. In LEO, this is atmospheric drag, compressed into the single constant B* term: residual errors in the modelled energy loss appear as a semi-major axis error and, accumulate along-track at typically

[1]The SATCAT Boxscore maintains a record of objects tracked and catalogued by the US SSN, and is updated daily via the SpaceTrack website: `https://www.space-track.org/#boxscore` (accessed April 2025)

1 km day^{-1} (Levit and Marshall, 2011), faster for high area-to-mass-ratio objects and by an order of magnitude during geomagnetic storms (Parker and Linares, 2024). In GEO, drag is negligible and the dominant lunar-solar and radiation-pressure perturbations are adequately modelled by SDP4 (Simplified Deep-space Perturbations), so element sets show little secular growth and are limited instead by un-modelled station-keeping manoeuvres. LEO element sets therefore require refreshing at least daily, whereas GEO sets remain useful for several days. In the particular case of the prototype Bluewalker-3 constellation satellite, the decay rate of the TLE was measured to be 0.4 ± 0.2 arcmin h^{-1} (Nandakumar et al., 2023). TLE mean orbital elements are publicly available via the Space-track website.[2]

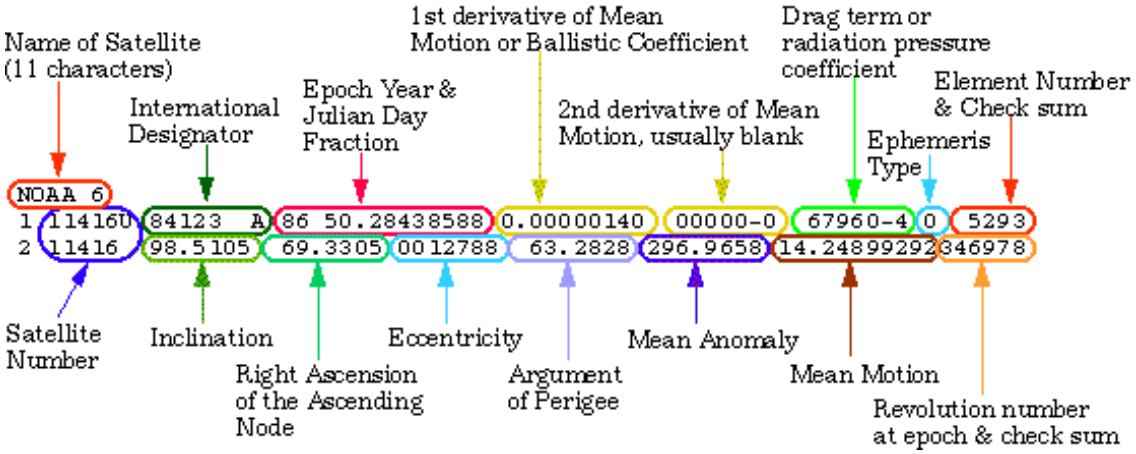

Figure 1.4: Example TLE of a satellite with key orbital elements and segments annotated. This example is taken from Mirghani et al. (2015).

[2] A publicly-available version of the catalogue of objects detected orbiting the Earth by the SSN is accessible via https://www.space-track.org/

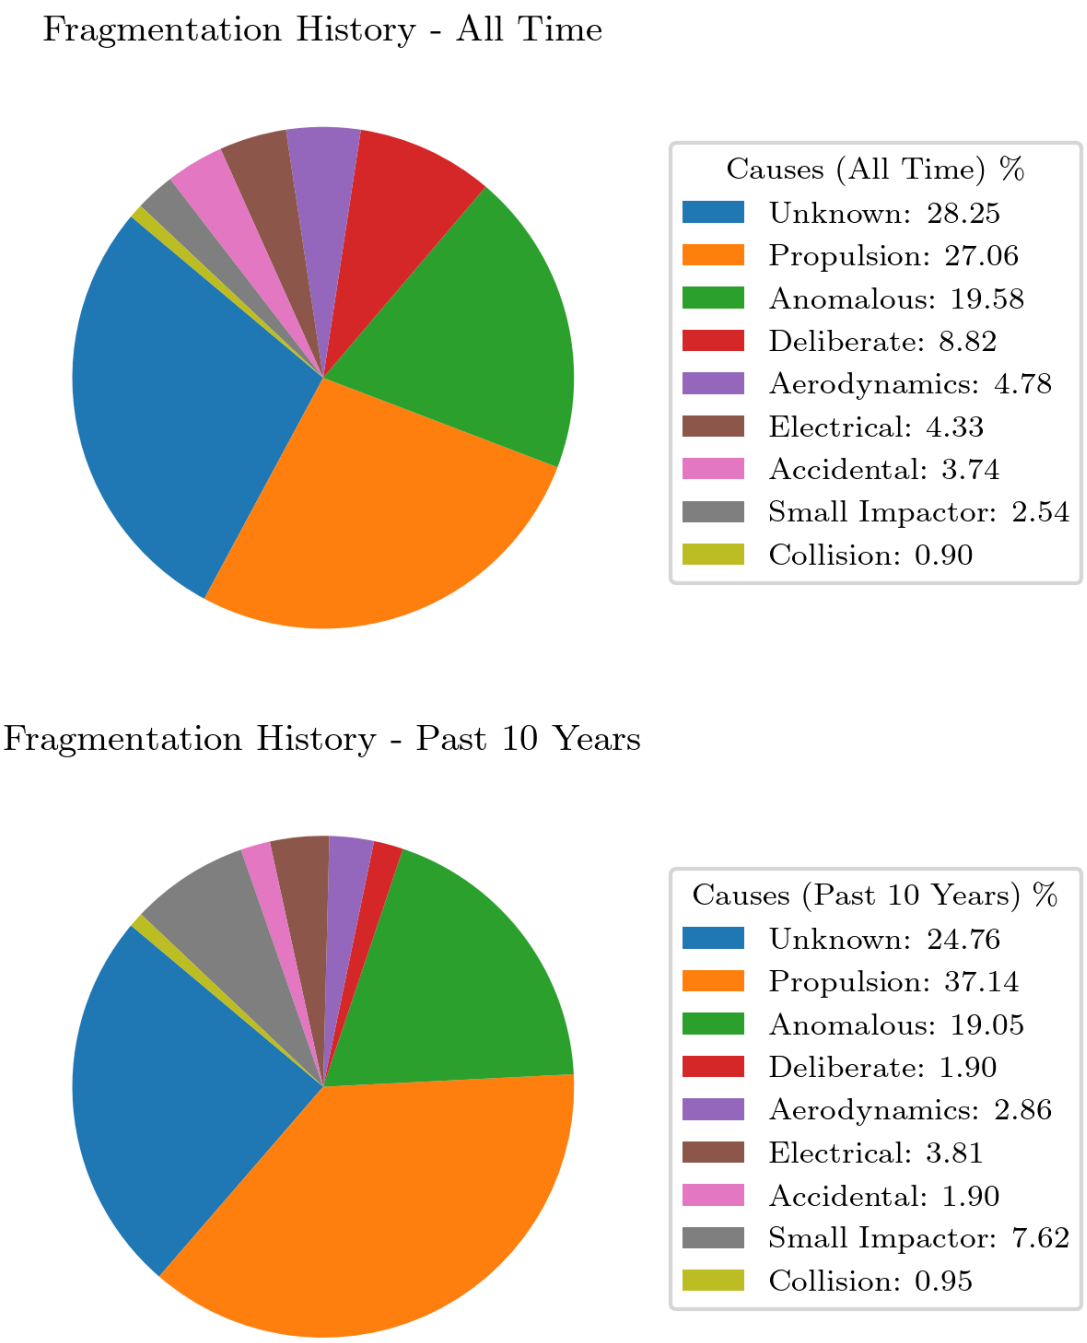

Figure 1.5: Fractional representation of fragmentation history. Upper pie chart displays the fragmentation history since the start of the space-age. Lower pie chart displays the fragmentation history over the past decade. Data adapted from European Space Agency (2025a).

Category	Definition
Accidental	Subsystems that showed design flaws ultimately leading to breakups in some cases.
Aerodynamics	A breakup most often caused by an overpressure due to atmospheric drag.
Collision	A breakup associated from the resultant collision between objects. A sub-class are so-called small impactors.
Small Impactor	Caused by a collision, but without explicit evidence for an impactor. Changes in the angular momentum, attitude and subsystem failures are, however, indirect indications of an impact.
Deliberate	All intentional breakup events i.e., Anti-satellite tests, detonated disposal etc.
Electrical	Most of the events in this category occurred due to an overcharging and subsequent explosion of batteries.
Propulsion	Stored energy for non-passivated propulsion-related subsystems might lead to an explosion, for example due to thermal stress.
Anomalous	Defined as the unplanned separation, usually at low velocity, of one or more detectable objects from a satellite that remains essentially intact. This may include debris shedding due to material deterioration, which includes insulation material or solar panels all of which have been observed from ground in the past.
Unknown	Assigned whenever there is lacking evidence to support a more specific classification.

Table 1.2: Categories and definitions of the main satellite breakup causes as defined by ESA (European Space Agency (ESA), ESOC).

1.2.2 Challenges and Trends

The rapid expansion of space activities over recent decades has resulted in an unprecedented growth in the population of artificial objects in Earth orbit, particularly within Low Earth Orbit (LEO) and Geostationary Earth Orbit (GEO). These orbital regimes are of high strategic, scientific, and commercial importance, yet their intensive utilisation has introduced a range of technical, operational, and environmental challenges that threaten their long-term sustainability (European Space Agency, 2023a; Inter-Agency Space Debris Coordination Committee, 2007).

1.2.2.1 Space Debris: The Current Picture

Typically, the challenge at the forefront of SDA is that of space-debris. Over the decades, the number of debris objects has grown immensely. One mode of growth of space-debris can be attributed to fragmentation events yielding fragmentation debris, these are events where a number of debris fragments are produced either due to a collision between two objects or an internal breakup of a solitary object. The European Space Agency (ESA) classifies fragmentation events into 9 categories; unknown, propulsion, anomalous, deliberate, aerodynamics, electrical, accidental, small impactor and collision (see Table 1.2).

There have been some notable examples of this occurring (see e.g., Anz-Meador et al., 2018; Cerutti-Maori et al., 2025):

- 1960–1980: Ablestar, OV2-1, Nimbus-4 and Delta-2910
- 1980–2000: Parus, Solwind, SPTOT-1, STEP-2 and CBERS-1
- 2000–2015: TES, Fengyun-1C, KOSMOS-2421, Iridium-33, KOSMOS-2251, NOAA-16
- 2015–2024: Fregat Tank, KOSMOS-1408, Long March-6A x2, KOSMOS-2499, Intelsat-29e, Intelsat-33e

Over the decades, the evolution of fragmentation events has introduced a number of trends relevant to the understanding of long-term space environment sustainability. Historically, fragmentation events of unknown origin have constituted the largest proportion, accounting for approximately 28.25% of all events recorded. However, in the past decade, propulsion-related fragmentation has become increasingly dominant (37.14%). This may be indicative of the greater deployment of spacecraft in the past decade. The relative stability observed in categories such as anomalous, electrical, and collision-related events further emphasises the emerging dominance of propulsion and impactor-related fragmentation in the current decade. A breakdown of these categories is shown in Figure 1.5. As of April 2025, the total number of fragmentation events in history exceeds 650 (European Space Agency, 2025a).

DISCOSweb[3] (Database and Information System Characterising Objects in Space) functions as a centralised reference hub, providing launch data, object registration records, launch vehicle details, and spacecraft characteristics—such as dimensions, mass, configuration, mission purpose, and ownership—for more than 40,000 trackable, unclassified space objects, whose orbital data are publicly releasable rather than withheld for national security reasons. Figure 1.6 displays the distribution of different object classes within the database, it is evident that the number of trackable fragmentation debris (~ 38%) exceeds the number of payloads (~ 30%) (both active and inactive).

If it is considered that not all debris fragments are tracked due to technological limitations, the consequences of an ever growing debris field in combination with frequent launches will give rise to many more challenges across various orbital regimes. In particular, according to the ESA MASTER[4] (Meteoroid And Space debris Terrestrial Environment Reference) statistical models, the orbital environment is estimated to contain more than 54,000 objects larger than 10 cm, around 1.2 million debris fragments between 1 cm and 10 cm, and approximately 140 million particles ranging from 1 mm to 1 cm (European Space Agency, 2025b). Although the orbital volume involved is vast, the hazard is poorly described by mean number density. Collisions in LEO occur at relative velocities averaging ~10 km s^{-1} (National Research Council, 2011), at which a fragment need only exceed ~1 cm to defeat the shielding carried by even heavily protected spacecraft, while objects smaller than ~10 cm are generally too small to be tracked and avoided (NASA Orbital Debris Program Office). The hazardous population is therefore dominated by objects that can neither be manoeuvred around nor shielded against, and a single catastrophic fragmentation permanently raises the background flux for every other user of that regime; the resulting risk of a self-sustaining collisional cascade is discussed more prominently in the next subsection.

[3] https://discosweb.esoc.esa.int
[4] https://sdup.esoc.esa.int/master/

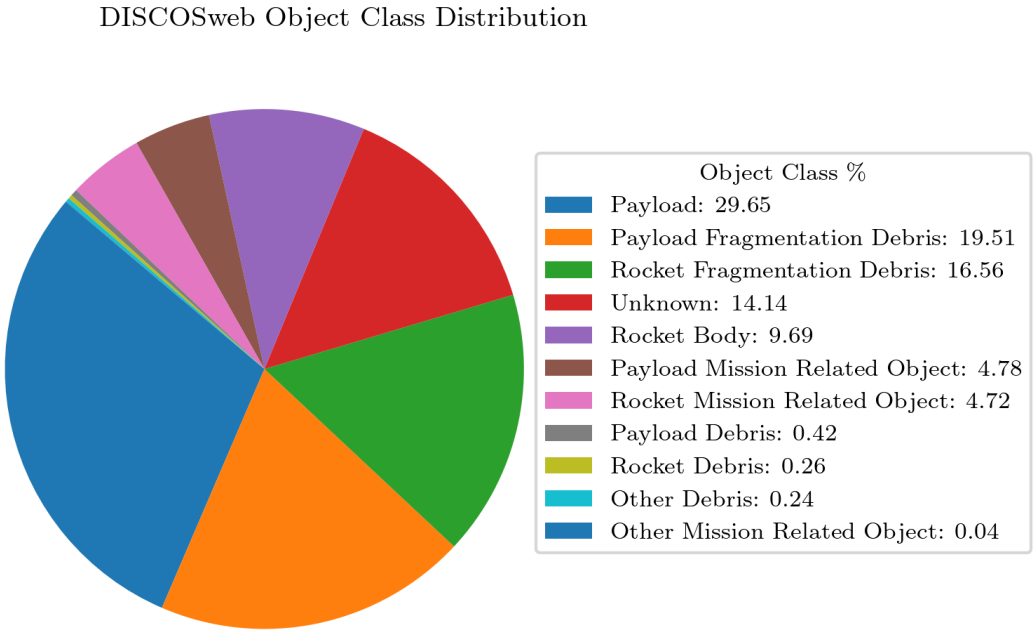

Figure 1.6: Fractional representation of the object classes scraped from the DISCOSweb database.

1.2.2.2 Challenges at LEO & GEO

In LEO, the deployment of large satellite constellations for communications, navigation, and Earth observation has led to a sharp increase in spatial density, especially within commonly utilised orbital shells between approximately 400 km and 1,200 km (European Space Agency, 2023a). Since the early 2000s, the number of operational satellites in LEO has increased to over ten thousand, a trend projected to continue as additional mega-constellations are launched (United Nations Office for Outer Space Affairs, 2019). This growth has significantly increased the frequency of close-approach events, resulting in a higher demand for collision avoidance manoeuvres and increasing operational burden on spacecraft operators; the largest commercial operator alone reported over 200,000 automated avoidance manoeuvres in the six months to May 2026 (Pultarova, 2026). The cost of this is not only operational. Avoidance burns draw on propellant budgeted for station-keeping and disposal, so a spacecraft that manoeuvres too often may be left without the reserves to de-orbit itself at end of life, thus shortening its operational lifetime. Each burn also invalidates the published ephemeris until a new orbit is fitted, worsening conjunction screening for everyone else in the interim. Since thresholds for manoeuvring are set conservatively to account for uncertainty in object state knowledge, better characterisation of the resident population reduces this burden directly. Concurrently, the accumulation of non-operational satellites, spent rocket bodies, and fragmentation debris has elevated the background collision risk throughout the LEO 'protected region', defined by the IADC as the spherical shell extending from the Earth's surface to an altitude of 2,000 km (Inter-Agency Space Debris Coordination Committee, 2007).

Of particular concern is the progressive increase in the orbital debris population and the associated risk of a collisional cascade, commonly referred to as the *Kessler Syndrome* (Kessler and

Cour-Palais, 1978). In such a scenario, collisions between orbiting objects generate debris at a rate exceeding natural removal mechanisms, primarily atmospheric drag. Long-term modelling studies conducted by space agencies indicate that certain densely populated LEO orbital shells may already be near or beyond critical density thresholds, such that debris growth could continue even in the absence of future launches (NASA Orbital Debris Program Office, 2018; European Space Agency, 2023a). Whether a runaway cascade develops on operationally relevant timescales remains debated: the modelled growth is slow, unfolding over decades to centuries, and depends strongly on future launch traffic and on compliance with post-mission disposal guidelines. The threat to sustained access to LEO is therefore less a sudden catastrophe than a gradual and effectively irreversible degradation of the most heavily used orbital shells.

In contrast, GEO is characterised by comparatively low object density but severe congestion due to the finite nature of orbital slots and radio-frequency allocations (International Telecommunication Union, 2020). Increasing competition for geostationary slots has reduced operational margins for station-keeping and increased the consequences of spacecraft anomalies. Although relative collision velocities in GEO are lower than in LEO, any fragmentation event at geostationary height would produce debris with unlimited orbital lifetimes due to negligible atmospheric drag (European Space Agency, 2023a). Furthermore, a growing population of ageing or uncontrolled spacecraft remains within or near the GEO 'protected region' (the domain bounded by $h_{GEO} \pm 200\,\text{km}$ and $\pm 15°$ in inclination (Inter-Agency Space Debris Coordination Committee, 2007)), exacerbating long-term congestion and risk. Figure 1.7 illustrates the longitudinal distribution of active geosynchronous satellites, highlighting the clustering of spacecraft at specific orbital slots along the geostationary belt. This distribution reflects commercial demand rather than orbital dynamics: each satellite can only serve users within a limited arc on either side of the longitude it sits above, so satellites concentrate over the Americas, Europe and Africa, and the Asia-Pacific, leaving minima over the Pacific and Atlantic oceans. The active geosynchronous satellite TLEs were obtained from the CelesTrak orbital element repository (Kelso, 2026).

Beyond spacecraft operations, the expanding satellite population has begun to measurably impact astronomical observations. Optical astronomy is increasingly affected by satellite trails crossing telescope fields of view and by a systematic increase in diffuse night-sky brightness caused by sunlight reflected from spacecraft and debris (International Astronomical Union, 2022). Radio astronomy faces similar challenges due to aggregate radio-frequency interference from large non-geostationary satellite constellations operating across wide bandwidths (Square Kilometre Array Observatory, 2021). These impacts are global in nature and cannot be mitigated solely through observatory siting or scheduling, raising concerns regarding the long-term viability of ground-based astronomical science.

1.2.2.3 Mitigation Strategies

The increasing congestion and collision risk within LEO and GEO have driven the development of a comprehensive suite of mitigation strategies aimed at ensuring the long-term sustainability of the orbital environment. These measures encompass international policy frameworks, spacecraft design and operational practices, and the emergence of active debris removal capabilities (Inter-Agency

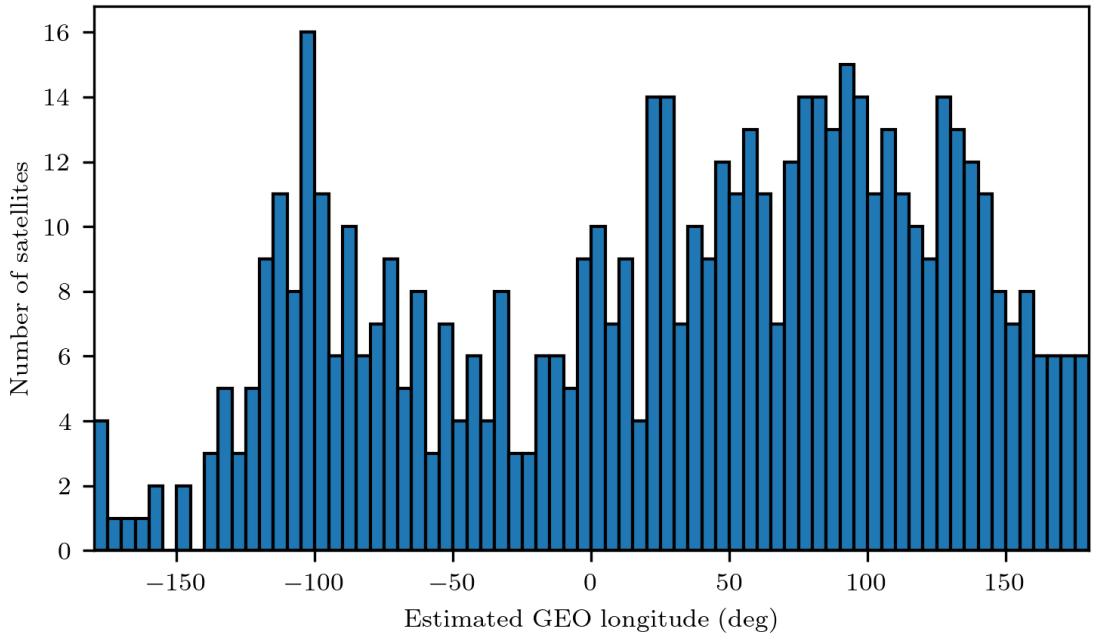

Figure 1.7: Distribution of estimated sub-satellite longitudes (the point on Earth's surface directly below the spacecraft) for active geosynchronous satellites derived from current TLE data at time of writing. Each bin spans 5° in longitude.

Space Debris Coordination Committee, 2007; United Nations Office for Outer Space Affairs, 2019).

At the international level, debris mitigation guidelines developed by the Inter-Agency Space Debris Coordination Committee (IADC) define *protected regions* in LEO and GEO and establish best practices for limiting the creation of long-lived debris (Inter-Agency Space Debris Coordination Committee, 2007). In LEO, post-mission disposal typically requires spacecraft to re-enter the atmosphere within a prescribed time period, historically set at 25 years but increasingly reduced by national regulators (Federal Communications Commission, 2022). In GEO, where natural orbital decay is negligible, end-of-life disposal requires spacecraft to manoeuvre into a disposal (graveyard) orbit above the geostationary protected region.

The minimum orbital height increment required for GEO disposal is commonly expressed as

$$\Delta H = 235 + C_r \frac{A}{m} \cdot 1000 \text{ [km]}, \tag{1.2}$$

where C_r denotes the solar radiation pressure coefficient and A/m is the spacecraft area-to-mass ratio (Inter-Agency Space Debris Coordination Committee, 2007). This formulation accounts for long-term perturbative effects, particularly solar radiation pressure, that may otherwise cause disposed objects to drift back into the operational GEO region.

In addition to regulatory measures, technological mitigation strategies are increasingly implemented at the spacecraft design stage. These include passivation of residual energy sources to prevent post-mission explosions, incorporation of propulsion or drag-augmentation devices to ensure reliable end-of-life disposal, and selection of operational orbital heights that enable natural decay following spacecraft failure (European Space Agency, 2023b). In response to concerns raised by the astronomical community, several operators (principally SpaceX) have adopted measures aimed

at reducing optical interference, including surface darkening treatments, deployable sunshades, and modified spacecraft attitudes during orbit raising (International Astronomical Union, 2022). Radio-frequency interference has been addressed separately, through downlink beam avoidance over radio observatories and bilateral coordination agreements (Square Kilometre Array Observatory, 2021). The effectiveness of these measures remains limited. Darkening coatings introduced thermal control penalties, and sunshades were abandoned where they conflicted with inter-satellite optical link performance; satellites in current constellations continue to exceed the magnitude 7 brightness threshold recommended by the International Astronomical Union (IAU) (Longa-Peña et al., 2026). Such measures are moreover adopted voluntarily, as no binding national or international licensing requirement presently constrains spacecraft optical brightness. They also address the observability of individual spacecraft rather than the underlying drivers of orbital congestion, and therefore leave unchanged the collision risk, debris generation and end-of-life disposal burden associated with constellation scale.

Active Debris Removal (ADR) represents a more direct intervention strategy, targeting large, high-risk debris objects for controlled removal from orbit. Missions flown to date have demonstrated component capabilities rather than end-to-end removal: ELSA-d validated magnetic capture of a client spacecraft fitted with a dedicated docking plate, this being a semi-cooperative demonstration rather than removal of an unprepared object (Astroscale Ltd., 2022); ADRAS-J achieved rendezvous and close-proximity inspection of an uncontrolled upper stage but did not attempt capture (Atarashi et al., 2024); and ClearSpace-1 has yet to fly (European Space Agency, 2021). The principal barrier, however, is economic rather than technical. The benefit of removing an abandoned object accrues to all users of the orbital regime while the cost falls on a single actor with no direct return, and no operator bears liability for debris it has already abandoned; ADR missions have consequently been procured by national agencies rather than commercial customers, and no self-sustaining commercial market yet exists for removal of legacy debris. Numerical simulations nonetheless indicate that removal of a small number of massive objects per year could significantly reduce the long-term probability of collisional cascading in LEO, provided removal is targeted at the objects contributing most to that risk (NASA Orbital Debris Program Office, 2018).

More broadly, the concept of *space sustainability* has emerged as a central principle guiding both government and commercial space activities, encompassing enhanced space situational awareness, improved data sharing for conjunction assessment, coordinated collision avoidance, and the development of space traffic management frameworks (United Nations Office for Outer Space Affairs, 2019). Progress remains uneven, however: these instruments are voluntary and implemented through divergent national licensing regimes (Inter-Agency Space Debris Coordination Committee, 2007; European Space Agency, 2023b), disposal compliance is incomplete particularly for rocket bodies (European Space Agency, 2025b) and the dominance of propulsion-related fragmentation noted in §1.2.2.1 indicates that passivation is not yet universal. Collision avoidance is likewise negotiated bilaterally, with no right-of-way convention and no obligation to manoeuvre. Underlying these limitations is a common dependence on knowledge of the objects themselves: safe approach to a tumbling target requires its spin-state in advance, prioritisation of removal targets requires discrimination of high-risk objects within the catalogue, and conjunction accuracy is limited by unmodelled area-to-mass variation arising from unknown attitude. These are observational problems

before they are regulatory ones, and it is to that high-precision observational capability that the present work is directed.

1.3 Thesis outline

The previous sections identified a recurring dependence: conjunction assessment, disposal verification, active debris removal, and the attribution of fragmentation events all require knowledge of what an RSO *is* i.e., attitude, surface materials, and operational condition, and none of this is available from catalogued orbital positions alone. For the RSOs, which are too distant and therefore too small to resolve, ground-based optical photometry is the principal source of such information: brightness, colour, and their temporal variation encode the geometry, reflectance, and rotational dynamics of the object through the physics of scattering light. Exploiting this signal at scale, however, requires wide-field instruments capable of high-cadence, multi-colour photometry of moving targets, together with analysis methods that can invert light curves for physical parameters with quantified uncertainties. The aim of this thesis is to develop and demonstrate that capability i.e., from instrument characterisation, through survey photometry, to spin-state inversion of individual objects.

To that point, this thesis explores how high-precision photometry can be utilised in such a way to draw insight into the characteristics of various RSOs. The primary datasets come from various instruments at the Roque de los Muchachos Observatory (ORM) on La Palma, Canary Islands. Chapter 2 outlines the sources of data, observing strategies, data reduction and various light curve analysis techniques. Chapter 3 outlines the characterisation of the STING instrument. Chapter 4 highlights the photometric results from a conducted survey of the GEO-belt with the aforementioned instrument. Chapter 5 highlights how photometry can be used to constrain the spin-state of a well-defined object such as Ajisai. Chapter 6 discusses the improved spin-state estimation through use of colour and polarimetry as complementary observations. Finally, Chapter 7 includes a summary of the main body of the thesis as well as a future outlook. Supplementary appendices can be found relating to various BRDF models, reference frame transforms and calculations relating to the Ajisai models in Chapter 5.

2 Methods

"Do you want to bet against us?"

Ron Saunders, 14th April 1981

This chapter describes the relevant technical background, sources of data used throughout this thesis, the observational strategies adopted, and a high-level overview of data reduction necessary to go from raw uncalibrated frames to high-precision photometric light curves. The technical background section establishes the reference frames, coordinate systems, time standards and reflectance models on which the later analysis depends. The emphasis is on outlining standard astronomical calibration methods and demonstrating how these are adapted to the specific requirements of ground-based optical satellite observations. A worked example simulation of a tumbling RSO light curve is explored. The chapter concludes with a description of various analytical techniques that can be applied to calibrated photometric light curves.

2.1 Technical Background

The measurements presented in this thesis are positions and brightnesses of RSO recorded from the ground over time. Interpreting either requires conventions and context to be established first. A position is meaningful only once the frame in which it is expressed, the coordinates used to describe it, and the time standard against which it is recorded have been defined, since the observer, the Sun and the target are all in continuous relative motion, and since the catalogue orbits against which detections are matched are themselves distributed in a particular reference frame and time system.

Similarly, a brightness can be made more meaningful once the illumination geometry has been specified and how a model adopted for the reflectance of sunlight from a satellite's surfaces has been defined, since an unresolved object provides no spatial information and its brightness can vary with time, aspect, phase angle and surface material. This section defines the frames, coordinate systems and time standards used, and the reflectance models that relate a satellite's shape, attitude and materials to its observed brightness.

2.1.1 Reference Frames, Geometries and Time

2.1.1.1 Fundamental Concepts

The accurate mathematical description of motion and position in space depends critically on the selection and transformation of reference frames. A reference frame is a well-defined coordinate system that enables consistent specification of spatial quantities—position vectors, velocities,

accelerations—and their time evolution. These frames provide the scaffolding upon which the equations of motion are constructed and interpreted, especially in multi-body systems where relative motion plays a dominant role.

A reference frame consists of three primary components: a defined origin (often the centre of mass (COM) of a celestial body); a basis of orthogonal axes (typically labelled $\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}$); and a convention specifying how those axes are oriented and how that orientation evolves in time. The orientation convention fixes the axis directions against physical references i.e., alignment with the Earth's mean equator and equinox at a stated reference epoch, with the positions of distant extragalactic radio sources, or with the instantaneous terrestrial pole and prime meridian. Because these physical references themselves move i.e., the Earth rotates, and its spin axis precesses and nutates, the convention must also state whether the axes follow that motion or are held fixed at an epoch, and, where they follow it, which model of the motion is adopted. It is this choice that distinguishes, for instance, an Earth-centred inertial frame from an Earth-centred, Earth-fixed one realised from the same origin. Finally, positions and velocities expressed in any frame must be timestamped against a defined time standard (e.g., UTC, TAI, TT), both to relate observations to one another and to evaluate the time-dependent transformations between frames.

Broadly speaking, reference frames are classified as either inertial or non-inertial. Inertial reference frames are those in which Newton's first law holds: a body not acted upon by a net external force will continue in a state of uniform motion. In practice, no physical frame is perfectly inertial. A frame attached to a rotating or accelerating body is non-inertial in the Newtonian sense, while even a non-rotating geocentric frame is in free fall about the Sun and is inertial only as much as tidal accelerations across the region of interest can be neglected. More fundamentally, general relativity admits no global inertial frame. Inertial frames exist only locally, for freely falling observers and so practical "inertial" frames such as the geocentric celestial reference system (GCRS) are conventional constructions: kinematically non-rotating with respect to distant extragalactic sources, with residual relativistic effects absorbed into the adopted IAU models (IERS Conventions, 2010). And to that point, many frames—such as those centred on the barycentre of the solar system or aligned with distant extragalactic sources—are sufficiently close to inertial for high-precision applications. Non-inertial frames, on the other hand, undergo rotation or acceleration relative to inertial space, and therefore require the inclusion of fictitious forces—such as Coriolis and centrifugal forces—in the dynamical equations. These frames are often necessary when describing motion relative to rotating planetary surfaces or in vehicle-fixed coordinate systems.

The selection of an appropriate reference frame is not arbitrary; it is determined by the requirements of the problem at hand. For instance, interplanetary trajectory design is most naturally carried out in a heliocentric inertial frame, while Earth observation missions are typically formulated in Earth-fixed or topocentric frames. Crucially, the ability to transform between frames underpins all practical usage. These transformations often involve time-varying rotation matrices, which account for physical phenomena such as Earth's precession, nutation, polar motion, and sidereal rotation.

2.1.1.2 Common Reference Frames and Coordinates

The International Celestial Reference System (ICRS) is the fundamental reference system adopted by the IAU for high-precision positional astrometry. Its origin lies at the barycentre of the solar system, and the directions of its axes are fixed relative to distant extragalactic sources. The International Celestial Reference Frame (ICRF) is the practical realisation of these fiducial directions, originally constructed from VLBI (Very Long Baseline Interferometry) observations of 212 defining extragalactic radio sources (Ma et al., 1998); the current realisation, ICRF3, contains over four thousand sources and provides a noise floor of ~30 μas (Charlot et al., 2020). Unlike its predecessors, optical fundamental-star catalogues such as FK5, anchored to the mean equator and equinox of J2000.0, whose accuracy was ultimately limited by Earth-orientation modelling and the proper motions of the reference stars themselves (Warren, 1990), the ICRS, adopted by the IAU in 1997, is kinematically non-rotating, with axes defined by sources of negligible proper motion at present accuracies, thereby decoupling the celestial frame from the dynamics of the Earth-Moon system entirely (International Astronomical Union, 1997).

The ICRS axes are intentionally consistent, to better than 0.1 arcsecond, with the equator and equinox of the J2000.0 epoch (January 1, 2000, at 12:00 Terrestrial Time, TT); that being said, the ICRS axes are meant to be regarded as fixed directions in space that have an existence independent of the precession and nutation of the Earth. The GCRS (International Astronomical Union, 2000) is the Earth-centred equivalent of the ICRS. Its primary direction is the *vernal equinox* i.e., the point at which the Sun crosses the equatorial plane moving northward. This makes the GCRS an example of an Earth-centred inertial (ECI) frame.

The equatorial coordinates given from the ICRS are Right Ascension (RA) and Declination (DEC), defined respectively as the angle measured eastward from the vernal equinox along the celestial equator, and the angular elevation from the equator to the north celestial pole.

The ICRS is the frame in which the astrometry throughout this thesis is ultimately expressed: satellite positions measured against catalogued reference stars inherit it directly, and its μas-level noise floor is negligible compared with the uncertainties of satellite astrometry. It likewise underpins the dynamical analysis of later chapters. In particular, the spin-state reconstruction of Ajisai discussed later in this thesis is formulated in the ICRF, with phase-angle-bisector vectors computed in that frame and the recovered spin-pole orientation reported in ICRF coordinates, while the astrometric validation against GPS satellites requires transforming precise ephemerides from the Earth-fixed frame in which they are distributed into the GCRS, the Earth-centred counterpart of the ICRS, for comparison with observed positions. A consistent celestial reference system is therefore not incidental background but the common frame linking the observational, astrometric, and dynamical components of this work.

A practical distinction must be drawn at this point between *astrometric* and *apparent* positions. Astrometric coordinates – such as the RA and DEC catalogued in the ICRS – refer to a fixed inertial frame and are the natural quantities for cataloguing and ephemeris work. Apparent coordinates, by contrast, refer to the true equator and equinox of date as seen from the observer's location, and incorporate precession, nutation, polar motion, stellar aberration, and (where relevant) light-time and atmospheric refraction. The hour-angle and horizontal coordinates introduced below are

inherently apparent quantities.

In contrast to the RA, a choice can be made to measure the azimuthal component of the spherical coordinates of an object in terms of Hour Angle (HA). The HA of a given point is defined as the angle between two planes: the first containing the Earth's axis and the local zenith (the local meridian), and the second containing the Earth's axis and the given point on the celestial sphere. This angle is measured westward of the meridian plane and can be defined by equation (2.1):

$$\mathrm{HA} = \mathrm{LST} - \alpha_{\mathrm{obj}}, \tag{2.1}$$

where the Local Sidereal Time (LST) is the apparent RA crossing the observer's local meridian and α_{obj} is the apparent RA of the object. The HA therefore describes the amount of sidereal time that has passed since the object crossed the observer's local meridian; a negative HA indicates the time until the next transit of the object across the local meridian. The relationship in equation (2.1) is only consistent when α_{obj} is the apparent RA at the observation epoch, not its J2000 catalogue value. The conversion between the two is routine in modern software, i.e., via the IAU Standards of Fundamental Astronomy (SOFA) routines (IAU SOFA Board), but the distinction is easy to overlook in practice.

Alternatively, the position of a satellite directly relative to the observer's local horizon can be represented, in a Horizontal Coordinate System (HCS). This system corresponds to two independent angular coordinates, altitude and azimuth: the altitude is measured from the local horizon to the zenith (0–90° for visible targets), and the azimuth is measured eastwards from true north. Like HA, both quantities are defined in the apparent frame and depend on the observer's location and the time of observation. The definitions of RA–DEC and Alt–Az are presented geometrically in Figure 2.1.

Another useful choice of reference frame could be an Earth-Centred Earth-Fixed (ECEF) frame. A specific example of this reference system is one in which GPS satellites utilise, WGS-84. Which independently defines the ellipsoid of the Earth (National Geospatial Intelligence Agency, 2021). The ECEF coordinate system can be defined by the following (see Figure 2.2):

- Origin lies at the centre of the chosen ellipsoid, for WGS84, this is the centre of mass of the Earth.
- The Z-axis lies on the line between the north and south poles, with positive values northward. In WGS-84, this is the international reference pole (IRP), which does not exactly coincide with the Earth's spin axis (Snay and Soler, 1999).
- The X axis lies along the plane of the equator, passing through the origin and extending from 180° longitude (negative) to the prime meridian. In WGS-84, this is the IERS (International Earth Rotation and Reference System Service) Reference Meridian, 5.3 arcseconds (102 metres) east of the Greenwich meridian at the latitude of the Royal Observatory, the discrepancy is due to the local gravity field at Greenwich not passing through the Earth's centre of mass (Malys et al., 2015).
- The Y axis is also in the plane of the equator, passing through the origin and extending from 90° West to 90° East longitude (positive).

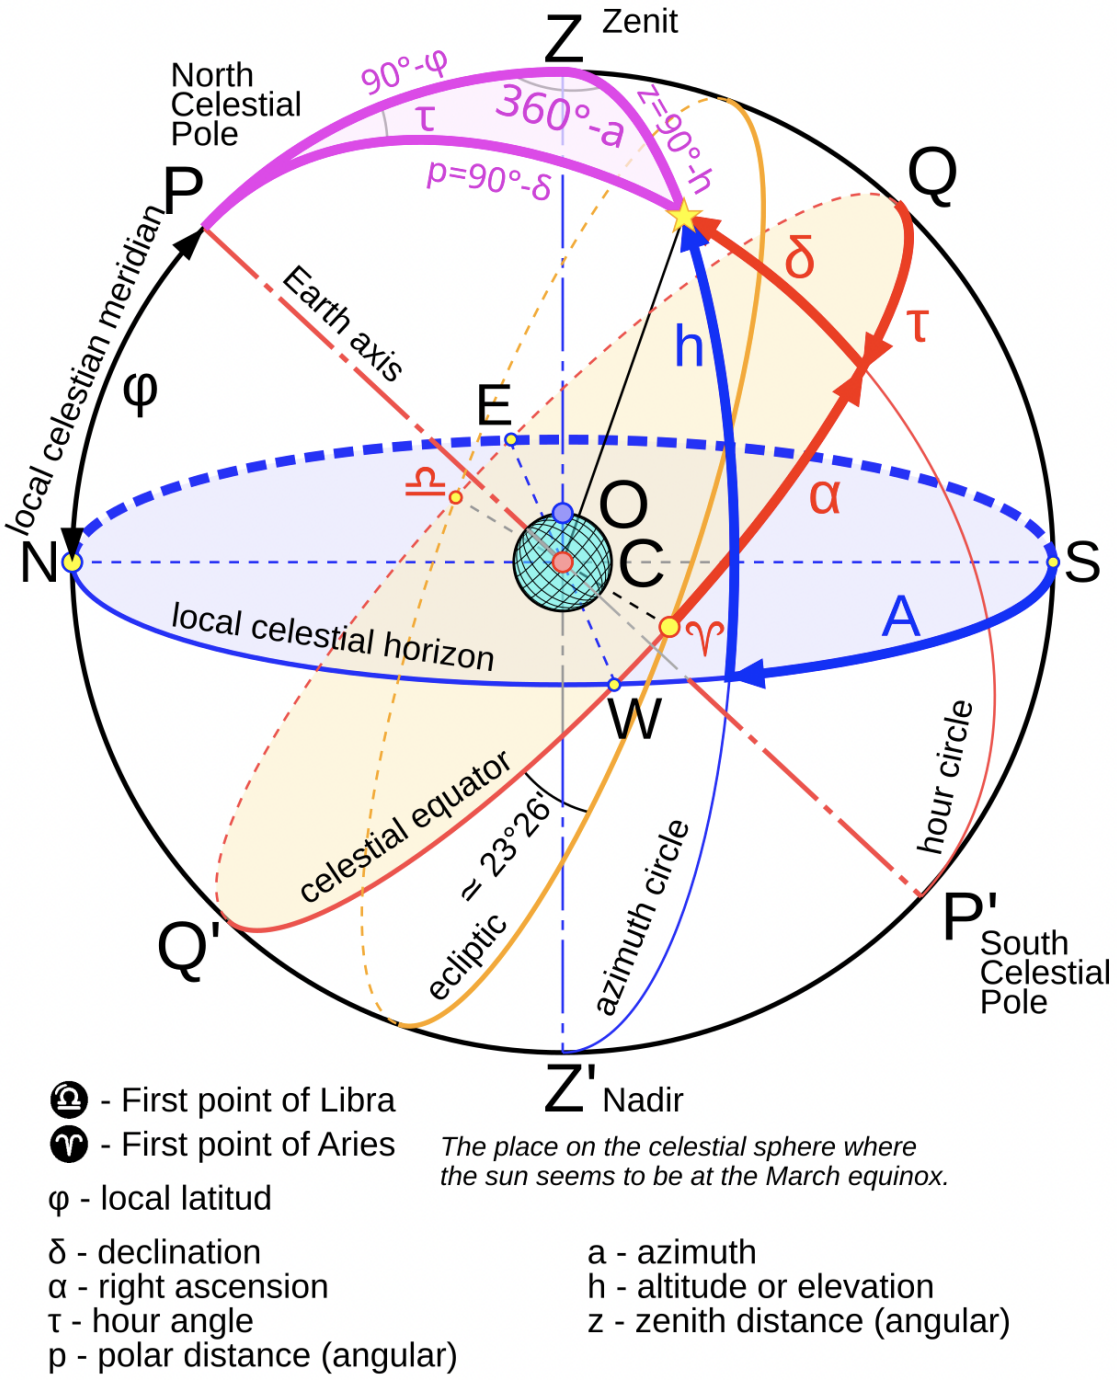

Figure 2.1: Representation of equatorial and horizontal coordinate systems with respect to the equator and equinox of J2000 and a local observer situated at the latitude of φ. *(CC by SA 4.0 - Jon Peli Oleaga)*

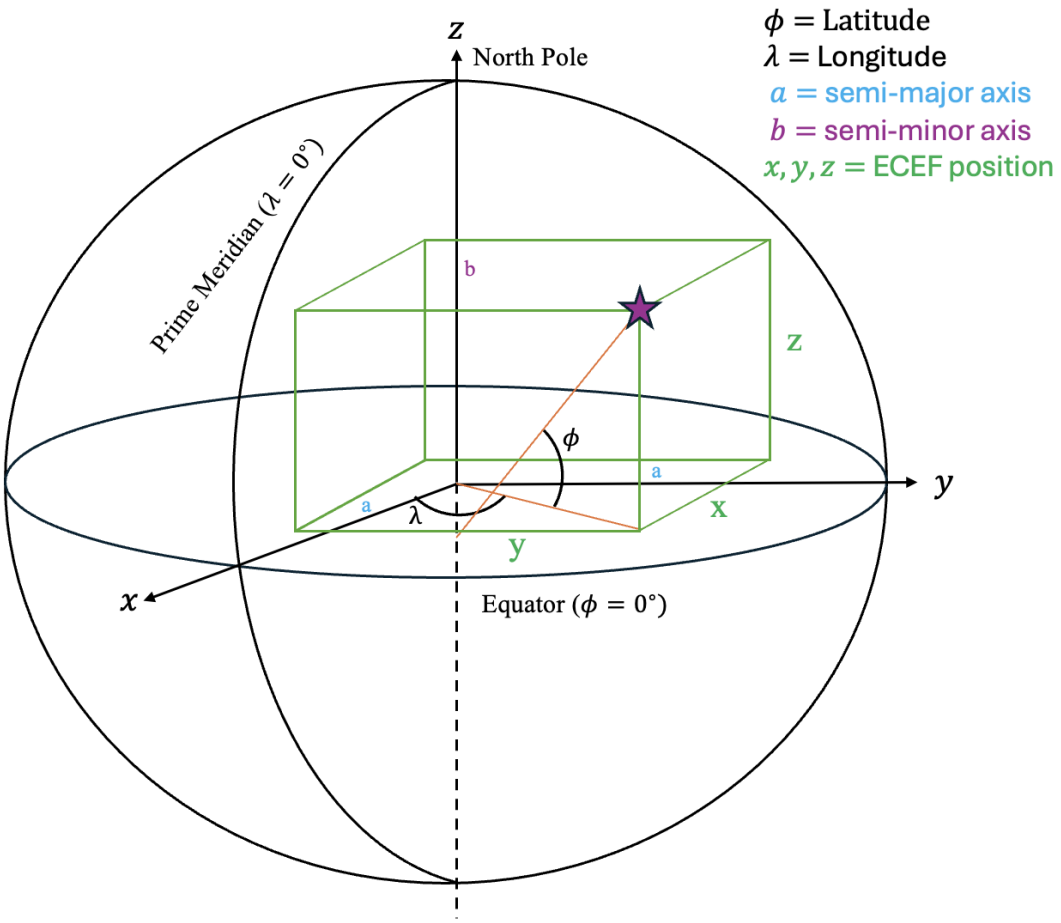

Figure 2.2: ECEF coordinates (x, y, z) shown in relation to a geodetic latitude and longitude position.

2.1.1.3 Useful Geometry

The concept of the phase angle bisector (PAB) is particularly useful in relation to the reflections from a satellite facet hit by sunlight and received by the observer. The PAB is defined as the unit vector halfway between the satellite-to-Sun ($\mathbf{S_{Sun}}$) and satellite-to-observer directions ($\mathbf{S_{obs}}$):

$$\hat{\mathbf{p}} = \frac{\vec{\mathbf{S}}_{\mathbf{obs}}(\mathbf{t}) + \vec{\mathbf{S}}_{\mathbf{Sun}}(\mathbf{t})}{\|\vec{\mathbf{S}}_{\mathbf{obs}}(\mathbf{t}) + \vec{\mathbf{S}}_{\mathbf{Sun}}(\mathbf{t})\|} \tag{2.2}$$

The PAB is useful for identifying potential specular reflections from flat facets. For an ideal specular surface, the surface normal must bisect the angle between the incident solar direction and the outgoing observer direction. Equivalently, a glint is expected when the surface normal, $\hat{\mathbf{n}}$, is closely aligned with the PAB vector, $\hat{\mathbf{p}}$, such that

$$\theta_{\mathrm{PAB}} = \cos^{-1}\left(\hat{\mathbf{n}} \cdot \hat{\mathbf{p}}\right) \tag{2.3}$$

is small. In practice, the tolerance depends on the finite angular size of the Sun, the surface roughness, and the degree to which the reflection is truly specular.

For an unresolved RSO, the reflected sunlight measured by an observer over time forms a light curve, defined here as a time series of the object's apparent brightness. The morphology of this light curve depends on the relative geometry between the Sun, object, and observer, as well as the object's attitude and surface reflectance properties.

The profile and features within observed light curves are the result of sunlight reflected from the satellite's surface in the direction of the observer. The most useful way to capture this geometry is through the concept of the solar equatorial phase angle (Payne et al., 2007). The solar equatorial phase angle is defined as the angle between the projection of Sun and observer on the celestial equator (Xiao-Fen et al., 2021) or moreover corresponds to the difference in right ascension between the satellite and anti-solar point (the point directly opposite the Sun on the celestial sphere from the observer perspective) i.e., $\phi_{eq} = 180 - \Delta\alpha$.

Figure 2.3 illustrates the geometry between observer, satellite and Sun and thus the construction of the PAB.

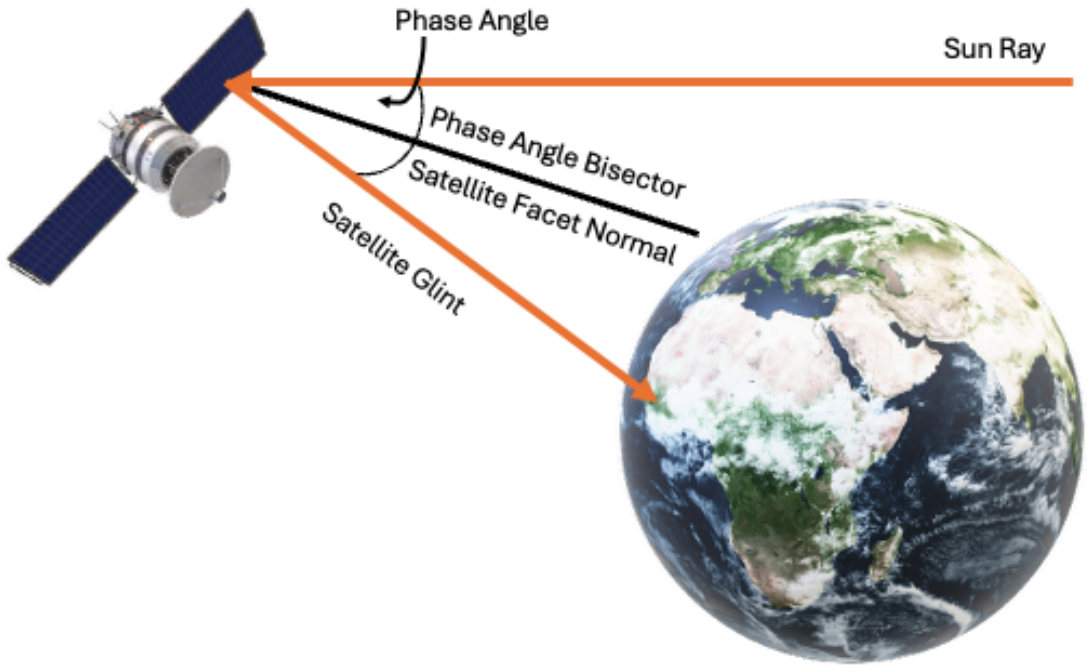

Figure 2.3: An illustration of the construction of the phase angle bisector vector with respect to the Sun and observer directions. This is with respect to the glint condition for a flat facet.

2.1.1.4 Time Systems in Space Applications

Accurate and unambiguous time representation is fundamental to orbital mechanics, satellite operations, astronomy and SDA. In data fusion applications, measurements originating from heterogeneous sensors and systems must be referenced to a common and well-defined time scale to ensure consistency.

Timekeeping in space applications is complicated by the fact that not all time systems are defined in the same way. Some, such as International Atomic Time (TAI), are continuous atomic time scales based on the SI second, whereas others, such as Coordinated Universal Time (UTC), are tied to the rotation of the Earth and therefore include leap seconds to remain close to mean solar time. Satellite navigation systems, such as GPS, use their own continuous time scales, while high-precision astronomical timing may require corrections for the changing position of the observer, such as heliocentric or barycentric light-travel-time corrections. Relativistic effects are also important for precise satellite navigation and timing, particularly when comparing clocks in different gravitational potentials or states of motion.

Several time systems are therefore routinely encountered in orbital mechanics and astronomical observations, including Modified Julian Date (MJD), Coordinated Universal Time (UTC), International Atomic Time (TAI), and GPS Time (GPST). Each serves a distinct role, and understanding their definitions and interrelationships is essential for correctly interpreting time-series data.

International Atomic Time (TAI) is a continuous and uniform time scale realised by a weighted ensemble of atomic clocks maintained by national metrology institutes and coordinated by the Bureau International des Poids et Mesures (BIPM). TAI progresses strictly in SI seconds and does not incorporate leap seconds, making it a monotonic time reference (Bureau International des Poids et Mesures, 2011). Due to its continuity, TAI is particularly well suited for high-precision applications such as orbital dynamics modelling, numerical integration, and state estimation, where discontinuities in time would introduce errors. Although TAI is not typically used for operational timestamping, it forms the physical basis from which other commonly used time scales are derived.

Coordinated Universal Time (UTC) is the primary civil time standard and is widely used for operational and human-readable timestamps. UTC is derived from TAI but is periodically adjusted through the insertion of leap seconds to maintain alignment with Earth rotation time (UT1) (McCarthy et al., 2008). These leap seconds introduce discontinuities, rendering UTC a non-uniform time scale. In satellite operations and SDA, UTC is commonly used for timestamping sensor observations, scheduling satellite and ground station activities, and exchanging data between international systems. However, because of its non-uniform nature, UTC is often converted to a continuous atomic time scale prior to use in orbit propagation, filtering, or data fusion algorithms. The relationship between UTC and TAI is given by

$$\mathrm{UTC} = \mathrm{TAI} - \mathrm{T}_{\mathrm{leap}}\ [\mathrm{s}], \tag{2.4}$$

where $\mathrm{T}_{\mathrm{leap}}$ denotes the cumulative number of leap seconds introduced since the establishment of UTC.

GPS Time (GPST) is the time scale used internally by the Global Positioning System and is also widely adopted by satellite and ground-based systems equipped with GPS-disciplined clocks.

GPST is a continuous atomic time scale that does not incorporate leap seconds. It was defined to coincide with UTC at the GPS epoch (00:00:00 UTC on 6 January 1980). At that epoch, UTC was offset from TAI by 19 seconds, and this offset was fixed into the GPS time definition (Duffett-Smith and Zwart, 2011). As a result, the relationship between TAI and GPST is constant:

$$\mathrm{TAI} = \mathrm{GPST} + 19.0 \ [\mathrm{s}]. \tag{2.5}$$

Since UTC continues to introduce leap seconds while GPST does not, the offset between GPST and UTC increases over time according to

$$\mathrm{GPST} = \mathrm{UTC} - 19.0 + \mathrm{T}_{\mathrm{leap}} \ [\mathrm{s}]. \tag{2.6}$$

The original Julian Day (JD) is used by astronomers and expressed in days since noon January 1, 4713 B.C. Modified Julian Date (MJD) is a time representation originating in astronomy and astrodynamics, designed to express epochs as a continuous count of days and fractional days. It is derived from the JD through

$$\mathrm{MJD} = \mathrm{JD} - 2\,400\,000.5. \tag{2.7}$$

The subtraction of the constant offset shifts the reference epoch to midnight on 17 November 1858 and significantly reduces the numerical magnitude of the date, making MJD more convenient for analytical work and numerical computation (Duffett-Smith and Zwart, 2011), MJD was adopted by the Smithsonian Astrophysical Observatory (SAO) in 1957 for the tracking of satellite orbits (Technische Universität Wien, 2007) due to this very reason. In astronomy, MJD is widely used to label observation times, ephemerides, and time series data, as it avoids the complexity of calendar-based representations while preserving continuity.

It is important to emphasise that MJD does not define a time scale by itself; rather, it is a representation that must be associated with an underlying time standard such as UTC or TAI (e.g., $\mathrm{MJD}_{\mathrm{UTC}}$ or $\mathrm{MJD}_{\mathrm{TAI}}$). Ensuring consistency of the underlying time scale is essential when interpreting MJDs, particularly in the presence of leap seconds.

In addition to atomic and civil time scales, the concepts of the *solar day* and the *sidereal day* provide important physical context for timekeeping in space applications. A mean solar day is defined as the average interval between successive transits of the Sun across a given meridian and has a nominal duration of 24 hours. It reflects the rotation of the Earth relative to the Sun and underpins the civil day used in everyday timekeeping (IERS Conventions, 2010; Seidelmann, 1992). In contrast, a sidereal day is defined as the interval between successive transits of a distant celestial reference direction, typically a fixed star or the vernal equinox, across the same meridian. The sidereal day has a duration of approximately 23 hours, 56 minutes, and 4 seconds, making it about four minutes shorter than the mean solar day (Vallado, 2013; U.S. Naval Observatory, b). This difference arises because the Earth advances along its orbit around the Sun while rotating, requiring an additional small amount of rotation each day for the Sun to return to the local meridian (Seidelmann, 1992). The definitions of a solar and sidereal day are illustrated by the diagram in Figure 2.4.

The distinction between solar and sidereal days is particularly relevant in orbital mechanics and satellite tracking. Earth rotation models used to transform between inertial reference frames and

Earth-fixed frames are fundamentally based on sidereal rotation, as satellites orbit relative to inertial space rather than the Sun (Vallado, 2013). Consequently, quantities such as Earth rotation angle, Greenwich sidereal time, and UT1 are central to precise orbit determination, ground station pointing, and sensor line-of-sight calculations (IERS Conventions, 2010). At the same time, the gradual and irregular variation of the Earth's rotation rate relative to atomic time leads to a divergence between solar time and uniform atomic time scales, which is the physical motivation for the introduction of leap seconds in UTC (IERS Conventions, 2010; McCarthy and Seidelmann, 2008).

In summary, UTC and GPST are commonly used for operational timestamping and system synchronisation, while TAI provides a continuous and uniform reference suitable for high-precision computations. MJD offers an astronomically motivated and compact way to represent time while remaining independent of calendar conventions. The concepts of the solar and sidereal day underpin the connection between atomic timekeeping and the Earth's rotation, and are fundamental to accurate modelling of satellite motion, reference frame transformations (see Appendix B), and observation geometry in space applications.

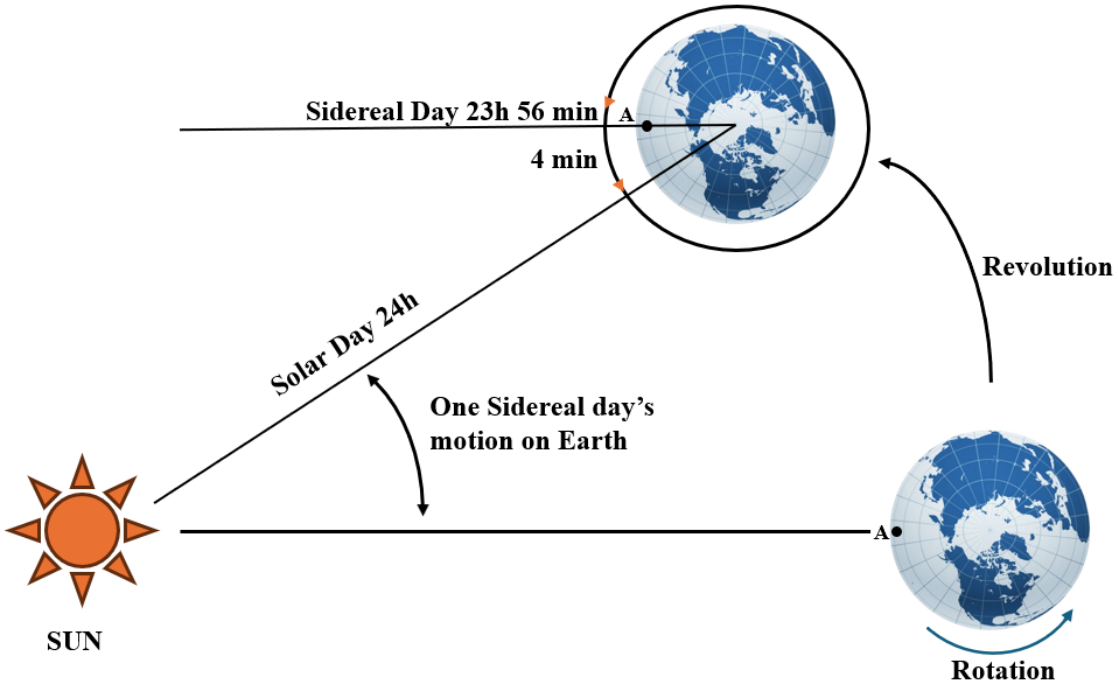

Figure 2.4: Relationship between the sidereal and mean solar day. One full rotation relative to distant stars defines a sidereal day (≃23 h 56 min). During this interval, Earth advances along its orbit, requiring an additional ≃4 minutes of rotation for the same meridian (point A) to realign with the Sun, producing the mean solar day of 24 h.

2.1.2 Optical Brightness Models

It is also useful to compare optical measurements to those derived purely from models describing the interaction between surfaces and the solar rays imparted upon it and the resultant scattered flux received at the observer. This can provide insight into the shape, relative orientation, size, dynamics and surface properties of a given satellite.

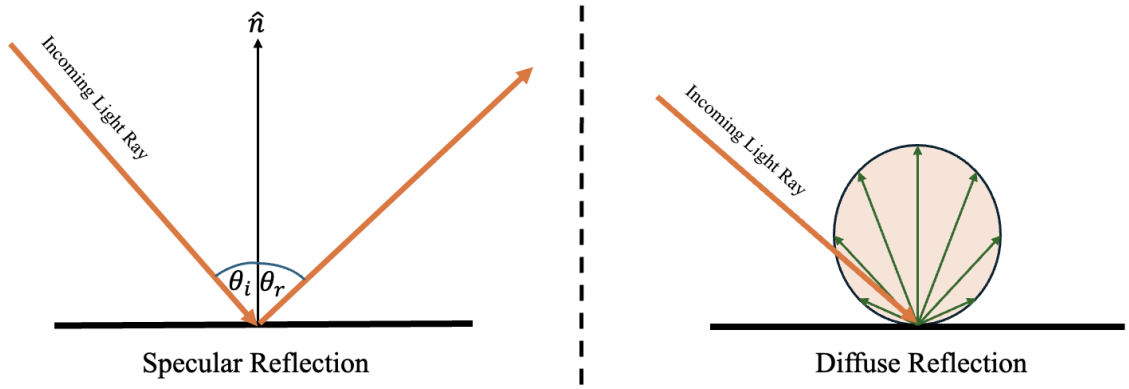

FIGURE 2.5: Demonstration of the difference between a specular reflection versus a diffuse (Lambertian) reflection from a surface.

2.1.2.1 FUNDAMENTALS

Light scattering from objects in space depends on both their geometric shape and their surface reflection properties. When sunlight illuminates an object, a fraction of the incident energy is reflected toward the observer, contributing to its apparent optical brightness. Two reflection behaviours are typically considered:

- **Specular reflection** — mirror-like scattering of photons concentrated in a single direction whereby the angle of incidence is always equal and opposite the angle of reflection ($\theta_i = \theta_r$) for all light rays.
- **Diffuse reflection** — scattering distributed more uniformly over many directions. An ideal diffuse reflecting surface is known as Lambertian. The apparent brightness of a Lambertian surface to an observer is the same regardless of the observer's angle of view. More precisely, the reflected radiant intensity obeys Lambert's cosine law, $I = I_0 \cos(\theta)$ (Lambert, 1760). Here each light ray has it's own angle of incidence and reflection.

Specular surfaces are smooth at the scale of the wavelength of the incident light, so when a beam of light hits them, it reflects in a single, well-defined direction. Quantitatively, the Rayleigh criterion states that a surface reflects specularly when its RMS height irregularities satisfy $\sigma < \lambda/(8\cos\theta_i)$, where θ_i is the angle of incidence (Beckmann and Spizzichino, 1987); in the optical ($\lambda \sim 550$ nm) this corresponds to $\sigma \lesssim 70$ nm at normal incidence. Diffuse surfaces, by contrast, have irregularities comparable to or larger than the wavelength, so even though the law of reflection applies locally on each microscopic facet, the incoming beam is scattered in many directions. Figure 2.5 demonstrates the difference between specular and diffuse reflections.

These two limiting cases also have different polarimetric consequences. Specular reflection from smooth surfaces can produce strongly geometry-dependent linear polarisation, with the polarisation angle tied to the scattering plane and the local surface orientation. Diffuse or multiply scattered reflection is generally smoother in phase and tends to reduce the degree of polarisation. This distinction is relevant later in the thesis, where simultaneous photometry and polarimetry are used to interpret glint-like features in the Gorizont-6 light curves.

For canonical shapes such as spheres, cylinders, and flat plates, analytical expressions exist to approximate how reflected light is distributed as a function of phase angle Φ (the Sun-object-observer angle). These relationships are commonly expressed as *phase functions*, $F(\Phi)$, which quantify the fraction of reflected light visible to an observer for a given geometry and surface behaviour.

The phase function appears directly in the apparent visual magnitude model for unresolved objects, which follows from expressing the observed flux from the object as a fraction of the flux received directly from the Sun. An object of effective illuminated area A and albedo ρ intercepts and reflects a fraction of the incident solar flux, redistributes it angularly according to the phase function $F(\Phi)$, with $F(0)$ considered as when the satellite is at full phase. The reflected flux received by the observer is diluted by the inverse square of the observer-object distance R, giving a flux ratio $F_{obj}/F_{\odot} = \rho\, A\, F(\Phi)/R^2$. Converting this ratio to magnitudes via $\Delta m = -2.5 \log_{10}(F_{obj}/F_{\odot})$ yields:

$$m_{vis} = m_{Sun} - 2.5 \log_{10}(\rho\, A\, F(\Phi)) + 5 \log_{10}(R), \tag{2.8}$$

The argument of the logarithm is dimensionless provided A and R are expressed in consistent units, the $\rho A F(\Phi)$ and $5 \log_{10}(R)$ terms being the single ratio $\rho A F(\Phi)/R^2$ separated using logarithm laws for clarity. For the case of a composite shape, the product $\rho\, A\, F(\Phi)$ is simply summed over the components.

Table 2.1 summarises commonly used specular and diffuse phase functions for simple shapes (McCue et al., 1971). These formulas enable first-order brightness modelling when detailed material properties or complex geometries are unknown.

Shape	Reflection Mode	Phase Function, F
Sphere	Specular	$\frac{1}{4\pi}$
	Diffuse	$\frac{2}{3\pi^2}[(\pi - \Phi)\cos(\Phi) + \sin(\Phi)]$
Cylinder	Specular	$\frac{\alpha}{4\Delta}$
	Diffuse	$\frac{[(\pi-\Phi)\cos(\Phi)+\sin(\Phi)]}{4\pi}$
Flat Plate	Specular	$\frac{4\cos(\frac{\Phi}{2})}{\pi\Delta^2}$
	Diffuse	$\frac{\cos^2(\frac{\Phi}{2})}{\pi}$

TABLE 2.1: Representative phase functions for idealised shapes using specular and diffuse scattering assumptions. Φ denotes the phase angle, α denotes a time-dependent flash parameter which describes the fact that during the flash of length t, one is seeing the integrated light across a segment of the Sun's disk, which was assumed to be of uniform brightness. Δ denotes the beam width (≥ 0.0093 rad – the angular diameter of the Sun).

Figure 2.6 highlights why such simple brightness estimates should be treated with caution. The apparent brightness is sensitive not only to intrinsic properties such as albedo and effective area, but also to phase angle, range, object shape, and attitude.

Furthermore, Figure 2.6 illustrates how each parameter of Equation 2.8 shapes the apparent brightness of an unresolved object, using the phase functions of Table 2.1 evaluated for representative geometries. Panel (a) shows the phase-angle dependence for idealised shapes at GEO: diffuse geometries fade steadily towards high phase angle as the illuminated fraction visible to the observer

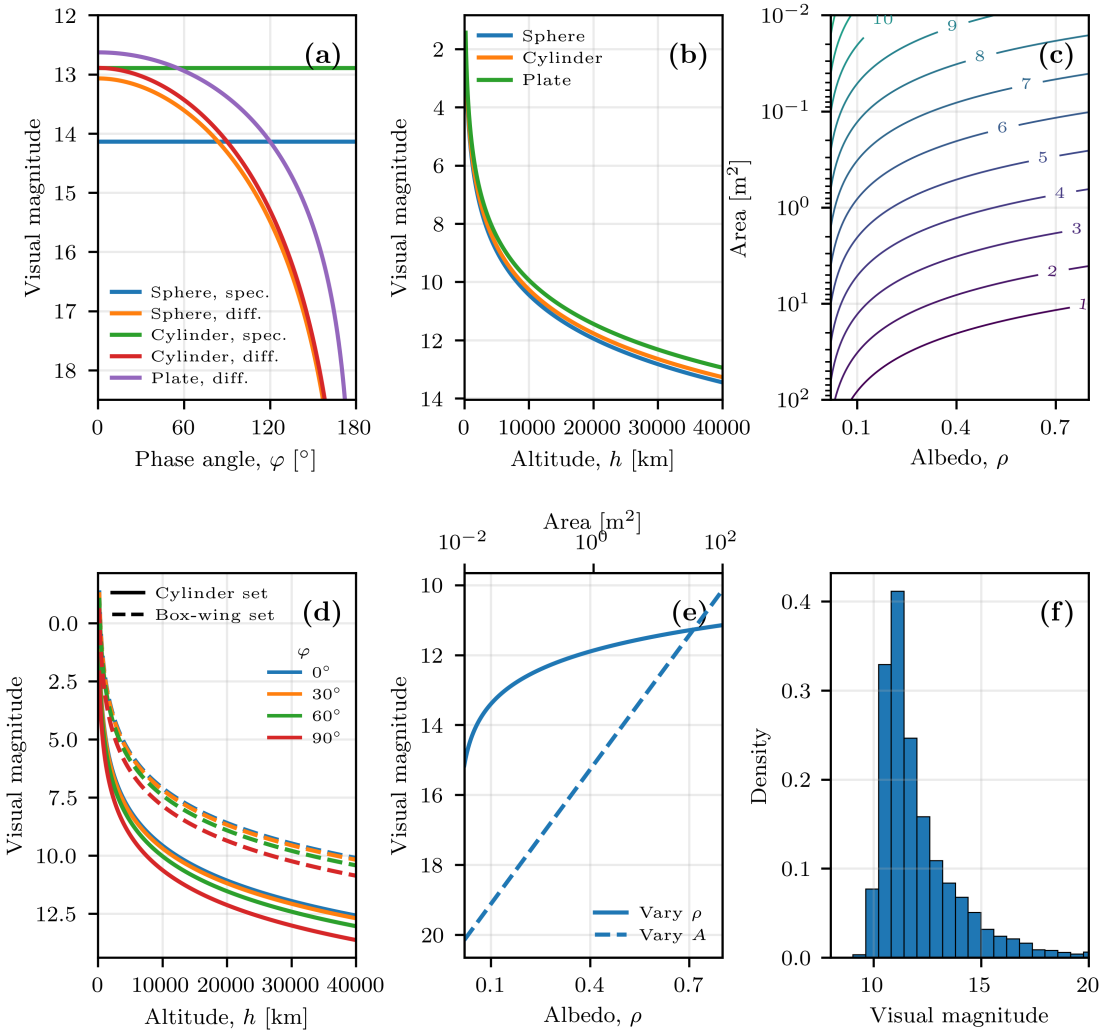

FIGURE 2.6: Parameter sensitivities in satellite visual brightness modelling using simple geometric shapes. **(a)** Apparent visual magnitude versus phase angle for idealised diffuse and specular shapes in GEO. **(b)** Apparent magnitude as a function of altitude for a selection of diffuse shapes at a fixed phase angle ($\varphi = 30°$). **(c)** Contours of apparent magnitude as a function of albedo and effective reflecting area for a diffuse plate in LEO (550 km). **(d)** Composite satellite configurations (cylinder and box-wing): magnitude versus orbital height for multiple phase angles. **(e)** One-dimensional parameter variations showing sensitivity to albedo and area for a diffuse sphere at GEO. **(f)** Monte-Carlo distribution of apparent magnitude at GEO for a representative composite geometry with randomized material and geometric properties.

shrinks, while specular geometries are phase-independent at this level of idealisation, appearing as constant brightness. Panel (b) isolates the orbital height term, showing the $5\log_{10}(R)$ dimming of the diffuse shapes at fixed phase angle from LEO out to GEO orbital heights. Panel (c) maps contours of constant apparent magnitude across the (ρ, A) plane for a diffuse plate in LEO, making the central degeneracy explicit: every point along a contour yields the same observed brightness, so albedo and area cannot be separated by a single measurement. This limitation is sometimes described informally as the spherical-cow problem: an unresolved object is approximated as a Lambertian sphere with a single effective albedo and area. Such a model is useful for order-of-magnitude detectability estimates, but it suppresses the very information that is most diagnostic for RSO characterisation, including panel glints, colour changes, non-convex geometry, and rotational modulation. Panel (e) presents one-dimensional cuts through this same dependence for a diffuse sphere at GEO, varying ρ at fixed A and A at fixed ρ. Panel (d) extends the model to composite bodies i.e., a cylinder-set configuration (bus, dish, and antenna) and a box-wing configuration (bus and solar panels). By summing the component contributions $\rho_i A_i F_i(\Phi)$, this panel shows that plausible spacecraft geometries are of comparable scale differing by only ~1-2 magnitudes across all orbital heights and phase angles. Panel (f) completes the picture with a Monte Carlo distribution of apparent magnitude for the box-wing composite at GEO under randomised albedo, area, and phase angle: the resulting spread of several magnitudes for a single fixed geometry demonstrates that a one-off brightness measurement constrains the product $\rho A F(\Phi)$ only, and taken together, this figure demonstrates the motivation for the time-resolved, multi-colour approach adopted in this thesis.

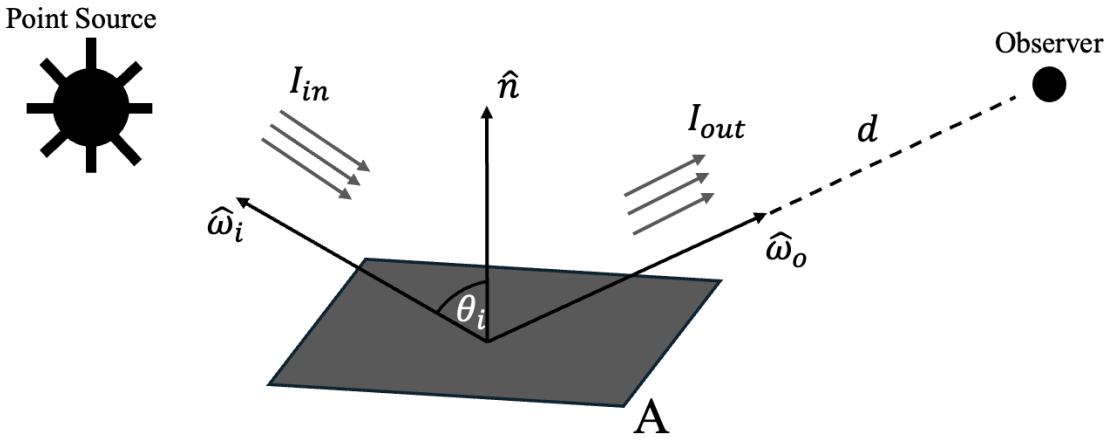

FIGURE 2.7: Illustration of the BRDF scattering geometry, relating the incoming light direction ω_i and outgoing direction ω_o with respect to a surface with normal $\hat{n}$ and area, A. I_{in} and I_{out} refer to incoming flux from a point source and scattered flux towards an observer at a distance, d, away.

2.1.2.2 BIDIRECTIONAL REFLECTANCE DISTRIBUTION FUNCTION

While phase functions give an efficient shape-level approximation, more detailed light-scattering behaviour is captured by the *Bidirectional Reflectance Distribution Function* (BRDF), first formalized by Nicodemus (1965). The BRDF expresses the ratio between the differential reflected

radiance $dL_r(\boldsymbol{\omega_o})$ in the outgoing direction to observer $\boldsymbol{\omega_o}$ and the differential incident irradiance $dE_i(\boldsymbol{\omega_i})$ arriving from direction to source, $\boldsymbol{\omega_i}$:

$$\text{BRDF} = f_r(\boldsymbol{\omega_i}, \boldsymbol{\omega_o}) = \frac{1}{L_i \cos\theta_o \cos\theta_i} \frac{\partial L_o}{\partial \hat{\boldsymbol{\omega}}_i}$$

Here, L_i denotes ingoing radiance (power per unit solid angle per unit projected area), L_o denotes outgoing radiance, θ_i is the angle between $\hat{\boldsymbol{\omega}}_i$ and the surface normal, $\hat{\mathbf{n}}$ and similarly, θ_o is the angle between $\hat{\boldsymbol{\omega}}_o$ and the surface normal. Unlike the idealised models in Table 2.1, the BRDF can capture realistic material effects such as roughness, coatings, and anisotropic scattering behaviour. The geometry is shown in Figure 2.7.

The aforementioned BRDF equation can be rearranged to describe the outgoing radiance:

$$L_o = \int f_r \cos\theta_i \cos\theta_o L_i \, d\boldsymbol{\omega_i} = \int f_r \cos\theta_i \cos\theta_o A I_{in} \delta(\boldsymbol{\omega_i}) \, d\boldsymbol{\omega_i} = f_r \cos\theta_i \cos\theta_o A I_{in}$$

The fraction of flux scattered at the surface from source to an observer is given by:

$$G(\text{Sun} \rightarrow \text{observer}) \equiv \frac{I_{out}}{I_{in}} = (\hat{\boldsymbol{\omega}}_i \cdot \hat{\mathbf{n}})(\hat{\boldsymbol{\omega}}_o \cdot \hat{\mathbf{n}}) f_r(\hat{\boldsymbol{\omega}}_i, \hat{\boldsymbol{\omega}}_o) \frac{A}{d^2}$$

To gather the total flux received by the observer from the satellite, the total is simply summed up over all N surfaces:

$$I_{obs} = I_{Sun} \sum_{s=1}^{N_s} G(\text{Sun} \rightarrow \text{observer})$$

$$G(\text{Sun} \rightarrow \text{observer}) = \frac{A f_r \eta}{||\mathbf{x_{obs}} - \mathbf{x_{sat}}||^2}$$

where:

$$\eta = (\hat{\mathbf{n}} \cdot \hat{\mathbf{v}}_{sat \rightarrow Sun})(\hat{\mathbf{n}} \cdot \hat{\mathbf{v}}_{sat \rightarrow obs})$$

To express the received flux in the AB magnitude system, the model must be converted to a spectral flux density, f_ν. The AB magnitude is then

$$m_{AB} = -2.5 \log_{10}(f_\nu) - 56.10, \tag{2.9}$$

where f_ν is measured in W m^{-2} Hz^{-1}. Equivalently, $m_{AB} = -2.5 \log_{10}(f_\nu / 3631\,\text{Jy})$. For broadband observations, f_ν should be interpreted as the bandpass-averaged spectral flux density rather than the bolometric flux divided by a single frequency.

Phase functions can be interpreted as a surface- and shape-integrated form of the BRDF, representing how much light from the Sun is ultimately reflected into the observer's direction. In other words:

- **Phase functions** — macroscopic, geometry-driven scattering behaviour
- **BRDF** — microscopic, material-dependent scattering behaviour

Thus, phase functions are valuable for rapid brightness prediction, while BRDF-based models enable high-fidelity optical characterisation of spacecraft and space debris as demonstrated in works such as Skuljan (2018); Fankhauser et al. (2023); Lu (2024) and BRDF estimates of materials

(Compean and Small, 2023; Cowardin et al., 2024a; Shah et al., 2024) can be applied to satellite 3D models in ray-tracing simulations as evidenced in the works of Cowardin et al. (2024b); Endo et al. (2023). Formulations/implementation of a selection of BRDF models can be found in Appendix A with `lumos-sat` (Fankhauser et al., 2023) remaining a useful python package for implementing a number of these models in the context of satellite characterisation. The simple diffuse and specular models introduced above can be interpreted as limiting cases of BRDF behaviour. A Lambertian surface has a constant BRDF, $f_r = \rho/\pi$, so the directional dependence of the observed brightness enters only through the illumination and viewing cosine factors. An ideal mirror-like surface instead concentrates the reflected light into a narrow direction satisfying the law of reflection. Real spacecraft materials lie between these limits: solar panels, multilayer insulation, painted surfaces, and antenna reflectors can each exhibit different mixtures of diffuse, specular, and wavelength-dependent response.

2.2 Data Sources

2.2.1 Instrumental and Observational Capabilities Required for RSO Characterisation

Characterisation sits at the centre of current national and international ambitions for SDA. The UK's cross-government SDA requirements (Ministry of Defence and UK Space Agency, 2026), maintained by the Ministry of Defence and UK Space Agency, include the characterisation of RSOs as an explicit user requirement, broken down into capabilities that closely align with the science goals of this thesis: deriving an RSO's effective rotation rate and rotation axis, its attitude, its bus type or class, its history of orbital manoeuvres, and the detection of faults and anomalies in its behaviour. Notably, these characterisation requirements are published only at the capability level, with no performance thresholds attached to them. Consequently, there is at present no formal specification of the photometric precision, cadence, or astrometric precision that an optical system must achieve to deliver these capabilities. The observational requirements must instead be derived from the underlying science and benchmarked against what previous studies have achieved.

Ground-based multi-colour photometry of GEO satellites has an established literature, and its observing modes define a clear capability envelope. Colour studies have generally used narrow-field telescopes tracking one object or cluster at a time, with colours assembled from sequential filter-wheel cycles (Zigo et al., 2023; Jolley et al., 2015; Di Cecco et al., 2023; Schmitt, 2020). Jolley et al. (2015), for example, obtained night-long BVRI light curves of six satellites with a 0.36 m telescope, taking a single image per filter per cycle with exposures of up to 45-60 s in B; the resulting light curves contain of order 200 points per band per night, an effective per-band cadence of two to three minutes, with each colour index built from measurements taken minutes apart. The AGO70 programme (Zigo et al., 2023) follows individual tracked targets through sequential BVRI cycles, deriving colour indices as campaign-level averages, while Schmitt (2020) reports colours of a representative GEO sample binned at particular phase angles. Two consequences follow. First, colours reconstructed from sequential filter observations are contaminated whenever the target varies on the filter-cycle timescale, while the most diagnostic features, such as specular

glints, can evolve on timescales of seconds. Indeed, Jolley et al. (2015) found that colour ratios varied so strongly with illumination and observation geometry that colour-based discrimination proved impractical with such data. Second, per-object tracking limits the sample size: campaigns spanning weeks to years have characterised only a handful to a few tens of objects, meaning that population-level conclusions have necessarily rested on small samples.

The premise of this thesis is that improving on the photometric precision, cadence, and simultaneity of these studies unlocks characterisation that was previously out of reach, and this defines four observational requirements. First, *photometric*: resolving the colour-index differences between bus classes demands per-exposure precision at the percent level, with the colour bands recorded simultaneously so that colours of time-varying sources are meaningful by construction; therefore there is a need to quantify the per-exposure precision each instrument delivers as a function of target brightness and sky conditions. Second, *temporal*: glints and rotational features of GEO targets evolve on timescales of seconds to minutes, setting the required survey cadence, while phase-angle-resolved behaviour requires coverage spanning full nights. Rapidly spinning targets such as Ajisai set a far more extreme requirement: it rotates with a period of ~2.4 s and its mirror glints last only ~10 ms, beyond any practical frame cadence. This cannot be met by fast framing. Streak photometry offers an alternative, in which the target's apparent motion converts position along the trail into time; there is a need to quantify the effective time resolution this provides, and to establish whether it suffices to resolve the *spacing* between successive glints, since it is this spacing, rather than the glint profiles themselves that encodes the rotation period. Third, *astrometric*: associating detections with catalogued objects requires positional accuracy well below the along-track error of propagated TLEs, which reaches tens of arcseconds at GEO within days of epoch (§1.2.1); arcsecond-level astrometry therefore suffices for reliable association, and a measure of whether the instruments can achieve it is required. Fourth, *spatial*: surveying the GEO belt requires a field of view wide enough to capture many orbital slots per pointing, rather than one object per campaign; the usable fraction of that field is limited by the mount, the enclosure and the local horizon, and therefore there is a need to map these visibility limits to establish the practical sky coverage of a survey.

These requirements drove the choice of instruments. STING records four passbands simultaneously through co-mounted cameras, satisfying the simultaneity requirement by construction; whether it meets the precision, cadence and astrometric requirements is an empirical question. The wide-fields that make the survey possible come at the price of a coarse plate scale compared with the narrow-field instruments above, and it is worth being explicit that this does not compromise the measurements. For spatially unresolved targets, photometric performance is determined by the noise budget, including photon statistics, scintillation, and sky background within the photometric aperture, rather than by angular resolution. Large pixels enter this budget only through the sky and read-noise terms, so it is expected that the penalty of coarse sampling is to be confined to the faint limit, where sky dominates, and to be negligible for the bright, scintillation-limited RSOs that form the survey sample. This can be verified by constructing a noise model that accounts for the pixel scale explicitly and comparing it with the measured light curve scatter under both bright- and dark-sky conditions. For astrometry, it is expected that the centroiding can recover positions to a small fraction of a pixel, so that arcsecond-level accuracy should be achievable despite the

coarse sampling; this can be verified by comparing measured positions against an independent, high-accuracy reference such as a GPS satellite ephemeris. The same reasoning applies a fortiori to CLASP: its native sampling is finer than STING's, and the 3×3 binning used for high-speed LEO photometry deliberately trades that angular resolution for the sensitivity and cadence the method requires, so the coarse effective plate scale is a design choice rather than a limitation. For streak photometry the coarse plate scale is expected to work in the method's favour, since with the target crossing pixels in tens of milliseconds the trail itself provides the time resolution that the frame cadence cannot.

These requirements are exercised throughout the thesis: photometric precision and colour simultaneity underpin the bus-class survey of Chapter 4, addressing the type/class requirement of the Ministry of Defence and UK Space Agency (Ministry of Defence and UK Space Agency, 2026); effective cadence and timing underpin the rotation-rate and rotation-axis determinations of Chapter 5 and Chapter 6; astrometric accuracy underpins catalogue identification throughout; and the change-detection requirement is addressed by the spin-period monitoring of Chapter 6.

2.2.2 SuperWASP

The SuperWASP (Super Wide Angle Search for Planets) project (Pollacco et al., 2006) was a ground-based wide-field survey originally designed to detect transiting exoplanets. Operating between 2004 and 2018, the survey discovered and confirmed nearly 200 Hot Jupiter-class planets. It consisted of two custom-built telescope arrays: SuperWASP-North, located on La Palma, and SuperWASP-South, located in South Africa. Each telescope used an array of Canon EF 200 mm f/1.8L USM telephoto lenses coupled to scientific CCD cameras, providing an instantaneous field of view of approximately 500 deg^2 per array.

By 2018, six of the original eight CCD cameras on SuperWASP-North remained operational. The surviving cameras were reconfigured using the available per-camera mount adjustments: four were aligned to form a square $16^\circ \times 16^\circ$ field of view, while the remaining two overlapped the central region to provide redundancy and enable cross-validation between cameras.

Between 2018 and 2022, this reconfigured SuperWASP-North system, hereafter referred to simply as SuperWASP, was used for several observational Space Domain Awareness (SDA) research projects. The instrument in this configuration is shown in Figure 2.8.



FIGURE 2.8: SuperWASP at the Roque de los Muchachos Observatory in August 2019. Image credit: P. Chote.

2.2.3 STING

The SuperWASP-North facility was substantially refurbished into the STING system between late 2022 and early 2023 (Airey et al., 2025a). Located at the Roque de los Muchachos Observatory (ORM) on La Palma, STING (shown in Figure 2.9) is a wide-field camera system featuring four co-mounted lenses that observe simultaneously in NIR (Near Infrared), R, G, and B passbands with an instantaneous field of view of about $10° \times 7.5°$. This upgrade replaced the original CCD cameras with QHY600M-PRO CMOS cameras to achieve an improved plate scale of 3.98 arcseconds per pixel, upgraded the mount to a direct-drive Planewave L350, and relocated the instrument into a 7-foot clamshell dome to eliminate the sky visibility restrictions of the original enclosure. A summary of the specifications of the instrument is displayed in Table 2.2.

The filters currently installed in STING are Baader CMOS R,G,B, and SDSS i. The non-standard RGB filters were selected due to their higher throughput and sharper wavelength transitions than Bessel BVR, and their better match to the sensor *QE* than SDSS ugr. This is particularly important for STING, where the four cameras operate in parallel using the same exposure time and would cause the SDSS u band to have an unacceptably poor signal to noise ratio compared to the other cameras. The RGB filters are also used by the adjacent GOTO (Gravitational-wave Optical Transient Observer) facility (Steeghs et al., 2022), making STING a complementary capability for bright astronomical targets.

The raw STING images are processed using an automated pipeline (Chote et al., 2025) that applies standard image corrections, performs astrometric and photometric calibration on each frame, identifies targets from an input TLE catalogue, and extracts photometric measurements with associated uncertainties.

Parameter	Value
FOV	$10.5° \times 7.5°$
Pixel Scale	3.98”/pixel
Aperture Diameter	11 cm
Readout Noise	2.61 e^-/pixel
Quantum Efficiency	Up to 0.92
Gain	0.43 e^-/ADU
Dark Current Rate	0.05 e^-/pixel/second
Operating Temperature	0°C

TABLE 2.2: Summary table of the specifications of the STING telescope. Here e^- denotes electrons and ADU is an analog-to-digital unit.

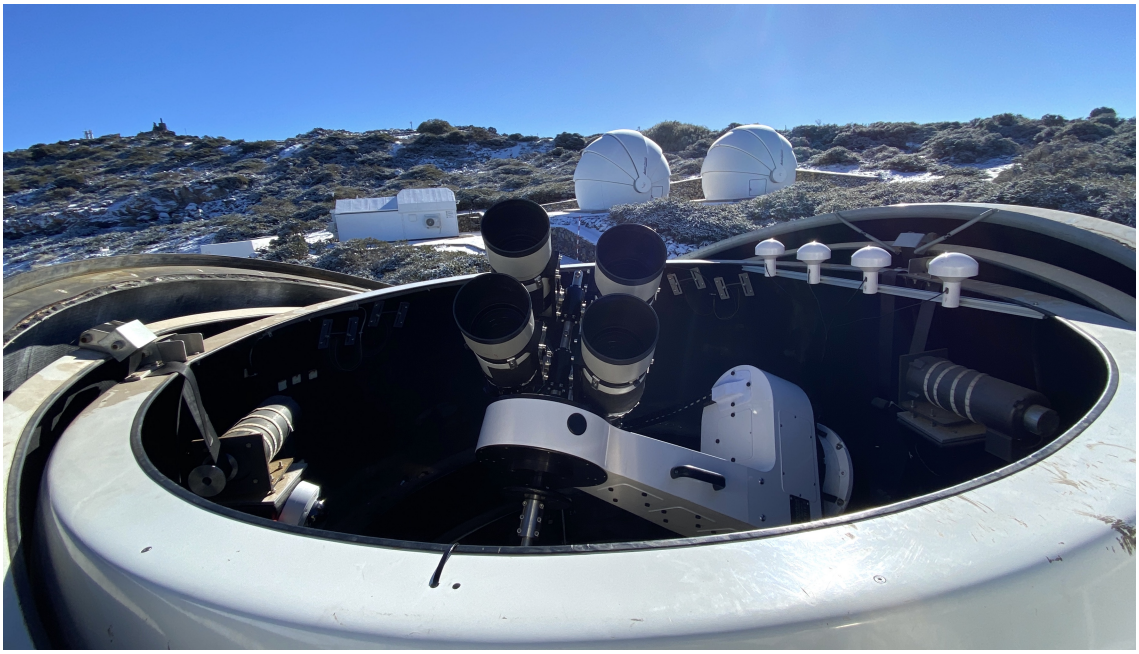

Figure 2.9: A photograph of the STING instrument, at the Roque de los Muchachos observatory on La Palma in the Canary Islands. The original SuperWASP enclosure and the two GOTO-North clamshell domes are visible in the background. Image credit: P. Chote.



Figure 2.10: CLASP at the Roque de los Muchachos observatory on La Palma, in the Canary Islands. Image credit: P. Chote.

2.2.4 CLASP

The CLASP telescope is also situated on the ORM. CLASP consists of two co-parallel RASA 36 cm aperture telescopes on a single Planewave L600 mount (see Figure 2.10). One is mounted with a QHY600PRO sensor and the other with a Princeton 1280SciCam and an 850 nm long-pass filter to reject EO wavelengths and collects in the shortwave infrared (SWIR) band. Further specifications of the instrument are displayed in Table 2.3 and the instrument throughput for both cameras is displayed in Figure 2.11 with comparison to the STING cameras and filters.

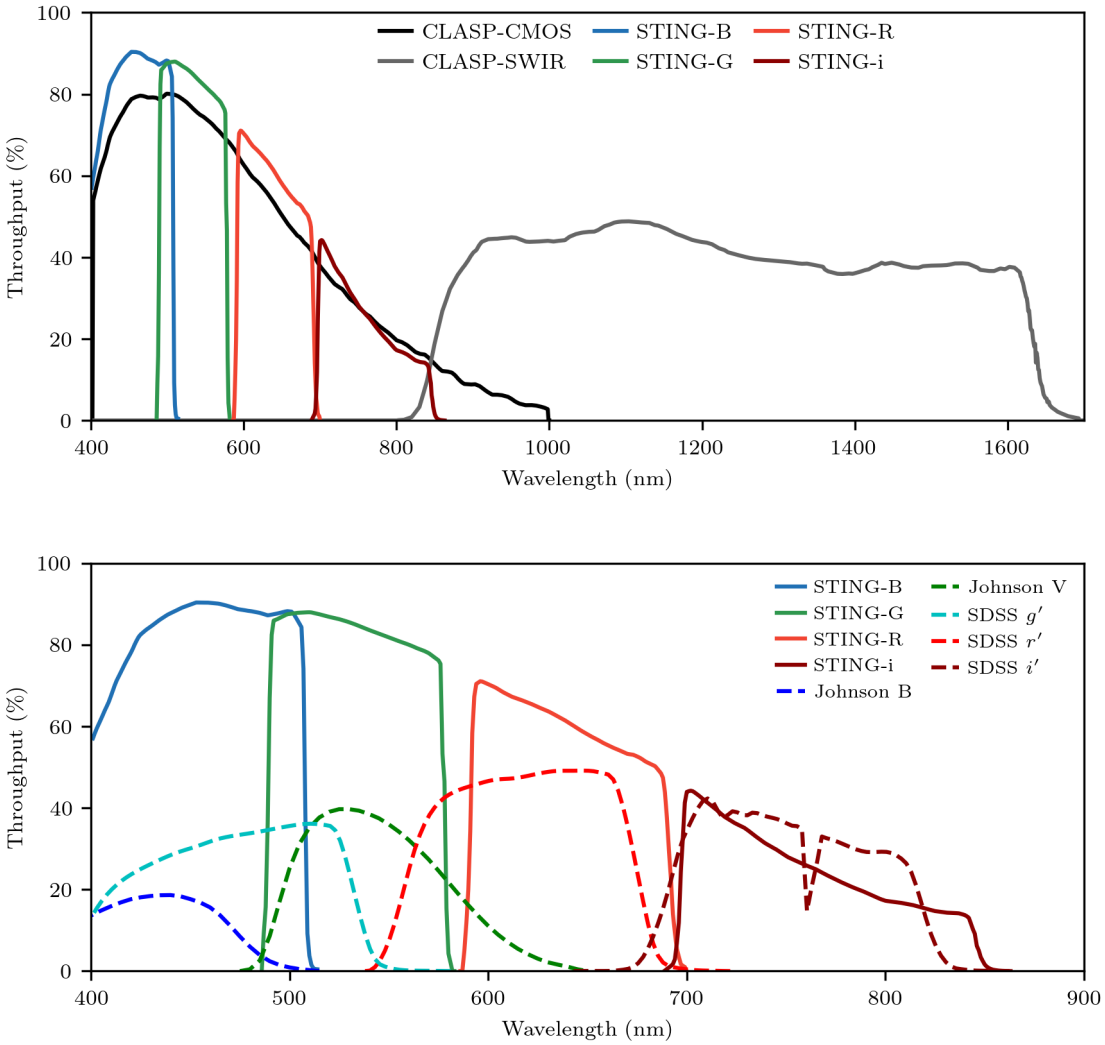

Figure 2.11: Top: Total throughput as a function of wavelength for the cameras and filters on CLASP. For comparison, filters used on STING are shown. **Bottom:** Total throughput as a function of wavelength for the cameras and filters used in STING (solid lines). For comparison, those used in the APASS AAVSO Photometric All Sky Survey (Henden et al., 2015) are shown (dashed lines).

Parameter	QHY600PRO	Princeton 1280SciCam
Telescope Aperture Diameter	36 cm	36 cm
Field of View	2.68° x 1.76°	1.14° x 0.89°
Pixel Width	2.94”/3px (0.98” binned 3x3)	3.12”/px
Passband	400-1000 nm	800-1600 nm
Quantum Efficiency	Up to 0.8	0.75
Bit Depth	16	14
Operating Temperature	-10°C	-60°C

Table 2.3: Summary table of the specifications of the CLASP telescope, reproduced from Macmanus et al. (2025).

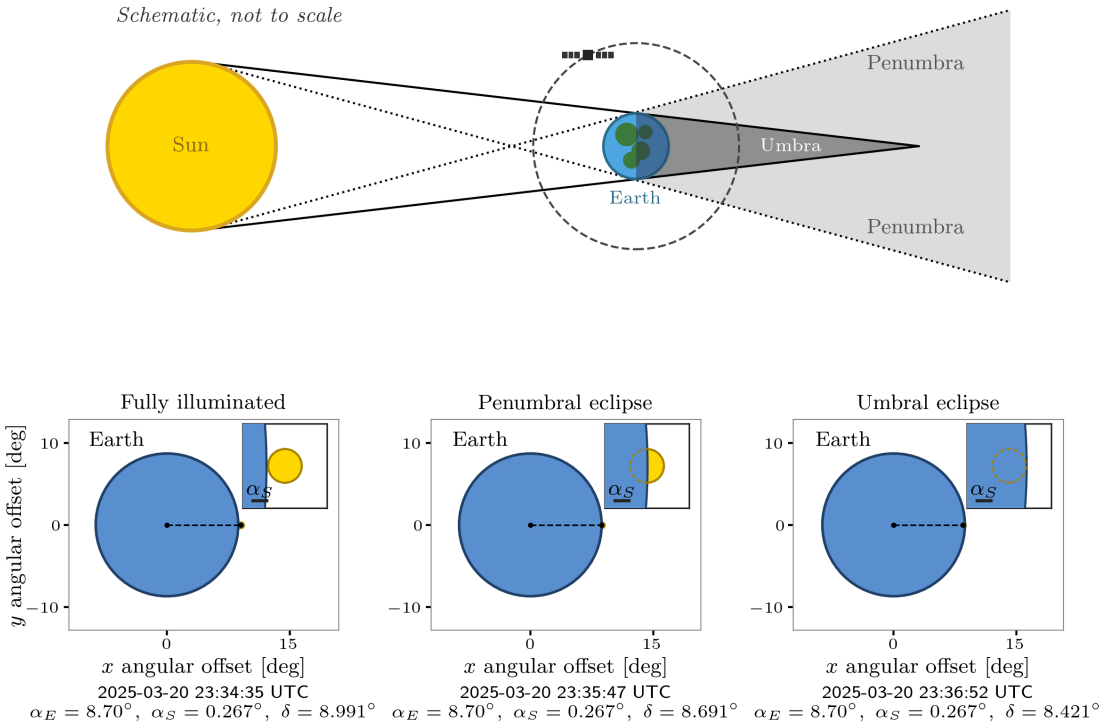

Figure 2.12: Top: Schematic of the Earth shadow geometry, showing the umbra (dark grey) bounded by the inner shadow tangent rays and the surrounding penumbra (light grey) bounded by the outer rays. The dashed circle indicates a representative orbit; the satellite and lunar symbols are illustrative and the diagram is not to scale. **Bottom:** Apparent-disk configurations seen from THOR 6 (GEO, NORAD 36033) at three representative epochs forming a chronological ingress sequence: fully illuminated, penumbral eclipse, and umbral eclipse. The blue disk represents Earth and the gold disk represents the Sun, both at their true apparent angular radii α_E and α_S as seen from the satellite, separated by their apparent centre-to-centre angular distance δ. The dashed line marks the Earth-Sun centre line. Because $\alpha_S \ll \alpha_E$ at these orbital heights, the solar disk is barely resolved at the main-panel scale; an enlarged inset (upper right of each panel) shows the same Sun-Earth limb geometry near the solar disk, with a scale bar of length α_S for reference, and a dashed circle indicating the occulted solar disk position when it lies behind Earth. Epochs are drawn from a propagation interval spanning 2025-03-20 21:00 UTC to 2025-03-21 05:00 UTC.

2.3 Observational Constraints and Strategies

A variety of factors can affect the ability of an optical telescope to observe objects in orbital regimes around the Earth:

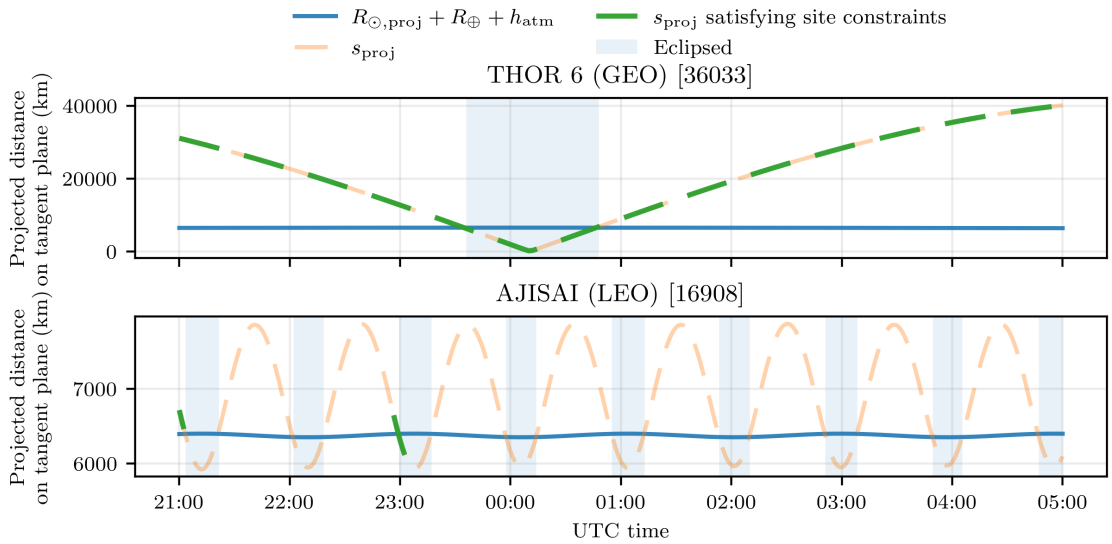

Figure 2.13: Projected separation s_{proj} between the satellite-Earth line and the satellite-Sun line on the plane tangent to Earth at the sub-satellite point, compared with the illumination threshold $R_{\odot,\mathrm{proj}} + R_\oplus + h_{\mathrm{atm}}$, where $R_{\odot,\mathrm{proj}}$ is the projected solar radius at the Earth tangent plane, $R_\oplus$ is Earth's mean equatorial radius, and h_{atm} is the assumed atmospheric grazing height. The satellite is geometrically illuminated when s_{proj} exceeds the threshold. Faint dashed traces show the full s_{proj} time series; bold dashed segments mark intervals during which both the site-elevation ($\geqslant 30°$) and solar-altitude ($\leqslant -12°$) constraints are satisfied at the observatory site. Lightly shaded bands indicate eclipsed intervals ($s_{\mathrm{proj}} <$ threshold); more lightly shaded bands indicate intervals during which the solar-altitude constraint is not met. **Top:** THOR-6 (GEO, NORAD 36033). **Bottom:** Ajisai (LEO, NORAD 16908).

- **Visibility** – As with any astronomical observation, the target must be above the observer's local horizon, whilst the altitude of the Sun $< 0°$. Higher elevation angles are preferable because light from objects near the horizon travels through more atmosphere, increasing airmass and atmospheric extinction (Young, 1994). Poor weather conditions will impact upon the ability to observe a target with cloud in particular affecting the stability of the photometric mean zero-point and if weather conditions are severe enough could pose an unsafe environment to the instrument itself, although precautions are automated in such cases for the instruments on La Palma.

- **Earth Shadow and Viewing Geometry** – For optical observations, the target must be illuminated by the Sun. Objects that pass into the Earth's shadow are eclipsed and are therefore not visible in reflected sunlight. This illumination condition depends on the Sun-Earth-object geometry, illustrated schematically in the top panel of Figure 2.12, while the observer's location determines the local viewing geometry, including whether the object is above the horizon and its apparent position on the sky. The Earth's shadow comprises a fully shadowed umbral cone bounded by the inner shadow tangent rays and a surrounding penumbral region of partial occultation bounded by the outer rays; an object in the umbra receives no direct sunlight, whereas an object in the penumbra sees a fraction of the solar disk occulted by Earth.

  For LEO objects, the shadow's projection covers a large fraction of the night sky at such low orbital heights. Consequently, the majority of LEO targets that are visible to ground-based

optical telescopes can only be observed during twilight periods shortly before sunrise and after sunset, when the observer remains in darkness while the target is still illuminated by the Sun. The Earth subtends an apparent angular radius of $\alpha_E \approx 54^\circ$, so the Sun is occulted by even a small misalignment from the anti-solar direction. At greater orbital heights, the effective cross-sectional radius of the umbra decreases with distance from the Earth; for comparison, the shadow's projection at GEO subtends an angular radius of approximately 10°, consistent with the value $\alpha_E \approx 8.7^\circ$ obtained for THOR 6 (NORAD 36033) in the bottom row of Figure 2.12. For GEO satellites, eclipses occur only during limited seasonal intervals around the equinoxes, when the Sun-Earth-orbit plane geometry causes the Earth's shadow to intersect the GEO belt; the epochs displayed in Figure 2.12 are drawn from one such season, around the March 2025 equinox.

The Sun's apparent size and its apparent offset from Earth as seen from the satellite can be calculated, and then both projected onto the plane at the satellite that is perpendicular to the Sun-satellite line and passes through Earth's centre. The Sun's angular radius as seen from the satellite is $\alpha_\odot = \arcsin(R_\odot / d_{\text{sat–Sun}})$. Projecting this angular radius to a physical radius on the tangent plane gives $R_{\odot,\text{proj}} = \tan(\alpha_\odot)\, d_{\text{sat-Earth}} \cos(\theta_{\text{Earth-Sun}})$, where $d_{\text{sat-Earth}}$ is the satellite-Earth distance and $\theta_{\text{Earth-Sun}}$ is the Earth-Sun angle as seen from the satellite (denoted δ in Figure 2.12). The projected lateral separation between the centres of Earth and Sun on the same plane is $s_{\text{proj}} = d_{\text{sat-Earth}} \sin(\theta_{\text{Earth-Sun}})$. Finally, the satellite is considered fully illuminated (i.e., not in Earth's umbra/penumbra) when this projected separation exceeds the effective blocking radius of Earth (including an atmospheric height margin) plus the projected solar radius: $s_{\text{proj}} > R_{\odot,\text{proj}} + R_\oplus + h_{\text{atm}}$. If the inequality holds, the Sun's disk lies completely outside the Earth shadow cone on the projection plane. The corresponding apparent-disk geometry is shown in the leftmost column of the bottom row of Figure 2.12, where the apparent solar disk lies clear of the Earth limb ($\delta > \alpha_E + \alpha_S$); the middle and right columns show THOR-6 entering the penumbra ($|\alpha_E - \alpha_S| < \delta < \alpha_E + \alpha_S$) and the umbra ($\delta < \alpha_E - \alpha_S$) within tens of seconds, as summarised by the epoch labels under each panel. Note that the apparent solar radius $\alpha_S \approx 0.27^\circ$ is essentially independent of orbital height, while α_E varies by nearly an order of magnitude between GEO and LEO; the inset in each panel shows the limb crossing on an enlarged angular scale, with a scale bar of length α_S for reference.

Figure 2.13 applies the projected criterion to the same two satellites over an eight-hour interval spanning the March 2025 equinox night. For THOR 6 (top), the projected separation s_{proj} remains above the illumination threshold for almost the entire window, dropping below it only during a narrow eclipse near local midnight UTC; for AJISAI (bottom), the rapid orbital motion produces multiple short eclipses, and the bold dashed segments — those satisfying both the local site-elevation ($\geqslant 30^\circ$) and solar-altitude ($\leqslant -12^\circ$) constraints at the observatory — capture only a small subset of the illuminated arcs.

Observations of high-altitude objects are therefore most favourable when conducted near the longitude of the anti-solar point while remaining outside the shadowed region. By minimising the solar phase angle under these conditions, the apparent brightness of the target

is maximised, as a greater fraction of incident sunlight is reflected toward the observer along the line of sight.

- **Galactic Plane** – Observations of crowded stellar fields such as those near the Galactic plane increase source blending within the FOV. This can adversely affect target detection, centroiding, and background subtraction and ultimately affect the derived astrometric solution.
- **Moon and Bright Sources** – Light reflected from the Moon increases the sky background level, reducing the effectiveness of detecting faint objects in an image frame. More importantly, it increases the noise contribution from the background: while background signal can be subtracted during processing, the associated noise cannot. As a result, faint sources are more likely to be lost within the noise floor. The lunar contribution is strongly phase and direction dependent: the sky brightness rises steeply toward full Moon and with decreasing angular distance from it, following the scattering behaviour described by sky-brightness models (Krisciunas and Schaefer, 1991). When the Moon is above the horizon it is typically the dominant background source, and it therefore enters scheduling directly: target fields are chosen to maintain a sizeable angular separation from the Moon, and the brightest lunar phases are reserved for correspondingly bright targets. Other bright sources, such as planets or bright stars, can also saturate parts of a raw frame, making target detection and extraction more difficult.
- **Telescope/Enclosure Limits** – Further observing constraints can be caused by the physical limits of the telescope and enclosure it may lie within. Ultimately these can be optimised, for example, with an instrument with a multi-camera setup, it makes the most sense to point the most southern camera array towards the northern fields and vice-versa to increase observation arcs. Horizon limits are are therefore important to be empirically calculated, as is done within this thesis.

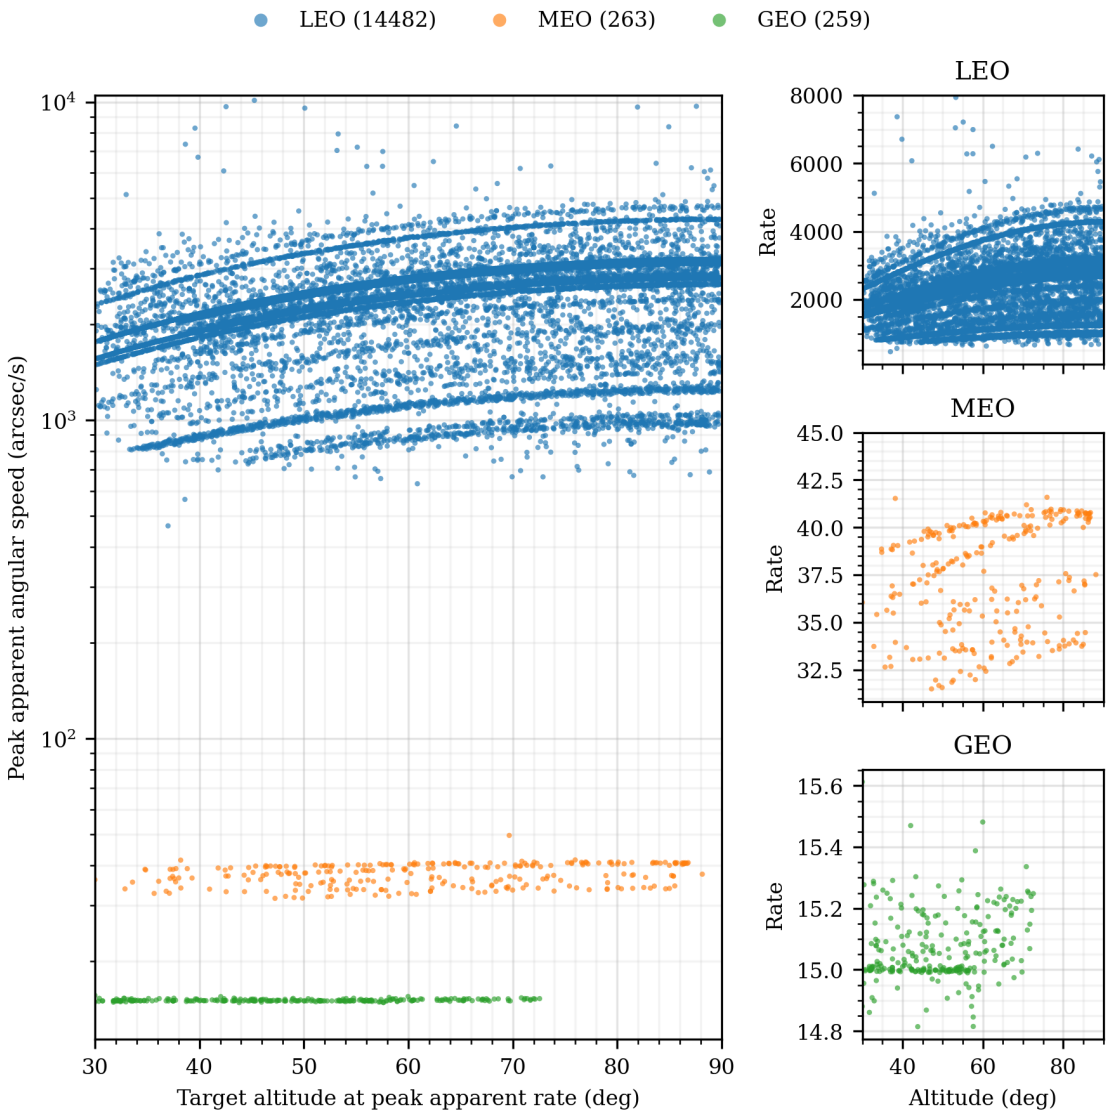

Figure 2.14: Peak apparent angular rates of catalogued satellites as seen from the ORM site, and evaluating the apparent angular rate of the observer-satellite line of sight at each time sample. For each object the peak is taken over a 24 h window beginning at its own TLE epoch, restricted to samples above 30° elevation; TLE epochs span 2026 February 11 to March 13, so the figure characterises the population rather than a single night. The left panel shows the full catalogue on a logarithmic scale, with object counts per orbital regime given in the legend; the right panels provide linear-scale views for the LEO, MEO, and GEO populations.

The apparent angular velocity of Resident Space Objects (RSOs) as seen from the ground varies significantly across orbital regimes, with direct implications for how these objects are detected and tracked. These angular rates, typically expressed in arcseconds per second ($''\,s^{-1}$), ultimately inform exposure time choices, telescope tracking modes, and data reduction techniques to maximise the SNR. The peak apparent angular rates of catalogued satellites visible from the ORM are shown in Figure 2.14, illustrating the distinct motion regimes of LEO, MEO, and GEO objects.

LEO objects move extremely fast across the sky. For example, Nomen Torres et al. (2017) show that a 1000 km circular LEO crossing zenith can attain ~1500$''\,s^{-1}$ angular rate, falling to only ~200$''\,s^{-1}$ near the horizon. Even lower orbits move faster: a 300 km LEO already reaches ~1500$''\,s^{-1}$ at just 20° elevation, and Envisat (~790 km) briefly exceeded 2500$''\,s^{-1}$ at zenith (Nomen Torres et al., 2017). To handle such motion, LEO surveys use very short exposures and rapid tracking. In practice, exposure times are kept as short as possible (on the order of 0.01–0.1 s) to minimise trailing (Kouprianov, 2008; NASA Orbital Debris Program Office, 2023). Telescopes usually slew along predicted orbits to keep the object in frame, acquiring many frames per pass. For example, Nomen Torres et al. (2017) found optimal exposures of ~0.03 s for their EMCCD (Electron Multiplying Charge-Coupled Device) systems. The spread of angular rates in LEO is primarily a spread in orbital height: near culmination the apparent rate is approximately the orbital speed divided by the slant range, so lower orbits appear faster both because they move faster and because they are closer. The dense bands correspond to large constellations occupying discrete shells (i.e., Starlink at ~550 km, Iridium at ~780 km, OneWeb at ~1200 km); each traces a curve rather than a clump because the slant range, and hence the rate, depends on the elevation at which a satellite's best pass within the search window culminates. The diffuse points between the bands are individual satellites at intermediate heights, and the highest rates arise from objects at $\lesssim$300 km that are recently launched or decaying.

MEO satellites (e.g., GPS at ~20,000 km) have much slower apparent speeds. Typical observed MEO rates are of order tens of arcseconds per second. Żołnowski et al. (2021) report ~45$''\,s^{-1}$ for a GPS satellite during a 1 s exposure. Because MEO objects still move noticeably, exposures are still short but usually longer than in LEO (typically ~1 s), and the mount is guided at the predicted satellite rate. In practice, a 1 s integration yields manageable trailing, allowing accurate astrometry and photometry. While some advanced pointing strategies have been explored (e.g., structured dithering patterns), most tracking approaches aim to minimise drift within each exposure frame.

Although GEO objects appear nearly stationary to an observer on Earth, they move relative to the background stars at roughly the sidereal rate, ~ 15$''\,s^{-1}$, because the stars sweep across the observer's sky with Earth's rotation. A perfectly station-kept GEO satellite would maintain a fixed position in the rotating Earth frame, but small inclination, eccentricity, and orbital drift introduce additional apparent motion in right ascension and declination. Seitzer et al. (2017) report typical component rates of $|\Delta\alpha| \lesssim 2''\,s^{-1}$ and $|\Delta\delta| \lesssim 5''\,s^{-1}$ arising from these perturbations. Because observers commonly guide at the GEO rate or disable sidereal tracking entirely, the residual motion relative to the detector is small, allowing relatively long exposures (e.g., ~ 5 s) without significant trailing. The MODEST (Michigan Orbital Debris Survey Telescope) system, for example, routinely used 5 s R-band exposures to detect GEO objects down to 17.5 mag (NASA Orbital Debris Program Office, 2023).

2.4 Data Reduction

Optical measurements can provide a wealth of information about an RSO's physical properties. Numerous reduction steps are required to transform raw, uncalibrated image frames into high-precision photometric light curves, which are utilised throughout this thesis. In practice, these reductions are performed as part of the existing automated processing pipeline; however, they are described here for completeness. These steps are outlined in the following subsections.

2.4.1 Instrumentation

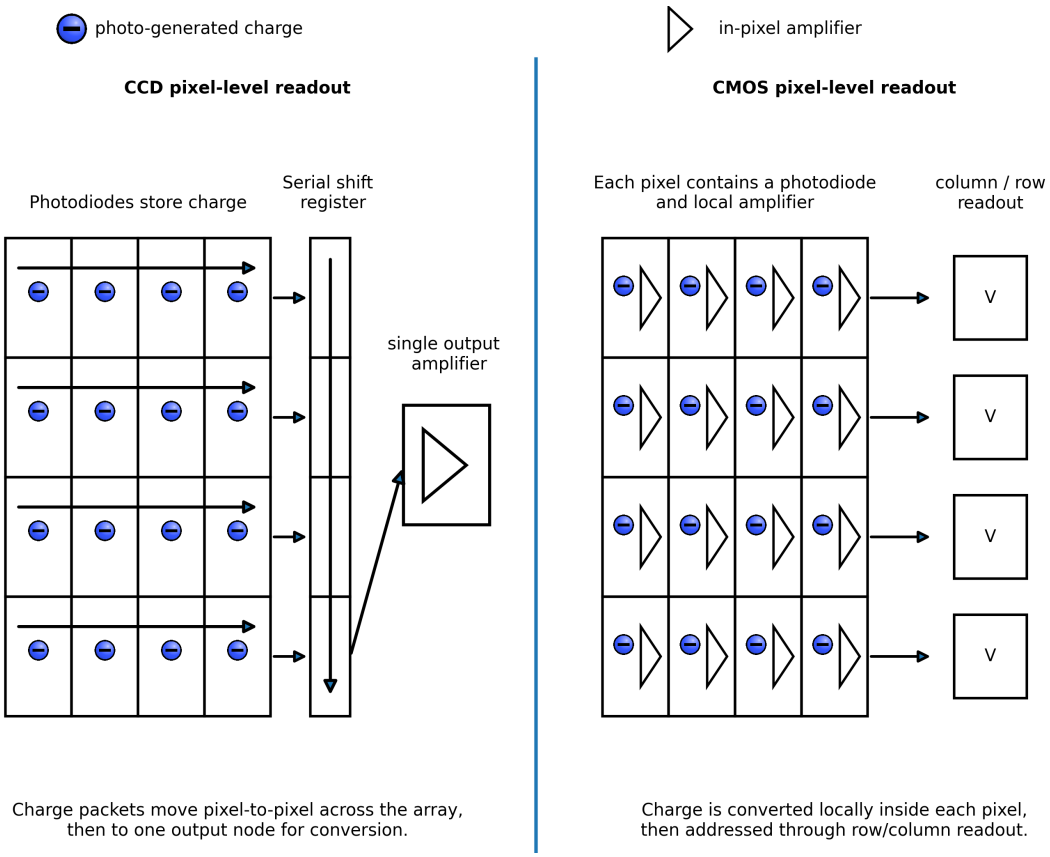

Figure 2.15: Comparison of pixel-level readout architectures in CCD and CMOS image sensors. In CCDs, photo-generated charge is shifted pixel-to-pixel to a single output amplifier, whereas in CMOS sensors, charge is converted locally within each pixel and read out through row/column addressing.

The majority of the analysis in this thesis took advantage of observations taken with CMOS sensors, which have become a more viable route for high-precision photometry for astronomical applications (see e.g., Karpov et al., 2020; Alarcon et al., 2023; Apergis et al., 2025). Among the datasets considered in this thesis, CCD sensors were used only for the historical SuperWASP observations of Ajisai (see Chapter 5).

CCD-based systems have been complemented by a broad range of optical configurations, spanning narrow-field instruments optimised for deep searches and wide-field cameras designed to maximise sky coverage and detection rates. More recent studies have also demonstrated the effectiveness of smaller-aperture telescopes equipped with modern, low-noise detectors for high-

cadence photometric monitoring and characterisation. Although CCDs have historically dominated astronomical imaging, CMOS detectors have become increasingly prevalent in astronomical instrumentation.

At the detector level, CCD and CMOS sensors differ primarily in how charge is transferred and read out from the array (see Figure 2.15). In CCD detectors, photons generate charge within each pixel via the photo-electric effect (Einstein, 1905). This charge is accumulated during the exposure and then transferred sequentially across the detector, pixel by pixel, to one or a small number of output amplifiers, where it is converted into a voltage and digitised (McLean, 2008). This serial charge-transfer architecture has historically provided excellent pixel-to-pixel uniformity, high sensitivity, and very low read noise, making CCDs particularly well suited to faint-object astronomical imaging. CCDs are also commonly associated with global-shutter operation, in which the full array is exposed simultaneously.

In CMOS detectors, by contrast, each pixel, or in some designs each column, contains its own amplifier and readout electronics. This active-pixel architecture allows the light-generated charge to be converted locally into a voltage and read out in parallel across the detector. As a result, CMOS sensors can provide faster frame rates, pixel-level addressing, and lower power consumption than CCDs, making them attractive for high-cadence observations. Typical CCD sensors have readout noise levels ranging from 5 to 20 electrons per pixel. In comparison, modern CMOS sensors typically achieve readout noise levels in the 1 to 3 electrons per pixel range (George, 2024). CMOS sensors can also exhibit more complex detector behaviour due to their distributed readout electronics, and many are commonly operated using rolling-shutter readout (39.027 µs per image row in our case), although global-shutter CMOS implementations are available in more advanced designs. The studies outlined in Chapter 4 and Chapter 6 demonstrate the capabilities of CMOS sensors for the characterisation of RSOs.

2.4.2 STANDARD ASTRONOMICAL CALIBRATIONS

Astronomical images contain both astrophysical signal and detector-dependent components, so accurate calibration requires an understanding of the detector architecture. The distinction between CCD and CMOS readout is important because detector-dependent effects, such as bias structure, dark current, read noise, non-linearity, and pixel-to-pixel sensitivity variations, influence the calibration procedures required to recover accurate astronomical signals.

Three types of calibration frames are traditionally used in optical astronomy (Howell, 2006):

- **Bias frames** measure the additive electronic offset introduced during detector readout. In traditional shuttered CCD systems these are often obtained as nominally zero-exposure, closed-shutter images. For shutterless CMOS systems, a true closed-shutter zero-exposure bias frame is not available in the same sense, so the electronic offset is instead treated through matched detector calibration frames, such as dark or dark-flat frames, taken with the same camera configuration as the science data.

- **Dark frames** are unilluminated exposures obtained with the same detector configuration as the science frames. They contain thermally generated charge as well as the additive electronic

offset. For the cooled CMOS cameras and short exposure times used in the observations analysed in this thesis, the thermal dark-current contribution is generally negligible, but matched dark frames remain useful for removing the electronic offset and fixed-pattern detector structure, as well as identifying hot/defective pixels that have abnormally high dark current. In the reduction pipeline used in this thesis, the electronic offset is instead carried within the master dark frame. A single master dark frame per camera together with its median bias level are constructed, and the dark contribution to a frame of exposure time t is reconstructed and thus no separate bias frames are required to calibrate the science images.

- **Flat-field frames** are uniformly illuminated images used to measure and correct for pixel-to-pixel variations and non-uniformity across the pixels arising from the optical path and detector response. Dividing by a master flat corrects for these multiplicative effects. Flat fields are acquired by taking exposures of a uniform light source which in the case of our observations is the sky during twilight at dusk and dawn. Over the $10.5^\circ \times 7.5^\circ$ STING field of view, the twilight sky is not uniform, so a median-combined sky flat retains some residual large-scale illumination gradient. Its effect is bounded using the calibration stars: the scatter of the individual stellar zero-points shows no significant trend against detector position. Because targets are tracked and remain near a fixed detector position during an observation, any residual flat-field error would enter primarily as a zero-point offset rather than as a time-dependent signal in the light curve.

In CCD-based photometry, calibration frames are traditionally applied independently: a master bias is subtracted from dark, flat, and science frames prior to dark subtraction and flat-fielding. This approach relies on the high linearity and temporal stability of CCD bias signals at very short exposure times (see e.g., Janesick, 2001; Howell, 2006). In contrast, CMOS sensors exhibit non-linear behaviour at short exposures and are often affected by amplifier glow, rendering dedicated bias frames less reliable (see e.g., Apergis et al., 2025; Andor Technology / Oxford Instruments, 2025). However, it is important to note, that the CMOS sensors used in this thesis do not show amplifier glow at a measurable level. In short, this effect can be tested for by acquiring a long exposure dark frame, which can capture the glow structure. The dark signal can be then subtracted thus only leaving the glow contribution (Apergis et al., 2025).

Importantly, dark frames inherently include both the thermal signal and the electronic offset normally isolated in a bias frame, while flat-field frames similarly contain the bias component. This is true for both CCD and CMOS detectors. As a result, the bias signal is correctly removed when science frames are calibrated using corresponding dark frames, and flat-field frames using matching dark-flats where required (PixInsight Development Team, 2022). Modern CMOS detectors further differ from CCDs in offering substantially lower read noise, enabled by per-pixel amplification and highly parallel readout architectures, with typical values of only a few electrons or below in contemporary astronomical devices (see e.g., Janesick, 2007; Andor Technology / Oxford Instruments, 2023, 2025). Following combination of nightly dark and flat calibration data with rejection of outliers, each science frame R is calibrated as

$$S = \frac{R - D(t_{exp})}{F}, \tag{2.10}$$

where $D(t_{exp})$ is the scaled master dark frame and F is the normalised master flat, constructed from the twilight frames of each night by dark-correcting each frame in the same way, normalising it by its sigma-clipped mean over the imaging region, and median-combining the stack. The variance of each calibrated pixel is modelled as:

$$\sigma_i^2 = \sigma_{bg,i}^2 + S_i, \tag{2.11}$$

where $\sigma_{bg,i}$ is the local background RMS.

A bad-pixel map is used to identify unstable, saturated, or hot pixels, which are subsequently excluded from astrometric and photometric extraction. For both CCD and CMOS cameras, calibration frames should be taken using the same gain, offset, binning, and readout mode settings as the science frames (American Association of Variable Star Observers, 2022).

2.4.3 ASTROMETRIC AND PHOTOMETRIC CALIBRATIONS

Following the instrumental calibration described above, the reduction pipeline (Chote et al., 2025) establishes an astrometric and photometric reference frame for each science image.

A spatially varying background model is estimated for each frame to account for sky brightness gradients, scattered light, and residual instrumental signatures. The background is evaluated on a coarse grid of 256×256 pixels across the detector and interpolated to produce a smooth two-dimensional model, which is subtracted from the calibrated image (Bertin and Arnouts, 1996; Barbary, 2016). Sources are rejected from each grid cell by iterative 3σ clipping, and the grid is median-filtered before bicubic-spline interpolation. This procedure preserves compact astrophysical/satellite sources while suppressing low-frequency background variations. The background estimation also yields a per-pixel noise model, from which the local background root-mean-square (RMS) is computed and used to define detection thresholds and photometric uncertainties.

Astrometric calibration establishes the mapping between detector pixel coordinates (x, y) and celestial coordinates (α, δ), including optical distortion, which is modelled as a third-order Simple Imaging Polynomial (SIP) correction to the tangent-plane projection. When available, a blind astrometric solution is derived with `astrometry.net` (Lang, 2009) by matching detected stellar sources to an external reference catalogue (APASS for STING), and is accepted if at least ten stars agree in position to within 6 pixels and in brightness to within 3σ. Otherwise if this fails, an external astrometric solution is interpolated across the observing sequence by fitting the solutions of neighbouring frames with 3σ outlier rejection, to account for small temporal variations in pointing and optical distortion. The STING instrument used throughout this thesis has a wide FOV, sometimes this means that images are partially blocked by terrain or the enclosure. To mitigate the impact of this, the astrometric solution is constrained using multiple subregions of each frame (4×4 tiles) calculated independently, and are subsequently combined into a global transformation by fitting the SIP model to the matched stars of all successful tiles, so that obscured tiles are simply ignored. This significantly improves the robustness of the calibration (Chote et al., 2019).

The resulting world coordinate system (WCS) allows the apparent motion of each RSO or the background stars during an exposure to be quantified. For a source with pixel positions (x_1, y_1) and (x_2, y_2) corresponding to the start and end of an exposure, the apparent trail length L and orientation

θ are given by

$$L = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2},$$
$$\theta = \arctan\left(\frac{y_2 - y_1}{x_2 - x_1}\right).$$

These quantities are used to inform both source detection and photometric aperture selection.

Photometric calibration is performed on a per-frame basis using field stars with known catalogue magnitudes. When the mount follows a target the field stars trail by the length L defined above; stars trailed by more than 3 pixels are measured through rectangular apertures of length $L + 10$ px and width 10 pixels aligned with θ, otherwise through circular apertures, and stars with a catalogue neighbour within 2.5 mag inside $45''$ of the trail are excluded. Instrumental magnitudes are defined as

$$m_{inst} = -2.5 \log_{10}\left(\frac{F}{t_{exp}}\right),$$

where F is the background-subtracted aperture flux in detector counts and t_{exp} is the exposure time. The photometric zero-point for each reference star is then

$$ZP_i = m_{cat,i} - m_{inst,i}.$$

To account for colour-dependent differences between the instrumental passband and the catalogue system, a polynomial model is fitted to the zero-point as a function of stellar colour,

$$ZP(C) = \sum_{k=0}^{n} a_k C^k,$$

where C is a catalogue colour index and a_k are fitted coefficients. The fitted zero-point model and its associated scatter provide a quantitative measure of frame-to-frame photometric stability.

2.4.4 TARGET DETECTION AND EXTRACTION

Source detection is performed on calibrated, background-subtracted images using a threshold-based segmentation approach (Bertin and Arnouts, 1996; Barbary, 2016). Objects are identified as contiguous groups of at least 9 pixels exceeding a 3σ significance level above the local background noise (raised stepwise to at most 9σ where de-blending fails in crowded fields). Detection thresholds are scaled by the estimated background RMS to ensure uniform sensitivity across frames obtained under varying sky conditions.

Imaging of satellites can differ from conventional imaging of stars in that RSOs which are tracked in a sidereal manner will appear as elongated trails rather than point sources. In the case of the data acquisition in the work described in this thesis, most of the RSOs observed were point sources with the stars trailed by the length L and angle θ. In any case, the detection process will therefore need to account for the expected morphology of each source. When the apparent trail length during an exposure is small compared to the point-spread function ($L < 3$ px), objects are treated as point-like. When the apparent motion is significant, the detection algorithm incorporates an oriented matched filter , consisting of a rectangle of length $0.75L$ and unit width at angle θ, smoothed with a 1-pixel Gaussian, aligned with the predicted trail direction, with detection performed at 2σ. This enhances sensitivity to elongated sources while suppressing false detections

due to noise fluctuations or background residuals. Candidates are further required to have a uniform flux profile along the trail, which rejects hot pixels and point-like satellites from the stellar catalogue.

Detected sources are characterised by their centroid positions, integrated flux estimates, and shape parameters; their positional accuracy is given by the astrometric residuals described in the prior subsection and their flux uncertainties by Equation 2.11. These measurements are used both to support astrometric calibration and to construct a reference catalogue of sources in the field.

Targets themselves are detected in short stacks aligned on their predicted motion. Each target's TLE is propagated to the mid-time of every frame and projected through the frame WCS, and groups of 15 consecutive frames are median-combined on these positions, suppressing the trailed stars while reinforcing the target. Sources are extracted from the stack at 4σ, and the detection closest to the predicted position is assigned to the target, carrying its NORAD catalogue identifier. The detections are then fitted per night as low-order polynomials in the offset from the TLE, with outlier rejection and eclipse intervals excluded; the fitted track defines the aperture positions used onwards.

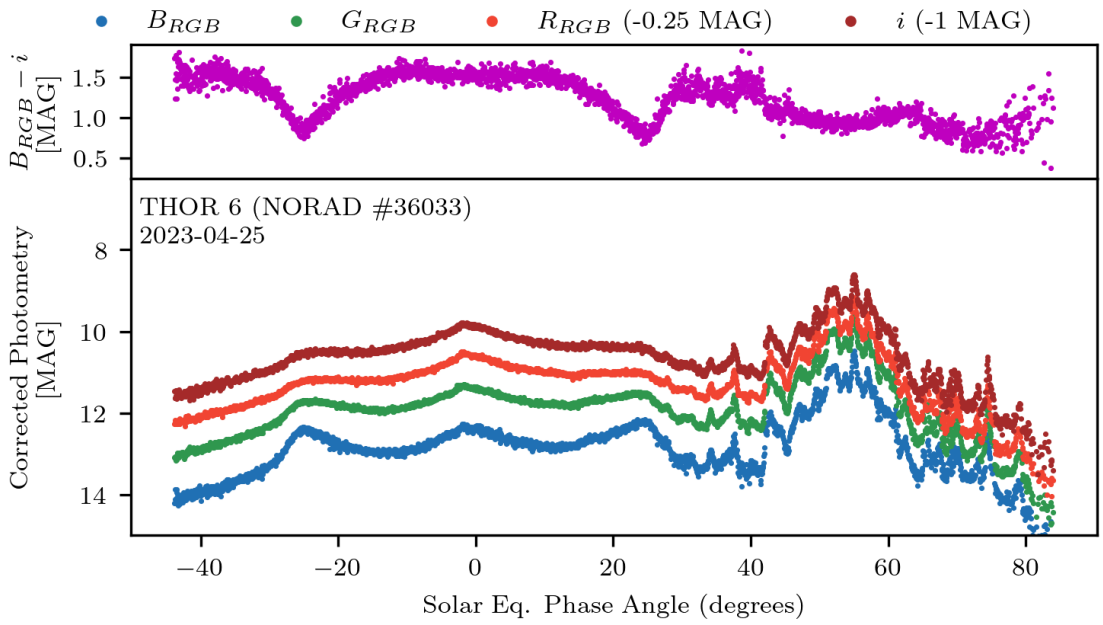

Figure 2.16: Illustrative example of a multi-colour calibrated STING light curve as a function of solar equatorial phase angle, with the top panel showing the $B_{RGB} - i$ colour index. The offsets applied to the R_{RGB} and i bands are to provide a visual clarity and separation between each light curve.

2.4.5 Aperture photometry and light curve extraction

Photometric measurements are obtained using aperture photometry applied to each calibrated frame. Target positions are determined by projecting celestial coordinates into detector space using the frame-specific astrometric solution. The apparent motion of each RSO during the exposure is used to classify the source as either point-like or trailed.

For point-like sources, circular apertures of fixed radius r are employed. For trailed sources,

rectangular apertures aligned with the direction of motion are used, with length

$$\ell = L + 2r \tag{2.12}$$

and width $2r$, where L is here the trail length of the target. Since the trail is the point spread function convolved with a line of length L, this encloses the same fraction of the flux as a circular aperture of radius r does for a point source, ensuring that the full flux of the trail is captured. Photometry is performed for a range of aperture radii ($r = 2$-6 pixels for STING) to assess the trade-off between signal-to-noise ratio and background contamination. Following Chote et al. (2025), the radius adopted is the one that minimises the scatter of the light curve: smaller apertures suit faint targets, where sky and read noise dominate, while larger apertures protect bright targets against seeing and focus variations. For the STING light curves analysed in this thesis $r = 5$ pixels was adopted typically throughout. The offset introduced by measuring the target through a different aperture from the calibration stars was found in Chote et al. (2025) to be $\lesssim 0.05$ mag and is neglected. If a bad pixel (pixels flagged as unreliable or affected by detector defects) is found to be within the aperture, the measurement is flagged and excluded from the light curve. A masked pixel biases the aperture sum by an unknown amount, and interpolating across it would require a model of the source profile that is not available for a target of rapidly varying brightness; given the high observing cadence, discarding the measurement is preferable to correcting it.

Photometric uncertainties are estimated by propagating the per-pixel variance of Equation 2.11 through the aperture,

$$\sigma_F = \sqrt{F_{ap} + A\sigma_{bg}^2}, \tag{2.13}$$

where F_{ap} is the background-subtracted aperture flux, A is the aperture area in pixels (πr^2 or $2r\ell$), and σ_{bg} is the background RMS per pixel, measured locally around the target.

Referring again to Chote et al. (2025) (see Figure 6 within this study), measurements contaminated by background stars streaking through the aperture were identified by placing a circular aperture at the target position in the target frame, and at the same RA/Dec coordinates in two selected reference images: the nearest preceding and following frames from the same camera in which the target has moved far enough that the two apertures no longer overlap, so that the reference aperture contains only the background stars. For each measurement, the smaller of the two reference-image aperture fluxes was adopted as the estimated contaminating flux, guarding against a reference on one side being contaminated by a different RSO.

$$f_{blend} = \frac{F_{ref}}{F_{target} + F_{ref}},$$

and only measurements with $f_{blend} < 0.05$ were retained.

The remaining measurements are combined to form time-resolved light curves for each RSO. These light curves form the primary observational dataset used in subsequent analyses of rotational variability, attitude dynamics, and surface scattering properties. Figure 2.16 displays an illustrative example of an RSO light curve. This particular light curve displays features which will be explored later within this thesis such as the strong blue colour features at $\sim \pm 25°$, the peak around $0°$ and the short timescale glinting regime between $\sim 40° - 80°$.

2.5 SIMULATING RSO LIGHT CURVES

The optical brightness of an unresolved resident space object (RSO) depends on a combination of illumination geometry, observer range, object shape, surface reflectance properties, and attitude state. Because none of these can be measured directly from an unresolved light curve, a forward model is needed to understand which features of an observed light curve can be attributed to which physical cause. The simulation described here serves two purposes in this thesis: it demonstrates how the attitude state alone alters the light curve morphology and it provides a noise-free light curve with a known rotation period against which period-determination methods can be tested against. To illustrate how these effects combine, a simplified forward model was constructed to simulate light curves from a box-wing satellite geometry. The purpose of this simulation is illustrative: it is not intended to provide a validated physical model of a specific spacecraft, but rather to demonstrate how different rotational states and surface-scattering assumptions can produce qualitatively different unresolved light curves.

The simulation consists of four main steps. First, the satellite, observer, and Sun positions are propagated to obtain the time-dependent illumination and viewing geometry. Second, a simplified faceted spacecraft model is defined in the body frame. Third, the attitude of the body is propagated using either a uniform-spin model or a torque-free tumbling model. Finally, the reflected flux from all illuminated and visible facets is computed using a Cook-Torrance BRDF (Cook and Torrance, 1982) before being converted into an illustrative magnitude.

The full mathematical form of this BRDF, including the GGX normal-distribution function, Schlick Fresnel approximation, geometric attenuation term, and half-vector definition, is given in Appendix A. The important distinction is that each surface facet is assigned two optical parameters: a roughness parameter, α, and a normal-incidence Fresnel reflectance, F_0. The roughness parameter controls the angular width of the specular response. Smaller values of α represent smoother, more mirror-like surfaces and can produce narrower glints when the Sun-facet-observer geometry approaches the specular condition. Larger values of α produce broader and weaker reflections. The parameter F_0 controls the strength of the specular component at normal incidence, with larger values increasing the relative contribution of specular reflection.

2.5.1 ORBITAL AND OBSERVING GEOMETRY

A representative geostationary observing geometry was generated by propagating a two-line element (TLE) set for THOR 6 using `skyfield` (Rhodes, 2019). The observer was placed at the Roque de los Muchachos Observatory (ORM) on La Palma, with latitude 28.76° N, longitude 17.89° W, and altitude 2300 m. The satellite was propagated over a two-hour observing window with a cadence of $\Delta t = 0.1$ s.

At each time step, three positions were calculated in an ECI frame: the satellite position, the observer position, and the solar position. These were used to define the unit vector from the satellite to the Sun,

$$\hat{\boldsymbol{\omega}}_i(t) = \frac{\vec{\mathbf{r}}_\odot(t) - \vec{\mathbf{r}}_{sat}(t)}{\left|\vec{\mathbf{r}}_\odot(t) - \vec{\mathbf{r}}_{sat}(t)\right|}, \tag{2.14}$$

and the unit vector from the satellite to the observer,

$$\hat{\omega}_o(t) = \frac{\vec{r}_{obs}(t) - \vec{r}_{sat}(t)}{\left|\vec{r}_{obs}(t) - \vec{r}_{sat}(t)\right|}. \tag{2.15}$$

The observer range was calculated as

$$d(t) = \left|\vec{r}_{obs}(t) - \vec{r}_{sat}(t)\right|. \tag{2.16}$$

Earth-shadow intervals were identified using the `skyfield` illumination test. Time steps for which the satellite was not sunlit were excluded from the reflected-flux calculation. This ensured that the simulated light curve only included epochs for which the object was illuminated by the Sun.

2.5.2 SIMPLIFIED SPACECRAFT GEOMETRY

The satellite was represented by a simplified box-wing geometry consisting of a cubic bus and two segmented solar-panel wings. The bus was modelled as six square facets, each with a fixed body-frame normal and area. The solar panels were represented by planar segments extending from either side of the bus. Each facet was assigned a surface normal, area, roughness parameter, and Fresnel reflectance value.

The body-frame surface model is deliberately simpler than a true spacecraft model. In particular, small structures such as antennas, multi-layer insulation, panel hinges, and detailed bus protrusions are not included in the photometric calculation. However, the model retains the most important geometric distinction for this demonstration: compact bus-like facets and extended panel-like facets. This is sufficient to illustrate how surface orientation and attitude state affect the observed brightness of an unresolved object.

The bus and panels were assigned separate Cook-Torrance material parameters. The bus was given a larger roughness value, representing a broader and less mirror-like scattering response. The panels were assigned a smaller roughness value, allowing sharper glint-like features to occur when the panel normals approach the specular reflection geometry.

The parameters used for the illustrative simulation are summarised in Table 2.4. The chosen values were selected to produce a representative box-wing satellite light curve with both broad brightness variations and glint-like features. They should not be interpreted as measured or validated parameters for THOR 6 or any other specific spacecraft.

Category	Value
Simulation configuration	
Observer location	ORM (28.76° N, 17.89° W, 2350 m)
Cadence	$\Delta t = 0.1$ s
Start time (UTC)	2026–02–03 22:00:00
End time (UTC)	2026–02–04 00:00:00
Orbital regime	GEO
Representative TLE	THOR 6
Spacecraft geometry	
Bus dimensions	$2.00 \times 2.00 \times 2.00$ m
Solar-panel wings	length = 12.00 m, width = 2.80 m
Panel segmentation	6 segments per wing
Cook-Torrance material parameters	
Bus	$\alpha = 0.250$, $F_0 = 0.600$
Panels	$\alpha = 0.080$, $F_0 = 0.250$
Uniform spin case	
Spin period	$P = 200.0$ s
Spin axis, $\hat{\mathbf{s}}$	$[+0.142,\ +0.906,\ -0.400]$
Angular velocity, $\vec{\omega}$	$[+0.0045,\ +0.0284,\ -0.0126]$ rad s^{-1}
Torque-free tumbling case	
Principal moments of inertia	$I_x = 3330.5$, $I_y = 1717.7$, $I_z = 3922.8$ kg m^2
Initial $\vec{\omega}_0$	$[+0.0051,\ -0.0205,\ +0.0232]$ rad s^{-1}
Mid-interval $\vec{\omega}_{\text{mid}}$	$[-0.0285,\ +0.0036,\ +0.0073]$ rad s^{-1}

Table 2.4: Model geometry, material properties, observer configuration, and attitude parameters used in the illustrative Cook-Torrance BRDF light curve simulation. These values define the simulation setup only and are not intended as a validated physical model of THOR 6 or any other specific spacecraft.

2.5.3 SURFACE-SCATTERING MODEL

In the simulation, the bus and solar-panel facets were assigned different Cook-Torrance parameters. The bus was assigned a larger roughness value to represent a broader, less sharply peaked scattering response, while the solar panels were assigned a smaller roughness value to allow stronger glint-like features. These choices are intended to produce representative qualitative behaviour rather than to reproduce the measured surface properties of a specific spacecraft.

At each time step, the BRDF was evaluated using the incident unit vector from the satellite to the Sun, $\hat{\omega}_i$, the outgoing unit vector from the satellite to the observer, $\hat{\omega}_o$, and the inertial-frame surface normal of each facet, $\hat{n}_s(t)$. Facets for which $\hat{n}_s(t) \cdot \hat{\omega}_i \leqslant 0$ or $\hat{n}_s(t) \cdot \hat{\omega}_o \leqslant 0$ were excluded from the single-sided bus calculation, as they were either not illuminated by the Sun or not visible to the observer. The solar-panel facets were treated as two-sided surfaces, so the normal direction was reversed where necessary to represent reflection from the rear side of the panel.

2.5.4 FACET FLUX SUMMATION

At each time step, the body-frame normal of each facet was rotated into the inertial frame using the current attitude quaternion. The reflected contribution from each visible and illuminated facet was then calculated from its BRDF, projected area, and observer range. The irradiance received by the observer was approximated as

$$E_{obs}(t) = \sum_{s=1}^{N_s} E_{\odot} f_{r,s} \frac{A_s \left[\hat{n}_s(t) \cdot \hat{\omega}_i(t)\right] \left[\hat{n}_s(t) \cdot \hat{\omega}_o(t)\right]}{d(t)^2}, \tag{2.17}$$

where $E_{\odot}$ is the incident solar irradiance at the satellite, A_s is the area of facet s, $f_{r,s}$ is the Cook-Torrance BRDF evaluated for that facet, and $d(t)$ is the satellite-observer range.

This expression retains the key dependencies required for the illustrative simulation: illumination angle, viewing angle, inverse-square range dependence, surface area, and material-dependent scattering response. Mutual shadowing between bus and panel components was not included, so the simulation should be interpreted as a simplified forward model rather than a full ray-tracing calculation.

The resulting irradiance was converted into an illustrative AB-like magnitude by assuming an effective monochromatic wavelength of $\lambda = 550$ nm. Since this conversion does not integrate over a measured instrumental bandpass, the resulting values should be interpreted as relative AB-like magnitudes rather than calibrated synthetic photometry. The main quantity of interest in this simulation is the morphology and relative variability of the light curve, not its absolute photometric calibration.

2.5.5 ATTITUDE MODELS

Two idealised attitude states were compared: uniform spin and torque-free tumbling. These cases were chosen because they represent two qualitatively different types of rotational behaviour that may be encountered when interpreting unresolved RSO light curves.

The attitude of the simulated spacecraft was represented using unit quaternions. A quaternion is a four-component quantity,

$$\bar{q} = [q_0,\ q_1,\ q_2,\ q_3], \tag{2.18}$$

where q_0 is the scalar component and (q_1, q_2, q_3) form the vector component. For attitude propagation, only unit quaternions are used, such that

$$|\bar{q}| = 1. \tag{2.19}$$

A unit quaternion can represent a rotation by an angle θ about a unit axis $\hat{u}$ as

$$\bar{q} = \left[\cos\left(\frac{\theta}{2}\right), \hat{u}\sin\left(\frac{\theta}{2}\right)\right]. \tag{2.20}$$

Quaternions were used instead of Euler angles because they provide a compact way to compose rotations and avoid singularities such as gimbal lock (Kuipers, 1999). To rotate a body-frame vector $\vec{v}_{\text{body}}$ into the inertial frame, the vector is first written as a pure quaternion,

$$\vec{v} = \left[0,\ \vec{v}_{\text{body}}\right], \tag{2.21}$$

and transformed as

$$\vec{v}_{\text{inertial}} = \bar{q} \otimes \vec{v} \otimes \bar{q}^{-1}, \tag{2.22}$$

where $\otimes$ denotes quaternion multiplication and $\bar{q}^{-1}$ is the inverse quaternion. For a unit quaternion, the inverse is simply the conjugate. This operation was used at each time step to rotate every facet normal from the spacecraft body frame into the inertial frame before evaluating the illumination and viewing geometry.

2.5.5.1 Uniform spin

In the uniform-spin case, the spacecraft rotates at a constant angular speed about a fixed axis in the body frame. The angular velocity vector is

$$\vec{\omega} = \omega\hat{s}, \tag{2.23}$$

where $\hat{s}$ is the unit spin axis and

$$\omega = \frac{2\pi}{P} \tag{2.24}$$

for spin period P. In this simulation, $P = 200$ s.

For a constant spin axis and angular speed, the quaternion attitude update can be written analytically as

$$\bar{q}(t) = \bar{q}_\Delta(t) \otimes \bar{q}_0, \tag{2.25}$$

where $\bar{q}_0$ is the initial attitude, and $\bar{q}_\Delta(t)$ is the axis-angle quaternion corresponding to a rotation of angle ωt about $\hat{s}$.

Physically, this case represents a simple, dynamically regular rotation state. The same facets repeatedly return to similar Sun-satellite-observer geometries, so the light curve is expected to show more regular and repeatable brightness features. This is not intended to imply that all defunct satellites spin in this way; rather, it provides a controlled reference case against which more complex attitude motion can be compared.

2.5.5.2 TORQUE-FREE TUMBLING

The second case considers torque-free tumbling. This describes the free rotation of a rigid body when no external torque is applied. If the body has unequal principal moments of inertia and the angular velocity vector is not aligned with one of the principal axes, the spacecraft does not simply rotate about a single fixed body axis. Instead, the angular velocity vector evolves in the body frame, producing precession and nutation of the body orientation.

The principal moments of inertia were denoted by (I_x, I_y, I_z), and the angular velocity in the principal-axis body frame was written as

$$\vec{\omega} = \left(\omega_x, \omega_y, \omega_z\right). \tag{2.26}$$

In the absence of external torques, the rotational motion is governed by Euler's equations,

$$I_x\dot{\omega}_x = (I_y - I_z)\omega_y\omega_z, \tag{2.27}$$

$$I_y\dot{\omega}_y = (I_z - I_x)\omega_z\omega_x, \tag{2.28}$$

$$I_z\dot{\omega}_z = (I_x - I_y)\omega_x\omega_y. \tag{2.29}$$

These equations show that, when I_x, I_y, and I_z are unequal, the angular-velocity components can vary even though the total angular momentum and rotational kinetic energy are conserved. This produces a changing instantaneous rotation axis in the body frame.

The quaternion kinematic equation was integrated simultaneously with Equation (2.29). For a body-frame angular velocity ω, the quaternion derivative is

$$\dot{\bar{q}} = \frac{1}{2}\bar{q} \otimes \begin{bmatrix} 0 \\ \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}. \tag{2.30}$$

The coupled system was integrated using a fourth-order Runge-Kutta scheme. This is an explicit numerical integration method in which the derivative is evaluated four times per time step: once at the start of the interval, twice at intermediate half-steps, and once at the end of the interval. These four estimates are then combined in a weighted average to advance the state. For a generic state vector y satisfying $\dot{y} = f(t, y)$, the update from t_n to $t_{n+1} = t_n + \Delta t$ is

$$k_1 = f(t_n, \mathbf{y_n}), \tag{2.31}$$

$$k_2 = f\left(t_n + \frac{\Delta t}{2}, \mathbf{y_n} + \frac{\Delta t}{2}k_1\right), \tag{2.32}$$

$$k_3 = f\left(t_n + \frac{\Delta t}{2}, \mathbf{y_n} + \frac{\Delta t}{2}k_2\right), \tag{2.33}$$

$$k_4 = f\left(t_n + \Delta t, \mathbf{y_n} + \Delta t\, k_3\right), \tag{2.34}$$

followed by

$$\mathbf{y}_{n+1} = \mathbf{y}_n + \frac{\Delta t}{6}\left(k_1 + 2k_2 + 2k_3 + k_4\right). \tag{2.35}$$

In this case, the state vector contains both the attitude quaternion and the body-frame angular velocity,

$$\mathbf{y} = \left[\bar{q},\ \vec{\omega}\right]. \tag{2.36}$$

The time step used for the integration was the same as the photometric sampling cadence, $\Delta t = 0.1$ s, which is small compared to the imposed 200 s characteristic rotation period. The quaternion was renormalised after each integration step to avoid numerical drift away from unit length.

Physically, the torque-free tumbling case represents a non-principal-axis rotation state. Such behaviour can occur for defunct spacecraft, rocket bodies, or debris objects whose attitude is no longer actively controlled. Because the body orientation evolves about multiple axes, the same facets do not return to the same illumination and viewing geometry in a strictly periodic way. The resulting light curve can therefore contain more complex, less repeatable brightness features than the uniform-spin case.

2.5.6 Illustrative light curve comparison

Figure 2.17 shows the simplified wireframe geometry described in §2.5.2 and the resulting simulated light curves for the two attitude cases. In the uniform-spin case, the brightness variations are comparatively regular because the same reflecting facets repeatedly pass through similar geometries with respect to the Sun and observer. In the torque-free tumbling case, the instantaneous body-frame rotation vector evolves according to Euler's equations, causing the facet orientations to change in a less repetitive manner. This produces a more complex light curve morphology.

The comparison demonstrates why attitude state is a central ambiguity in the interpretation of unresolved RSO photometry. A change in rotational state alone, with the same simplified geometry, material parameters, orbit, and observing site, can substantially alter the observed brightness pattern. The inverse problem is correspondingly degenerate: it is well established that different combinations of shape, material, and attitude can produce similar light curve features (see e.g., Linares et al., 2014; Bradley and Axelrad, 2014). This motivates the use of time-resolved, multi-band observations and physically informed forward models when attempting to interpret RSO light curves.

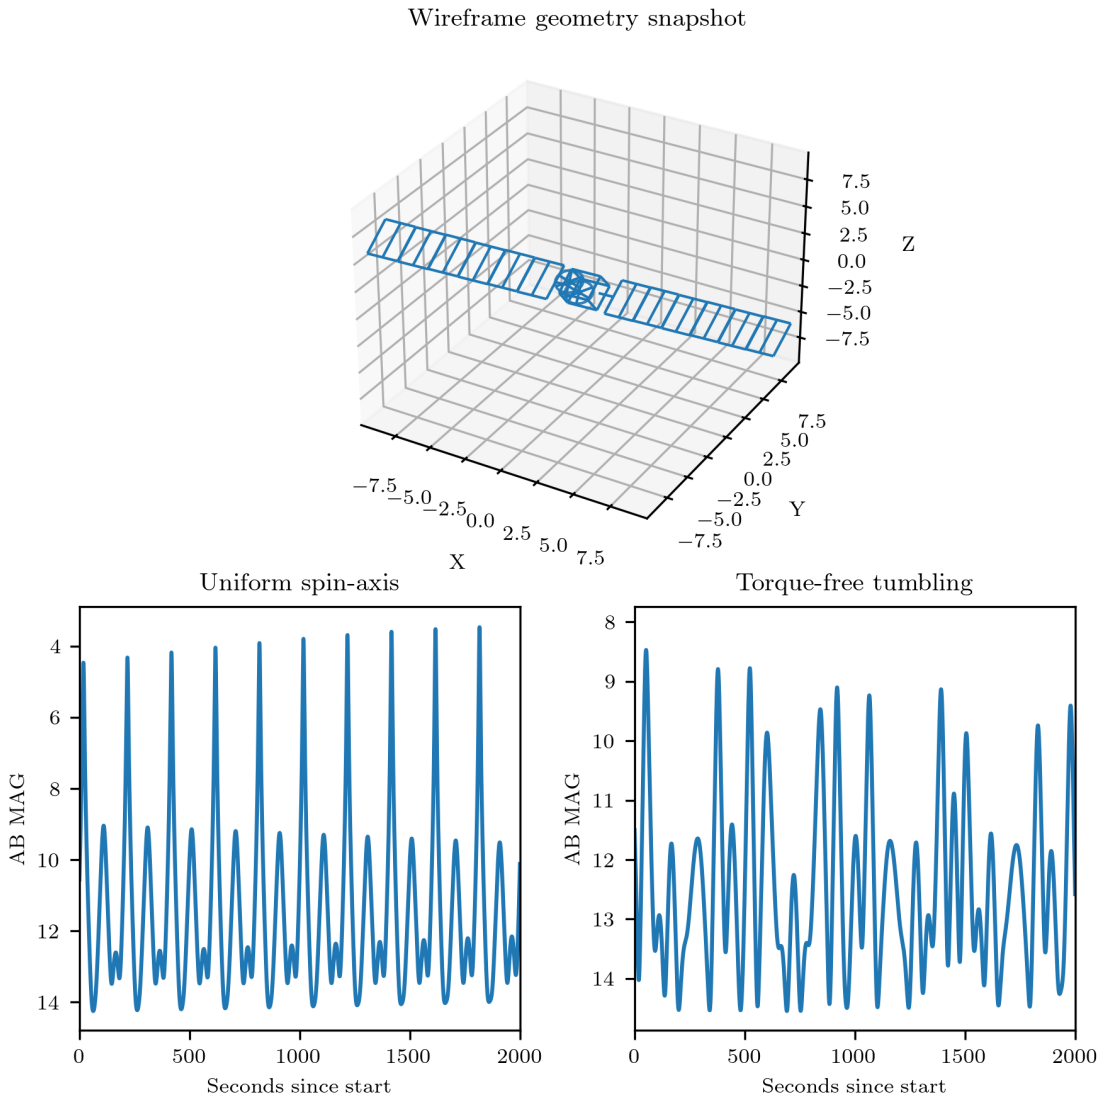

FIGURE 2.17: Top: 3D wireframe representation of the simplified box-wing satellite geometry used in the illustrative BRDF simulation. **Bottom left**: Simulated relative AB-like magnitude light curve for uniform spin about a fixed body-frame axis. **Bottom right**: Simulated relative AB-like magnitude light curve for torque-free tumbling of a rigid body with unequal principal moments of inertia. The two simulations use the same orbital geometry, observer location, spacecraft geometry, and material parameters; only the attitude evolution differs. The comparison illustrates how the rotational state of an unresolved RSO can strongly affect the morphology of its observed light curve.

2.6 Light Curve Analysis Techniques

2.6.1 Period Determination Methods

As demonstrated previously, optical light curves of defunct artificial satellites often exhibit periodic or quasi-periodic brightness variations caused by rotational motion, tumbling, and changing illumination and viewing geometries. Determining the characteristic period of these variations is fundamental for estimating satellite spin rates and diagnosing attitude stability. This task is complicated by irregular temporal sampling, limited observation durations, and non-sinusoidal light curve morphology.

To address these challenges, three period determination techniques are commonly employed: Phase Dispersion Minimisation (PDM), the Lomb-Scargle periodogram, and the Fast Fourier Transform (FFT). Each method relies on different assumptions and is sensitive to different features of the data.

Phase Dispersion Minimisation is a non-parametric technique designed to identify periodicity without assuming a specific signal shape (Stellingwerf, 1978). For a trial period P, the light curve is folded in phase according to

$$\phi_i = \frac{t_i \bmod P}{P}, \tag{2.37}$$

where t_i are the observation times. The phased data are divided into bins, and the variance within each bin is compared to the overall variance of the data. This comparison yields the PDM θ statistic,

$$\theta = \frac{\sum_{j=1}^{M}(n_j - 1)\sigma_j^2}{(N-1)\sigma^2}, \tag{2.38}$$

where M is the number of phase bins, n_j and σ_j^2 are the number of points and variance in the j-th bin, N is the total number of data points, and σ^2 is the variance of the full dataset. Periodic signals correspond to minima in θ.

Because PDM does not rely on sinusoidal fitting, it is well suited to signals with sharp features, asymmetries, or multiple maxima within a single rotation cycle. It is also robust to uneven sampling, making it particularly useful when the underlying variability departs from simple harmonic behaviour. However, resolving such structure requires a sufficiently large number of datapoints, so that the phase-folded signal is adequately sampled across multiple bins.

The Lomb-Scargle periodogram is a spectral analysis method developed to detect periodic signals in unevenly sampled time series (Lomb, 1976; Scargle, 1982). It estimates the power at angular frequency ω by fitting sinusoidal models of the form

$$y(t) = A\cos(\omega t) + B\sin(\omega t), \tag{2.39}$$

and computing the normalised power

$$P(\omega) = \frac{1}{2\sigma^2}\left[\frac{\left(\sum_i y_i \cos\omega(t_i-\tau)\right)^2}{\sum_i \cos^2\omega(t_i-\tau)} + \frac{\left(\sum_i y_i \sin\omega(t_i-\tau)\right)^2}{\sum_i \sin^2\omega(t_i-\tau)}\right], \tag{2.40}$$

where τ is a phase offset chosen to ensure invariance under time translation.

A key advantage of the Lomb-Scargle periodogram is the ability to assess the statistical significance of detected periodicities through the false-alarm probability (FAP). The FAP represents the probability that a peak of a given height could arise purely from random noise. For a maximum observed power z, the FAP may be approximated by

$$\mathrm{FAP} \approx 1 - (1 - e^{-z})^{M}, \tag{2.41}$$

where M is the number of independent frequencies sampled (Scargle, 1982). Low FAP values indicate a statistically significant detection, while higher values suggest that the peak may be noise-induced.

While Lomb-Scargle is highly effective for approximately sinusoidal signals, its reliance on harmonic fitting can lead to ambiguity when the true signal is strongly non-sinusoidal or dominated by higher-order harmonics.

The discrete Fourier transform decomposes a time series into sinusoidal frequency components. For uniformly sampled data with samples x_n taken at regular intervals, the DFT is commonly written as

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i k n/N}. \tag{2.42}$$

The Fast Fourier Transform (FFT) is an efficient algorithm for evaluating this transform, reducing the computational cost relative to direct summation (Cooley and Tukey, 1965).

The main limitation of the FFT is that it assumes samples lie on a regular time grid. For unevenly sampled observations, one may instead evaluate the Fourier transform directly using the actual sampling times,

$$X(f) = \sum_{n=0}^{N-1} x_n e^{-2\pi i f t_n}, \tag{2.43}$$

although this is computationally more expensive. Applying an FFT to unevenly sampled data therefore usually requires interpolation or resampling onto a uniform grid, which can introduce spectral artefacts, reduce sensitivity to short-period signals, and complicate interpretation of the resulting power spectrum (Press et al., 1992). For regularly sampled light curves such as those obtained with STING, where gaps are short and infrequent, this is a minor concern, and the FFT provides an efficient and direct estimate of the dominant frequency.

Each period determination method offers complementary strengths and characteristic failure modes. PDM is the most robust for non-sinusoidal and irregular signals, but produces minima at every integer multiple of the true period and at aliases of the sampling, so the deepest minimum need not be the fundamental; Lomb-Scargle provides efficient frequency detection with statistical significance estimation via the false-alarm probability, but concentrates power in the strongest harmonic of a non-sinusoidal signal; and FFT offers a computationally efficient frequency-domain representation when uniform sampling is available.

Method	Estimated Period (s)
Lomb-Scargle	99.9
PDM	199.8
FFT (Interpolated)	100.0

Table 2.5: Recovered period estimates following the application of the Lomb-Scargle, PDM and FFT methods on simulated BRDF light curve data.

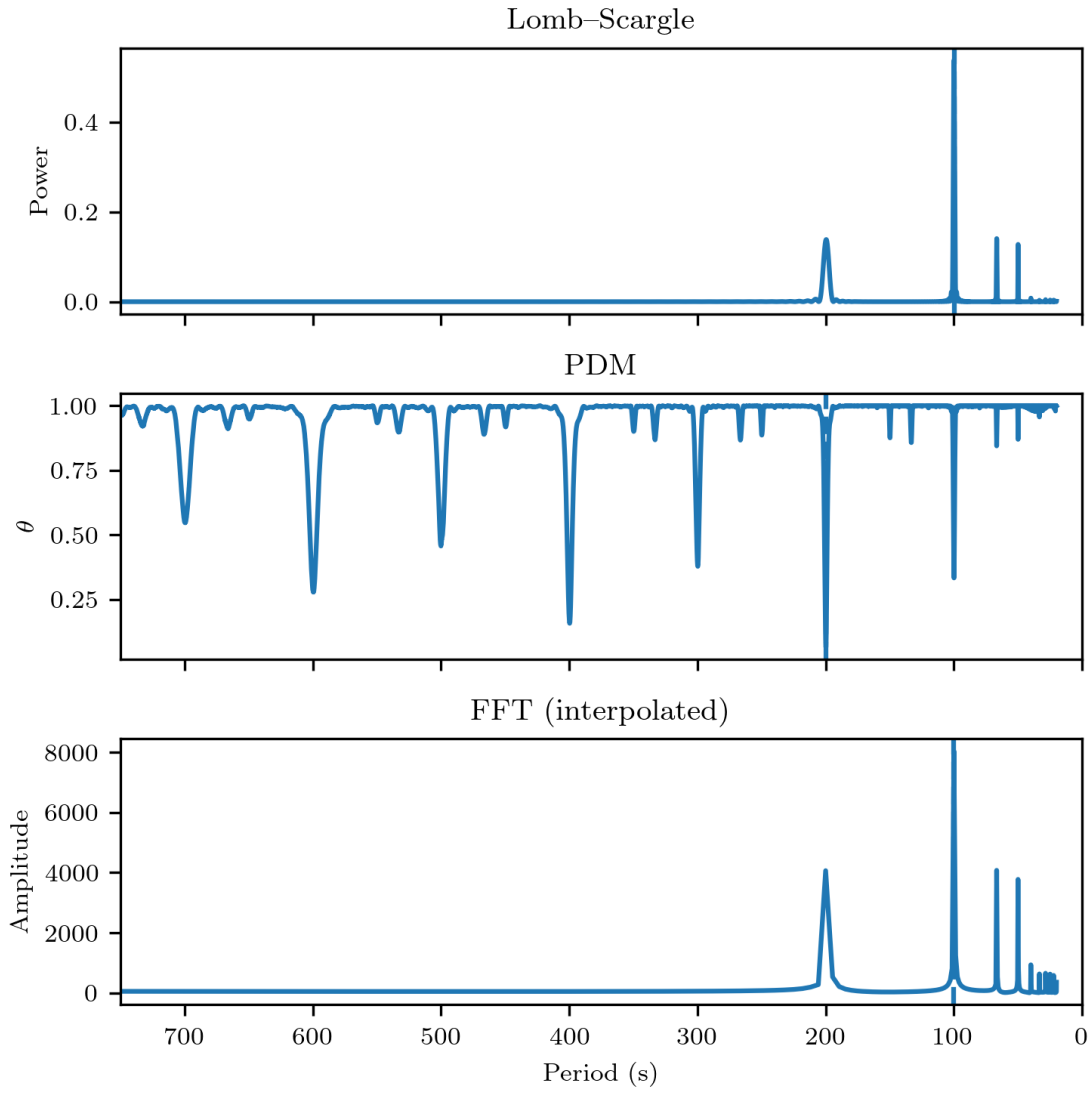

Figure 2.18: Periodograms following the application of the Lomb-Scargle, PDM and FFT methods on simulated BRDF light curve data.

To provide an example of the three period determination methods in action, here taking the simulated BRDF model light curve data from the uniform spin-axis mode (see Figure 2.17) and passing it to the different methods results in the periodograms displayed in Figure 2.18 and consequently the values given in Table 2.5. The light curve is simulated with a spin-period of 200.0 seconds and therefore there is an expectation of what should be derived. All three methods detect the same

signal but report it differently. The Lomb-Scargle and FFT spectra place their strongest peak at 100 s with a secondary peak at 200 s: the light curve is strongly non-sinusoidal, so its Fourier power is concentrated in the first harmonic of the spin period rather than the fundamental. PDM places its deepest minimum at the fundamental, but also produces minima at integer multiples of the period (400 s, 600 s and 700 s in Figure 2.18), since folding at nP also aligns the signal. Neither behaviour is specific to this simulation: a non-sinusoidal periodic signal will always concentrate Fourier power in its strongest harmonic, and PDM will always show sub-harmonic minima. The fundamental is therefore identified by combining the methods: the Lomb-Scargle peaks at 100 and 200 s form a harmonic series whose longest member coincides with the deepest PDM minimum. For real observations, gaps and finite duration additionally introduce sampling aliases into all three periodograms. It is important to note, that the artificial signal from the BRDF model don't come with any associated noise and hence application of these methods on real satellite light curves has an added layer of complexity.

In Chapter 6, the focus is on how de-trending whilst minimising θ and how additional colour and polarisation data can possibly help break the degeneracy in recovering the true rotation period using the PDM method.

2.6.2 RELATING LIGHT CURVE FEATURES TO PHYSICAL FEATURES

The ultimate goal of RSO characterisation is to take the observed light curves and link the features seen in them to the physical properties of the object.

Figure 2.19 illustrates the different interesting features after observation of various objects in different orbital regimes. Intelsat 10-02 is an actively 3-axis stabilised communications satellite in GEO. The inset axes for Intelsat 10-02 show a zoomed in regime between -15 to 5° of solar equatorial phase angle. The features displayed here are that of double peaked glint, due to the geometry between observer-Sun-satellite, a glint close to 0° of solar equatorial phase angle is expected since the solar panels are directed towards the Sun for power needs, and no other large surface on a stabilised GEO satellite is held near the Sun direction. The same panel glint appears as a single peak in the Thor-6 light curve of Figure 2.16, where the two arrays are offset together rather than relative to each other. Optical light curves can be used to identify and calculate how far a satellite's solar arrays are offset from the Sun (Payne et al., 2006) using this geometry. This is useful as these solar panel offsets have a direct implication on the status and health of a given GEO satellite and thus can be used to characterise them. The primary causes for this offset pointing can be:

- Ensuring the satellite has the required electrical generation from launch until its designated end of operation, the offsets are likely to be implemented intentionally at the beginning of operations as panels would generate too much power.
- Uneven degradation of the light sensors fitted on the solar panels can lead to unintended pointing when trying to balance the voltages across the panels (Payne et al., 2006).
- To balance the impact of the solar radiation pressure.

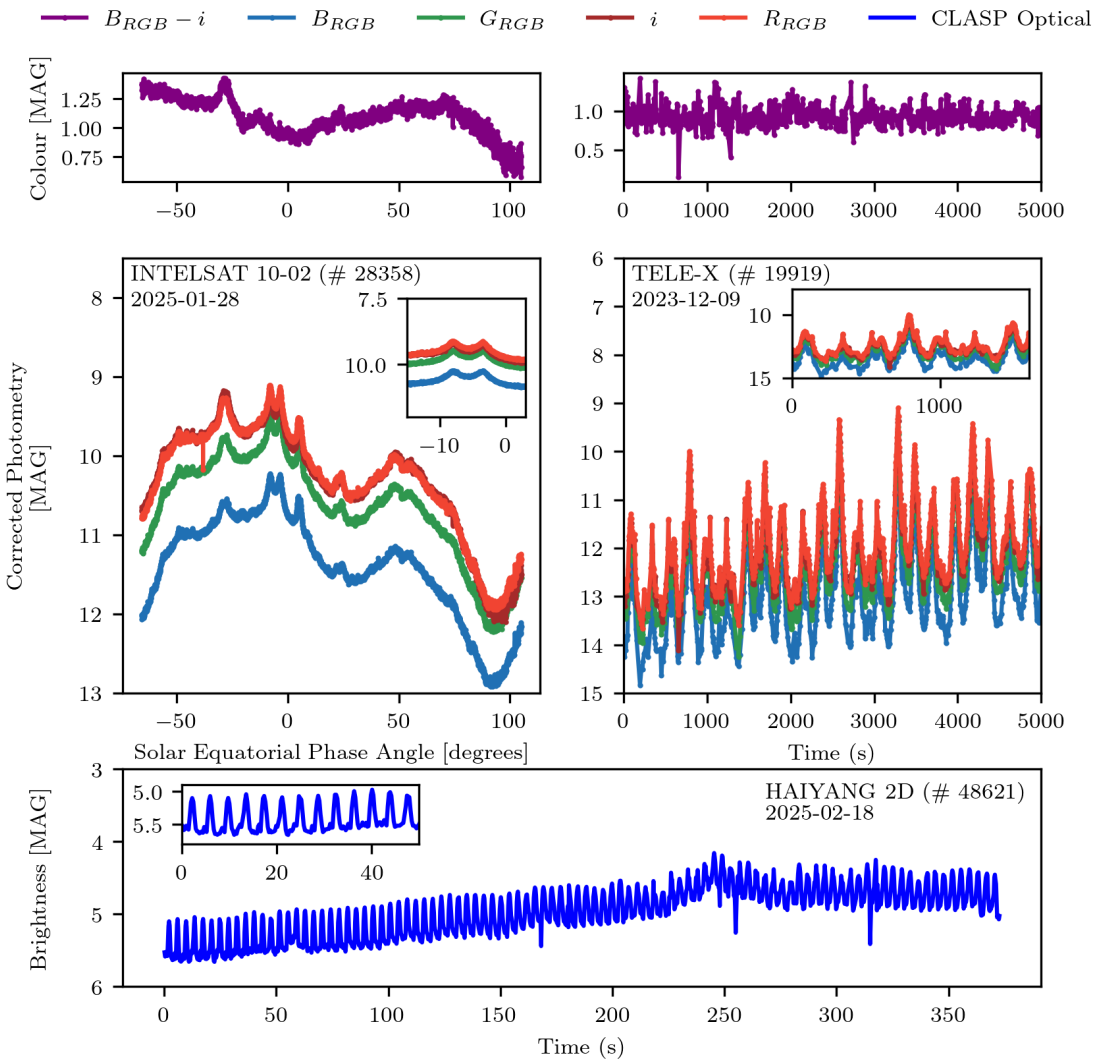

Figure 2.19: Examples of differing light curve features of satellites in different orbital regimes. The Intelsat 10-02 (GEO) and TELE-X (graveyard) multi-colour light curves were obtained through observation with the STING instrument and the HAIYANG 2D (LEO) light curve was obtained through observation with the CLASP instrument.

In this case, the double peak that lies away from 0° indicates that the solar panels are offset from the Sun slightly as well as offset from each other. Chapter 4 explores the measurement of solar panel offsets from a sample of observed active GEO satellites as well as the colour of solar panel glints.

TELE-X is a retired GEO communications satellite placed into graveyard orbit on the 16th of January 1998[1]. It is evident from the light curve that the object is under the influence of external torques causing it to tumble. The zoomed in regime of the light curve between 0 and ~ 1750 seconds illustrate at least 1 spin cycle. Chapter 6 explores how colour measurements can be utilised to break the degeneracy between spin period and harmonics derived from PDM measurements.

HAIYANG-2D is an active LEO satellite in the Chinese HY-2 series, designed for marine

[1] https://space.skyrocket.de/doc_sdat/tele-x.htm

dynamic environment monitoring[2]. Although the optical signature could initially be mistaken for the light curve of a uniformly rotating satellite, as illustrated in Figure 2.17, the measured periodicity is more likely associated with the satellite microwave radiometer antenna. The HY-2D scanning microwave radiometer uses a 1.2 m diameter offset parabolic reflector mounted at a fixed 40° half-cone angle, corresponding to an Earth incidence angle of 47.7° . The instrument rotates continuously about an axis parallel to the satellite body with a mechanical period of 3.79 s (Zhao et al., 2019).

This interpretation is strongly supported by a Lomb-Scargle periodogram, which gives a photometric period of 3.7998 s, in almost exact agreement with the radiometer scan period from the aforementioned documentation. The optical light curve is a direct detection of the spinning microwave radiometer antenna, rather than a signature of the bulk spacecraft attitude motion. The antenna is forward-viewing and scans $\pm 70°$ relative to the satellite nadir track, producing a 140° Earth-facing scan and a swath width of approximately 1600 km (Zhao et al., 2019). Importantly, this $\pm 70°$ range does not imply an oscillatory back-and-forth motion; the antenna rotates continuously, with only part of each rotation viewing the Earth. The observed 3.7998 s optical modulation should therefore be interpreted as the apparent synodic photometric signature of the continuously rotating antenna.

This demonstrates the importance of external documentation in linking the physical characteristics of a satellite to what is observed on the ground. It is also a concrete example of the degeneracy noted in the prior sections: the periodic signal of HAIYANG-2D superficially resembles that of the uniformly spinning body in Figure 2.17, yet has an entirely different physical origin.

2.7 Markov Chain Monte Carlo

Bayesian inference provides a rigorous statistical framework for estimating model parameters from observational data while explicitly accounting for measurement uncertainties and correlations between parameters. In this thesis, Bayesian methods are employed to infer physical parameters from photometric light curves of RSOs. Parameter estimation is performed using Markov Chain Monte Carlo (MCMC) techniques, which enable efficient sampling of posterior probability distributions in high-dimensional and potentially degenerate parameter spaces.

Bayesian inference is based on Bayes' theorem (Bayes and Price, 1763), which relates the posterior probability distribution of a set of model parameters θ, given observed data D, to the likelihood function and prior information:

$$p(\theta \mid D) = \frac{p(D \mid \theta)\, p(\theta)}{p(D)}. \tag{2.44}$$

Here, $p(\theta \mid D)$ is the posterior probability distribution, $p(D \mid \theta)$ is the likelihood of obtaining the data given a particular parameter set, $p(\theta)$ represents the prior distribution encoding existing knowledge or physical constraints, and $p(D)$ is a normalisation constant. For parameter estimation, the evidence term $p(D)$ does not need to be evaluated explicitly, as it is independent of θ (Lee, 2012).

[2] https://space.skyrocket.de/doc_sdat/hy-2c.htm

Assuming that individual measurements are statistically independent and that their uncertainties are well approximated by Gaussian errors, the likelihood function may be written as

$$p(D \mid \theta) \propto \exp\left(-\frac{1}{2}\chi^2\right), \tag{2.45}$$

where the chi-squared statistic is defined as

$$\chi^2 = \sum_i \frac{[m_i - M(t_i;\theta)]^2}{\sigma_i^2}. \tag{2.46}$$

Here, m_i is the observed measurement at time t_i, $M(t_i;\theta)$ is the model prediction evaluated at that time, and σ_i is the estimated photometric uncertainty. This likelihood formulation is commonly adopted in astronomical parameter estimation and provides a natural means of incorporating measurement uncertainties into the inference process (Lee, 2012; Sharma, 2017).

Prior distributions encode physically motivated constraints and existing knowledge of model parameters. When prior knowledge is limited, weakly informative priors are adopted to minimise bias while excluding unphysical solutions. All prior choices are explicitly defined to ensure transparency and reproducibility of the inference process (Lee, 2012).

Posterior distributions are sampled using an affine-invariant ensemble MCMC algorithm, as implemented in the `emcee` package (Foreman-Mackey et al., 2013). This class of sampler evolves an ensemble of walkers that explore the parameter space simultaneously, rather than relying on a single Markov chain. The affine-invariant property ensures efficient sampling of highly anisotropic or correlated posterior distributions.

At each iteration, the position of a given walker is updated based on the positions of other walkers in the ensemble, allowing collective exploration of the posterior distribution. This approach reduces sensitivity to initial conditions and improves robustness in complex parameter spaces compared to traditional single-chain MCMC methods.

The ensemble of walkers is initialised by drawing parameter values from the prior distributions or from small perturbations around an initial estimate. An initial burn-in phase is discarded to allow convergence toward the stationary posterior distribution.

For comparison, least-squares fitting methods aim to identify a single parameter set that minimises the squared residuals between model predictions and observations. While computationally efficient, such methods provide only point estimates and rely on local curvature approximations to estimate uncertainties. They do not naturally capture parameter degeneracies, non-Gaussian uncertainties, or multi-modal solutions.

Bayesian MCMC sampling provides a complete probabilistic description of the parameter space, enabling robust uncertainty quantification and explicit identification of degeneracies. The ensemble-based MCMC approach is also less sensitive to local minima and initial conditions, making it more suitable for complex models (Lee, 2012; Sharma, 2017).

2.7.1 Assumptions and limitations

The MCMC framework adopted in this thesis assumes that photometric errors are independent and approximately Gaussian. While generally valid for background-limited aperture photometry, this

assumption may not fully capture correlated noise arising from atmospheric variability or residual systematics. In addition, posterior inferences depend on the choice of prior distributions, which may influence results when the data have limited signal-to-noise ratio or temporal coverage.

MCMC methods are computationally intensive, requiring repeated evaluation of the likelihood function and limiting the complexity of models that can be explored. Furthermore, convergence diagnostics such as the integrated autocorrelation time test only whether the sampler has converged within the region it has explored; they cannot guarantee that all modes of a multimodal posterior have been found, particularly in the presence of strong degeneracies. Despite these limitations, Bayesian MCMC provides a robust and transparent framework for parameter estimation.

3 Characterisation of the STING Instrument

"Have you tried turning it off and back on again?"

Paul Sean Ernest Hunter, 2026

3.1 Motivations

STING (§2.2.3) was newly commissioned during my PhD project as a wide-field, simultaneous multi-colour instrument for observing resident space objects. Before its measurements can be used to interpret satellite colour, glint behaviour, and astrometric motion, the instrument must be characterised under real observing conditions. This chapter therefore establishes the performance envelope of STING (refer back to §2.2.1) in three areas: its photometric precision, its end-to-end timing and astrometric accuracy, and the practical visibility limits imposed by the mount, enclosure, and local horizon. These characterisations provide the basis for trusting the STING observations used in the later science chapters.

3.2 Photometric Performance

The photometric performance of STING was characterised by constructing a theoretical noise model for each camera following the standard formulation for detector noise in astronomical photometry (Howell, 1989). The various contributory sources of a constructed theoretical noise model were calculated using Equations 3.1 - 3.6 for each of the four filter bandpasses.

f_{target} refers to the synthetic target flux in analog-to-digital units (ADU), the raw detector counts, which are converted to detected electrons through the camera gain, $g = 0.43$ e^- ADU^{-1}, captured within the entire aperture and f_{sky} refers to the sky background within the aperture, $f_{sky} = n_{pix}\,\bar{s}$, where $\bar{s}$ is the mean background per pixel from the model fitted to each frame by the reduction pipeline (§2.4.3), averaged over all frames of a night for each camera. The sky background varies strongly with lunar phase, Moon-target separation and altitude, so two nights were chosen to bracket the range: 12th April 2023, with the full Moon crossing the meridian ("bright"), and 18th May 2023, near new Moon ("dark"). Real observations lie between these limits.

The noise model includes the following terms:

- N_{photon}: photon shot noise from the target flux, arising from the Poisson statistics of detected photons.

- N_{sky}: noise from the sky background within the photometric aperture.
- N_{DC}: dark-current noise from thermally generated charge in the detector.
- N_r: read noise introduced during detector readout.
- N_{scint}: atmospheric scintillation noise, estimated using the modified Young approximation.
- N_{total}: the quadrature sum of the relevant noise contributions.

The modified Young approximation used to calculate the scintillation noise, i.e., Equation 3.5 and consequently the empirical coefficient C_Y are taken from (Osborn et al., 2015). The scintillation term depends on the target's position in the sky through the airmass, with the noise amplitude scaling as $(AM)^{3/2}$; the model curves use the mean airmass of the survey sample, so the bright-end limit in Figure 3.1 is representative rather than exact for any individual target. The approximation is also an estimate of the *median* scintillation for a site: (Osborn et al., 2015) show from turbulence profiling on La Palma that the night-to-night distribution is log-normal, with first and third quartiles at roughly 0.8 and 1.3 times the median, and that Young's formulation captures neither this spread nor its dependence on wind direction. The scintillation floor in Figure 3.1 should therefore be read as typical rather than an exact representation for a given night.

Writing the detected target signal as $F = g\, f_{target}$ electrons, the fractional noise contributions in a single exposure are:

$$N_{Photon} = \frac{\sqrt{F}}{F} \tag{3.1}$$

$$N_{Sky} = \frac{\sqrt{g\, n_{pix}\, \bar{s}}}{F} \tag{3.2}$$

$$N_{DC} = \frac{\sqrt{n_{pix}\, I_{dc}\, t}}{F} \tag{3.3}$$

$$N_r = \frac{\sqrt{n_{pix}}\, \sigma_r}{F} \tag{3.4}$$

$$N_{Scint} = \sqrt{10^{-5}\, C_Y^2\, D_{tel}^{-4/3}\, t^{-1}\, (AM)^3\, e^{-2h/h_0}} \tag{3.5}$$

$$N_{Total} = \sqrt{N_{Photon}^2 + N_{Sky}^2 + N_{DC}^2 + N_r^2 + N_{Scint}^2} \tag{3.6}$$

where $\bar{s}$ is the mean sky background per pixel in ADU, I_{dc} and σ_r are the dark-current rate and read noise per pixel in electrons, D_{tel} is the aperture diameter, AM the airmass, h the observatory altitude and h_0 the atmospheric scale height. $n_{pix} = \pi r^2 = 78.5$ is the number of pixels in the $r = 5$ px ($20''$) aperture used typically throughout the thesis. Because the aperture is fixed rather than scaled to the point spread function, and is several times larger than the seeing disc, n_{pix} does not depend on seeing and the fraction of target flux lost outside the aperture is negligible; the sky and read-noise terms are correspondingly always integrated over the full aperture area.

The photometric noise model of STING is shown to be consistent with the RMS scatter measured in the satellite light curves (see Figure 3.1). Satellite light curves which were measured during 'bright' and 'dark' nights (with respect to the level of illumination of the moon disk) were binned

with a sufficiently small bin width (10 minutes) such that a linear fit could be made, for which the RMS is calculated for each bin.

The typical photometric precision in a single 10 s exposure varies between roughly 1% for targets brighter than 9th magnitude (scintillation limited), and increases to 10% at 14th magnitude (limited by sky background and readout noise). Dark current is negligible when operating the cameras at 0 °C. Multiple exposures can be binned to improve the precision of faint targets.

These results show that STING achieves the percent-level per-exposure precision, with simultaneous colour bands, set as the photometric requirement in §2.2.1. For bright RSOs, the single-exposure precision is sufficient to resolve short-timescale glints (typically of 0.1 mag or greater) and colour changes substantially larger than the percent-level noise floor. For fainter targets, temporal binning can improve the precision where the intrinsic variability timescale permits. The instrument is therefore well suited to the broad colour and light curve morphology studies in later chapters, although subtle colour changes in individual faint exposures should be interpreted with the corresponding noise limits in mind. The expectation of §2.2.1, that coarse sampling penalises only the sky-limited faint end, is also borne out: on dark nights the n_{pix}-dependent terms in Equation (3.6) become significant only fainter than ~12th magnitude.

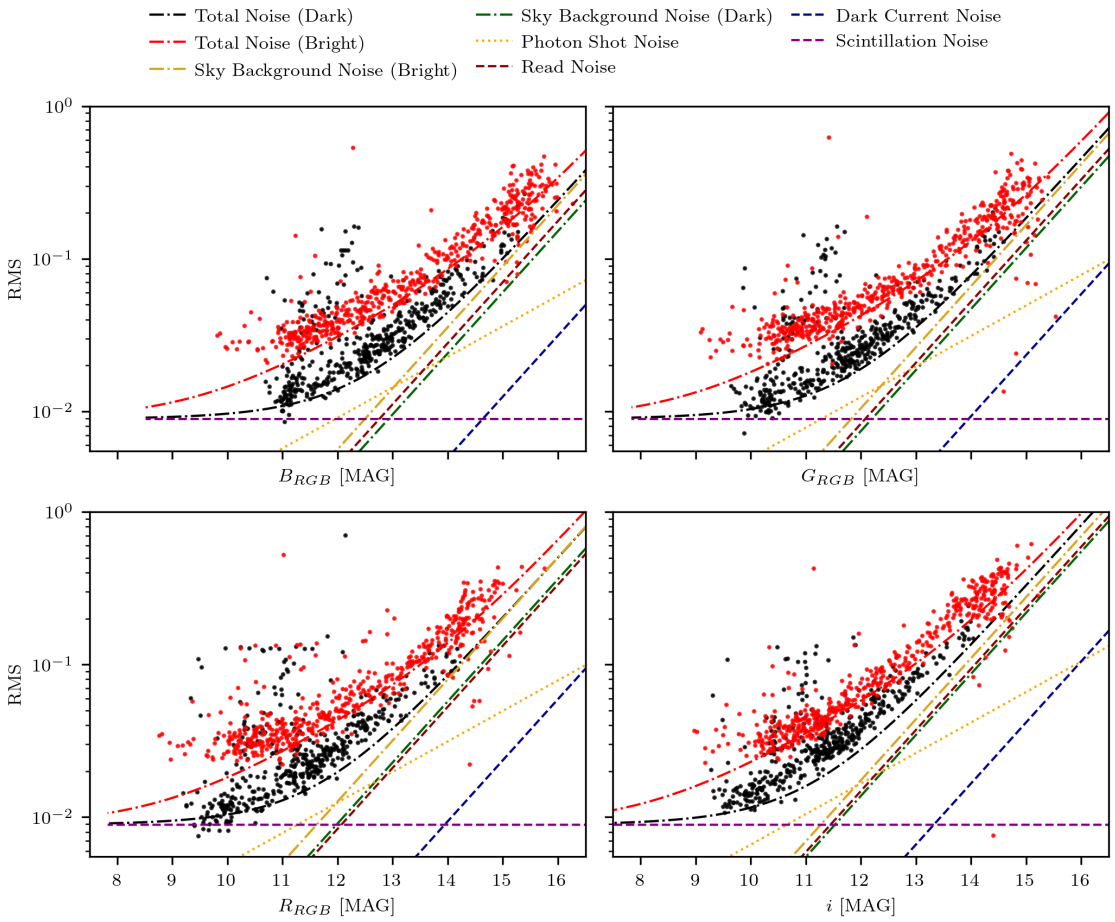

FIGURE 3.1: Fractional RMS of STING satellite light curves in single 10 s exposures as a function of calibrated magnitude, for the four cameras: CAM1 (B_{RGB}, top left), CAM2 (G_{RGB}, top right), CAM3 (R_{RGB}, bottom left) and CAM4 (i, bottom right). Magnitude increases (brightness decreases) to the right. The lower envelope of the measured scatter agrees well with the theoretically calculated noise model under both bright- and dark-sky conditions. Bright-time measurements (12th April 2023) are shown by the red markers and dark-time measurements (18th May 2023) by the black markers, with the corresponding total noise models plotted for each condition together with the individual photon shot noise, sky background noise, read noise, dark current noise, and scintillation noise components. The model is a lower envelope: points lying well above it, such as the group at $B_{RGB} \approx 10$–11 on the dark night, are targets exhibiting short-timescale glints within a 10-minute bin, which are not removed by the linear de-trending and measure real brightness changes rather than noise. Read and dark current noise are identical as functions of detected flux for the four cameras, but appear offset between panels because the curves are plotted against band-dependent calibrated magnitude, where each camera has a different photometric zero-point.

3.3 Horizon Limits

On 9 February 2026 and 12 March 2026, short automated survey sequences were executed to empirically constrain the telescope's practical visibility limits near the horizon. Each sequence consisted of time-ordered 10 s exposures while the telescope stepped through a grid in altitude and azimuth, sampling altitude bands at approximately 35°, 30°, 25°, 20°, 15°, 10°, and 5°. Three cameras were available on each night: CAM1–CAM3 on 9 February, and CAM1, CAM2, and CAM4 on 12 March. The missing camera on each night was due to a fault in the power distribution unit.

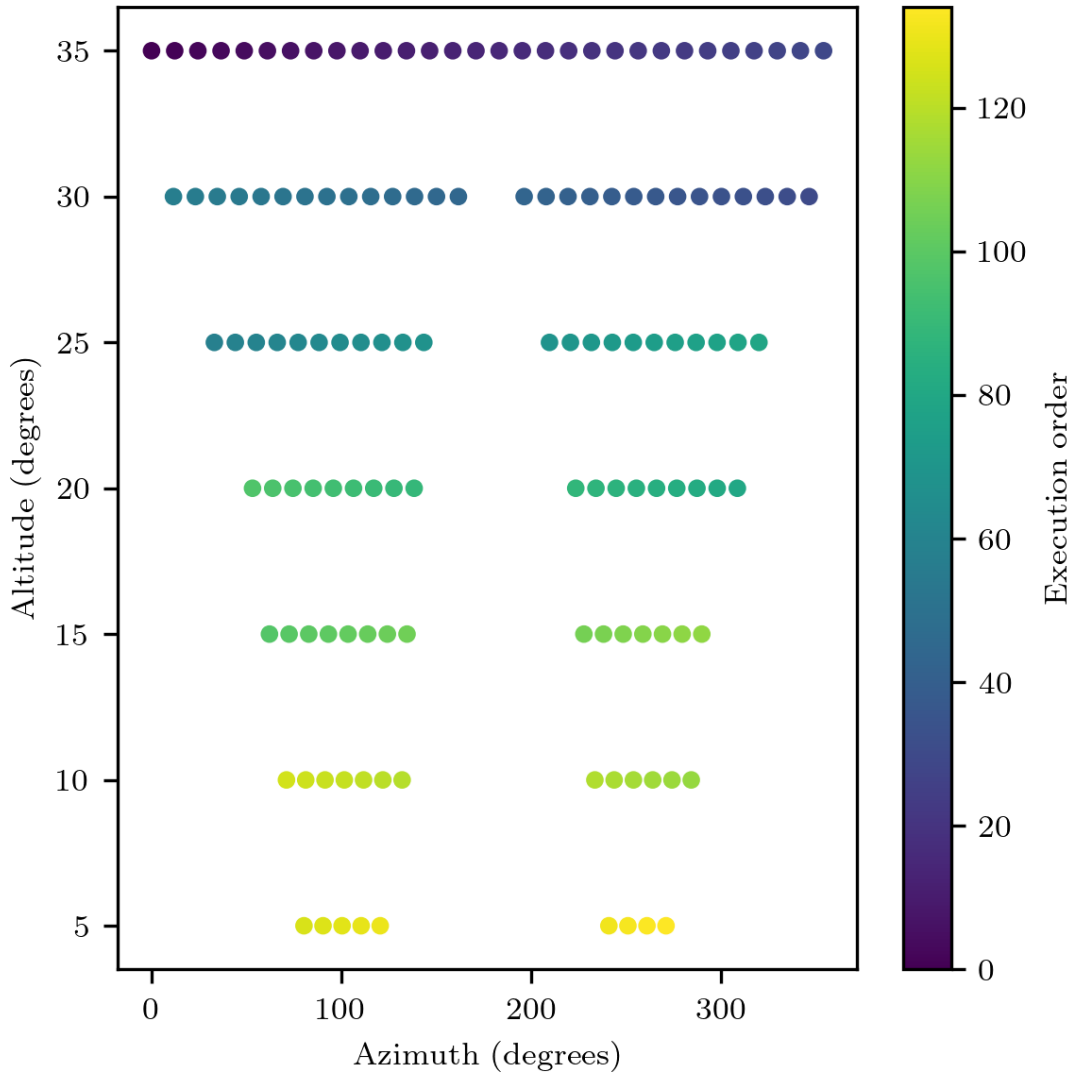

Figure 3.2: Example Alt–Az observing sequence for the 9 February 2026 horizon survey. Each marker represents a telescope pointing, with colour indicating execution order from early observations in dark/purple to later observations in yellow across successive altitude bands.

The survey was designed to probe progressively lower altitudes, where vignetting and dome or structural obscuration are expected to become increasingly significant. These observations therefore provide a data-driven basis for defining operational visibility constraints in the control system, ensuring that future science observations are scheduled only in regions of the sky that are both accessible and useful for STING. Figure 3.2 shows the executed pointings in Alt–Az space. Gaps due north and south correspond to requested pointings that fell outside the TCS limits of $-90° < \mathrm{HA} < 90°$ and $-30° \leqslant \delta \leqslant 85°$.

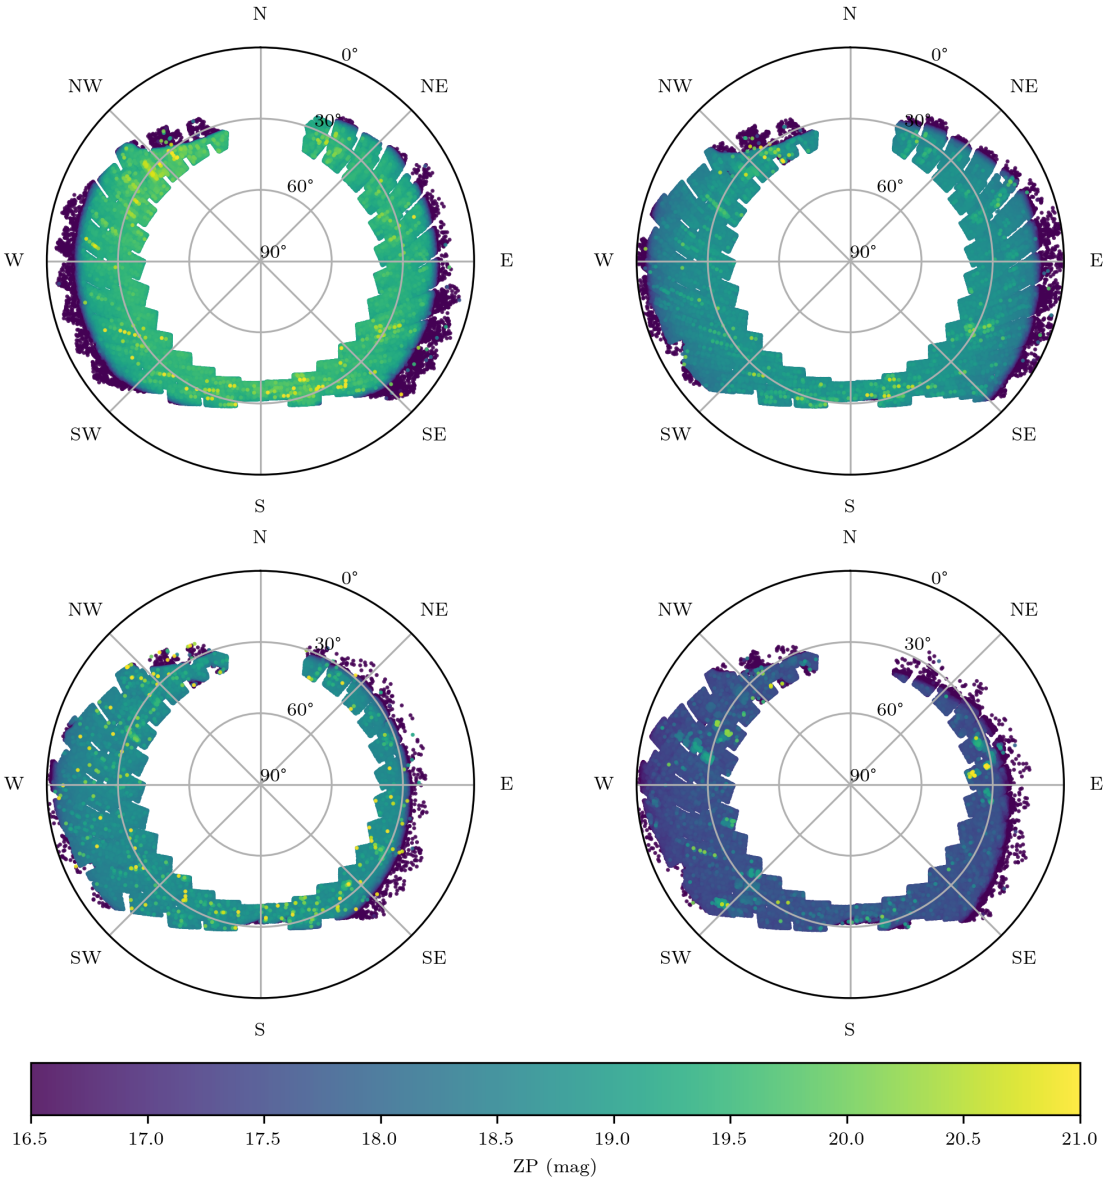

FIGURE 3.3: Alt–Az zero-point maps for the four cameras, used to identify horizon limits. Each panel shows catalogue-star zero-points measured from calibrated images and plotted in local Alt–Az coordinates, using a polar projection with radius = 90° − alt and angle equal to azimuth. Points are coloured by zero-point magnitude using a common colour scale. Panels show **top left**: Camera 1 (B_{RGB}), **top right**: Camera 2 (G_{RGB}), **bottom left**: Camera 3 (R_{RGB}), and **bottom right**: Camera 4 (i). The gaps due north and south are due to: $\delta \gtrsim -30°$ (transiting at altitude $\approx 31°$) to the south and $|\mathrm{HA}| < 90°$ to the north.

Photometric quality was assessed by measuring per-star zero-points in each calibrated exposure and examining their spatial dependence in local Alt–Az coordinates. Each raw image was dark- and flat-corrected before aperture photometry was performed on catalogue stars. For each accepted star, the instrumental zero-point was computed as

$$\mathrm{ZP} = m_{\mathrm{cat}} + 2.5 \log_{10}(F/t_{\mathrm{exp}}).$$

The resulting zero-points are shown in Figure 3.3, one panel per camera, using a common colour scale to allow direct comparison. The maps show a clear decline in zero-point as the instrumental horizon limits are approached. The zero-point is used in preference to raw flux because the aim is to locate the gradient, not the cutoff: complete blocking appears simply as the absence of stars at the edge of each map, whereas the decline in zero-point towards that edge measures the region of progressive vignetting that precedes it. Being independent of stellar brightness, the zero-point also allows stars of all magnitudes to be combined into a single map.

To summarise the derived obstruction limits, three complementary views are combined in Figure 3.4. First, the observed Alt–Az measurements are reduced to an azimuth-dependent minimum-altitude boundary, $\mathrm{alt}_{\mathrm{min}}(\mathrm{az})$. This is obtained by binning in azimuth, taking a low-percentile altitude in each bin, filling only short gaps, and applying circular smoothing to suppress outliers while preserving real structure. The same boundary is then mapped into the equatorial frame by evaluating the margin

$$m(\mathrm{HA}, \delta) = \mathrm{alt}(\mathrm{HA}, \delta) - \mathrm{alt}_{\mathrm{min}}\big(\mathrm{az}(\mathrm{HA}, \delta)\big)$$

on a dense (HA, δ) grid. The visible region is defined by $m \geqslant 0$, and the horizon limit is extracted as the $m = 0$ contour. Finally, $\mathrm{alt}_{\mathrm{min}}(\mathrm{az})$ is also shown in polar Alt–Az coordinates to provide an intuitive sky-footprint view. The GEO belt as seen from La Palma is overlaid in both coordinate systems for reference.

These horizon limits translate directly into operational scheduling constraints for STING. Their impact is strongly direction-dependent because the visible GEO belt from La Palma spans only a limited arc of the local sky. To the south, the derived boundaries lie well below the GEO belt for all four cameras. The belt reaches alt $\approx 56^\circ$ near due south, while $\mathrm{alt}_{\mathrm{min}}(\mathrm{az})$ remains $\lesssim 20^\circ$ over azimuths of approximately 100°–260°. Routine GEO survey observations are therefore unaffected by structural obscuration in this region.

Towards the north, by contrast, the boundaries rise sharply. At azimuths $\lesssim 30^\circ$ and $\gtrsim 330^\circ$, the dome and pier impose a hard cutoff at alt $\gtrsim 30^\circ$. This has little practical effect on GEO operations, because from La Palma the GEO belt is not visible above the horizon at azimuths north of approximately 95° or 265°. The loss of northern sky access therefore removes no usable GEO targets, although it does restrict access to high-declination sidereal calibration fields. Such fields can instead be scheduled close to transit.

The remaining operational concern is the east-west asymmetry in $\mathrm{alt}_{\mathrm{min}}$. To quantify this, the camera-by-camera horizon limits were compared with the longitude distribution of the catalogued GEO population. The master TLE list was propagated to the observing epochs and filtered for GEO-like objects using $0.9 \leqslant n \leqslant 1.1$ rev sidereal-day^{-1}, eccentricity < 0.2, and inclination $< 30^\circ$. This yields $N = 1036$ objects. As expected, the distribution shows pile-ups near the two

gravitationally stable longitudes: a dominant eastern peak at $+70^\circ \lesssim \lambda \lesssim +100^\circ$, reaching ~35 objects per 5° bin, and a secondary western peak around $\lambda \approx -105^\circ$, reaching ~22 objects per 5° bin.

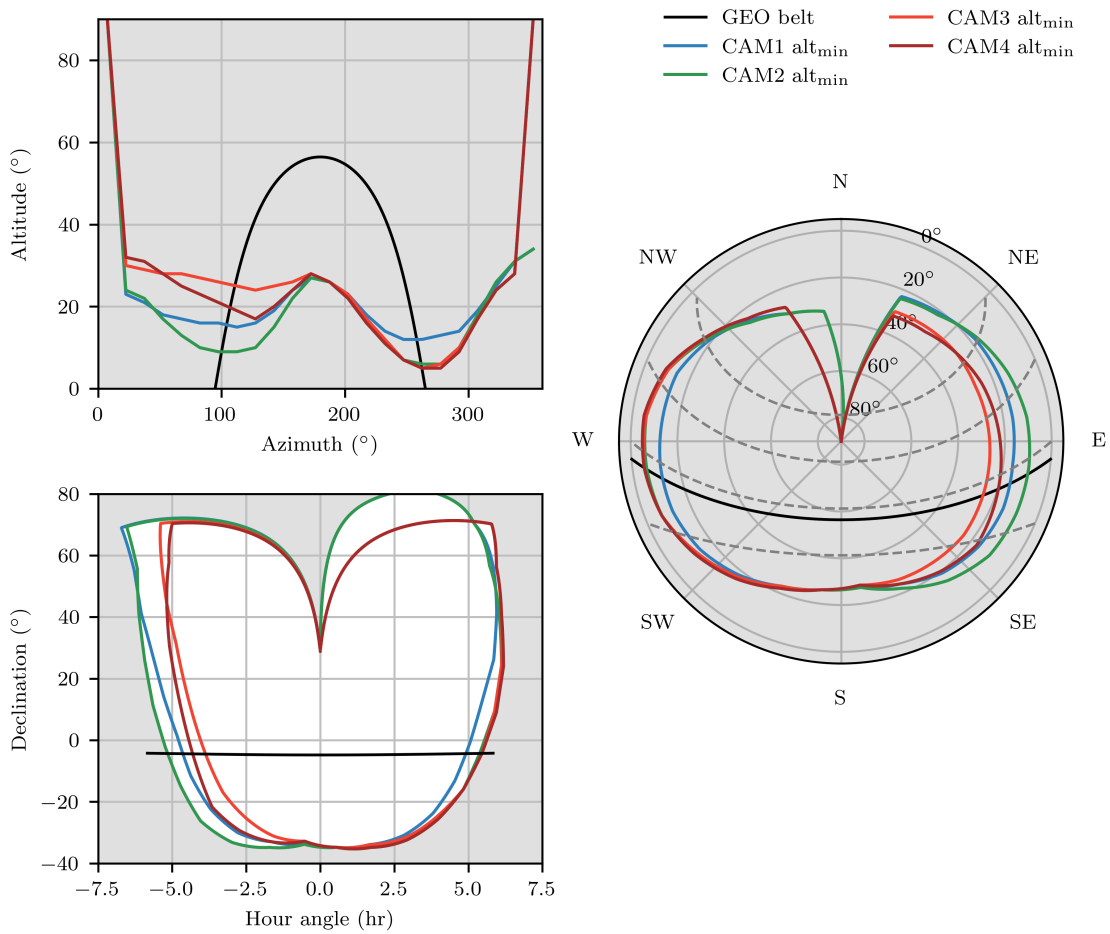

FIGURE 3.4: Combined obstruction and horizon-limit visualisation. **Top left:** Cartesian Alt–Az view showing the robust azimuth-dependent minimum-altitude boundary, $\mathrm{alt}_{\mathrm{min}}(\mathrm{az})$, for each camera, with the GEO belt overlaid. **Bottom left:** Equatorial view; the shaded region shows the union of points satisfying $\mathfrak{m}(\mathrm{HA}, \delta) = \mathrm{alt} - \mathrm{alt}_{\mathrm{min}} \geqslant 0$, while the curves show the extracted $\mathfrak{m} = 0$ horizon-limit contours for each dataset. The GEO belt is overlaid for reference. **Right:** Polar Alt–Az representation, with radius $= 90^\circ -$ altitude, showing the same $\mathrm{alt}_{\mathrm{min}}(\mathrm{az})$ boundaries together with the GEO belt. The dashed lines correspond to declination loci spanning from -20° to 40°.

From the ORM longitude of $\lambda_{\text{site}} \approx -17.9°$, geometric visibility, defined by alt $> 0°$, spans sub-satellite longitudes of approximately $-97°$ to $+62°$. This includes about 41% of the catalogue. Sub-satellite longitudes east of the site appear in the local eastern sky, while longitudes west of the site appear in the local western sky. Thus, the dense eastern stable point near $+75°$ is seen low in the east when first acquired, whereas the western peak near $-105°$ lies just beyond the western horizon and is not accessible from the ORM.

Using the recommended working limit of alt $\geqslant 20°$, the accessible sub-satellite longitude range narrows to approximately $-75° \lesssim \lambda \lesssim +40°$, containing 287 objects, or 28% of the catalogue. Extending down to the empirical floor of alt $\approx 5°$ increases the range to $-92° \lesssim \lambda \lesssim +56°$ and recovers a further 106 objects, corresponding to $\sim 10\%$ of the catalogue. Of these, 39 lie in the marginal $5° \leqslant$ alt $< 10°$ shells at the eastern and western limbs.

The east-west asymmetry is therefore most relevant at the eastern limb. Sub-satellite longitudes of $\lambda \approx +50°$ to $+56°$, which sample the outer wing of the Indian-Ocean cluster as it rises at azimuths of $\sim 85°$–$95°$, coincide with the strongest CAM1/CAM2 obstruction, where $\text{alt}_{\text{min}} \sim 10°$. CAM1 and CAM2 therefore lose access to part of this region, while CAM3 and CAM4 retain it. By contrast, the western limb, $\lambda \approx -92°$ to $-86°$, is observed as the eastward tail of the Pacific cluster sets at azimuths of $\sim 265°$–$275°$. This region falls in the cleaner western quadrant, where all four cameras reach $\text{alt}_{\text{min}} \approx 5$–$7°$.

For true geostationary objects, the Alt–Az position is nearly fixed, while inclined or eccentric GEO-like objects show only limited apparent motion over a night. The limb losses should therefore be treated as real camera-dependent incompleteness for marginal low-altitude targets, particularly on the eastern limb for CAM1 and CAM2. Since these losses occur only at the extreme visibility limits, they do not affect the core GEO survey region. The empirical horizon limits therefore provide conservative, camera-specific scheduling constraints. This closes the spatial requirement of §2.2.1. The wide-field is only useful to the extent that the mount, enclosure and horizon allow it to be pointed at the GEO belt, and the practical coverage is now quantified: of the 41% of the catalogued GEO population that is geometrically above the horizon from the ORM, 28% of the catalogued GEO population is accessible above the 20° working limit in all four cameras, rising to ~38% at the empirical floor of 5°, with camera-dependent losses confined to the eastern limb.

3.4 GPS Ephemeris Comparison

The astrometric requirement of §2.2.1 is arcsecond-level positional accuracy, well below the along-track error of propagated TLEs, and it was proposed there that this be verified against an independent high-accuracy reference. GPS satellites are the natural choice: they are the only resident space objects for which a centimetre-level precise ephemeris is publicly available, so any discrepancy between their measured and predicted positions is attributable to the instrument and pipeline alone. Although they occupy MEO rather than GEO, the test is of the measurement chain, not the orbital regime: the same detection, astrometric calibration and centroiding is applied to every target, and GEO objects, moving more slowly across the field, present an easier case than the trailed GPS streaks used here. The frame calibration, astrometric solutions and field stacking used in steps 1-3 of §3.4.3 are performed by the automated data reduction pipeline (Chote et al., 2025), developed by

Paul Chote. The GPS-specific analysis i.e., the observing plan, the SP3 interpolation and light-time solution, the streak rectification and centre-of-mass measurement, and the residual analysis (steps 4-8) was developed and carried out by myself.

3.4.1 GPS Ephemeris Data

Precise GPS satellite centre-of-mass ephemerides used in this analysis were obtained from the National Geospatial-Intelligence Agency (NGA) Office of Geomatics (National Geospatial-Intelligence Agency, 2025). These data are distributed in the SP3 precise ephemeris format, which provides satellite position and clock information at discrete time epochs. The SP3 format supports both position and velocity records and provides centimetre level positional precision and picosecond-level clock precision (Hilla, 2016). This data format differs from TLEs, which rely on a compact set of mean orbital elements and an analytical model to predict orbits with lower overall accuracy (see §1.2.1). The precise ephemerides are produced retrospectively rather than predicted: dual-frequency observations from a global network of ground monitoring stations are combined in a post-processed orbit determination, in which satellite positions and clock states are estimated simultaneously against dynamical and environmental models compliant with the IERS Conventions (2010) (IERS Conventions, 2010), including solid-Earth and ocean tidal corrections and tropospheric delay. NGA distributes the resulting product as the Department of Defense reference standard for GPS orbit and clock states, referenced either to each satellite's centre of mass or to its transmitting antenna phase centre; the centre-of-mass files are used here.

An SP3 file begins with a header containing metadata describing the dataset, including the start epoch, GPS week, epoch interval, coordinate reference frame, and the set of satellites included in the file. In the NGA products used here, satellite positions are provided in ECEF Cartesian coordinates referenced to the WGS-84 frame. The header also specifies the number of epochs and the sampling interval, which is typically 5 minutes for precise ephemeris products. An example of the SP3 data structure is shown below:

```
*  2012 10  1  0  0  0.00000000
P  1 -1677.901589 16285.584053 -20920.350325 275.849218
V  1 -25096.728698  8665.572021  8818.525202 0.015870
```

Satellite state information is provided for each epoch in a series of records. Each epoch begins with a header line identified by the symbol '*', which specifies the epoch time. The 'P' record contains the satellite identifier together with the satellite centre-of-mass position expressed in ECEF Cartesian coordinates in kilometres and the satellite clock offset, or clock bias, which represents the difference between the satellite's onboard clock and GPS system time in microseconds. The optional 'V' record provides the corresponding instantaneous satellite velocity components in the ECEF frame in kilometres per second and the satellite clock drift rate, defined as the time derivative of the clock offset. In this analysis, the velocity entries were used as endpoint derivatives for the cubic Hermite interpolation rather than as a constant velocity to propagate the satellite position across the full 5-minute SP3 interval.

For observation times that did not coincide exactly with a tabulated SP3 epoch, the satellite state was obtained by interpolation between the two bracketing SP3 epochs. Rather than assuming

constant velocity over the full SP3 sampling interval i.e., computing positions as $\mathbf{r}(t) = \mathbf{r}(t_0) + (t - t_0)\,\boldsymbol{v}(t_0)$, cubic Hermite interpolation was used, with the tabulated positions and instantaneous velocities at both endpoints of the interval, thus the Hermite interpolant uses the endpoint velocities to constrain the local cubic polynomial (Kreyszig, 2006). If the requested time coincided with an SP3 epoch, the tabulated state was used directly. Interpolation was only performed for times bracketed by available SP3 samples; no extrapolation outside the ephemeris time span was used, thus the interpolated position satisfies:

$$\mathbf{r}(t_0) = \mathbf{r}_0, \qquad \mathbf{r}(t_1) = \mathbf{r}_1,$$

and its derivative satisfies

$$\dot{\mathbf{r}}(t_0) = \boldsymbol{v}_0, \qquad \dot{\mathbf{r}}(t_1) = \boldsymbol{v}_1.$$

Writing

$$\tau = \frac{t - t_0}{t_1 - t_0}, \qquad \Delta t = t_1 - t_0,$$

the interpolated position is

$$\mathbf{r}(t) = h_{00}(\tau)\mathbf{r}_0 + h_{10}(\tau)\Delta t\,\boldsymbol{v}_0 + h_{01}(\tau)\mathbf{r}_1 + h_{11}(\tau)\Delta t\,\boldsymbol{v}_1,$$

where

$$h_{00} = 2\tau^3 - 3\tau^2 + 1, \qquad h_{10} = \tau^3 - 2\tau^2 + \tau,$$
$$h_{01} = -2\tau^3 + 3\tau^2, \qquad h_{11} = \tau^3 - \tau^2.$$

The interpolated velocity is obtained by differentiating the same polynomial with respect to time:

$$\boldsymbol{v}(t) = \dot{\mathbf{r}}(t) = \frac{h'_{00}(\tau)}{\Delta t}\mathbf{r}_0 + h'_{10}(\tau)\boldsymbol{v}_0 + \frac{h'_{01}(\tau)}{\Delta t}\mathbf{r}_1 + h'_{11}(\tau)\boldsymbol{v}_1.$$

Therefore the velocity samples are used explicitly through the $\Delta t\,\boldsymbol{v}_0$ and $\Delta t\,\boldsymbol{v}_1$ terms.

3.4.2 Observing Strategy

An observing plan was developed to track selected GPS satellites. Satellite positions were propagated using the `Skyfield` library. The observing sequence targeted four GPS satellites (NAVSTAR 62, NAVSTAR 64, NAVSTAR 81 and NAVSTAR 84). Each satellite was assigned a predefined observing window spanning the night of the 12th March 2026.

Satellite positions were evaluated at a cadence of one second to construct a high-resolution prediction of the satellite trajectory based upon the most up to date TLEs. TLE accuracy is required only for planning in order to keep the target within the field of view and to predict the streak geometry. The astrometric comparison itself is made against the precise SP3 ephemeris, which is independent of the TLEs; the one-second sampling simply ensures the planned trajectory is smooth on the timescale of the 2 s exposures. Observations were only permitted when the satellite altitude exceeded 30°, when the target was outside of Earth shadow and with a Moon separation $\geqslant 30°$. Time intervals that did not satisfy these constraints were discarded from the observing plan.

The observations using the `ObserveTLESidereal` action utilised 2 s exposures and were performed simultaneously using three cameras (CAM1, CAM2 and CAM4). Sidereal tracking was

used in preference to tracking on the TLE for three reasons: it keeps the field stars point-like, giving the best astrometric solution against which the streak is measured; the resulting streak converts position along the trail into time, so that the comparison tests the end-to-end timing chain as well as the astrometry; and the mount motion is then independent of the TLE, which enters only through the choice of field centre. CAM3 was not present in this analysis due to a fault in the power distribution unit at time of observation.

For each planned observation, the satellite position was evaluated both at the start and end of each exposure to determine the apparent trail vector in right ascension and declination. This allowed the expected streak length and orientation in the detector frame to be estimated.

To ensure that the satellite remained within the detector field of view (FOV), the predicted trajectory was divided into a sequence of observing fields. A new field was generated whenever the angular extent of the satellite path exceeded the detector footprint, defined by the FOV of STING. The telescope pointing for each field was defined by the midpoint of the satellite trajectory within the segment. Observations for each field began five seconds before the first predicted position and continued until the final position within the segment.

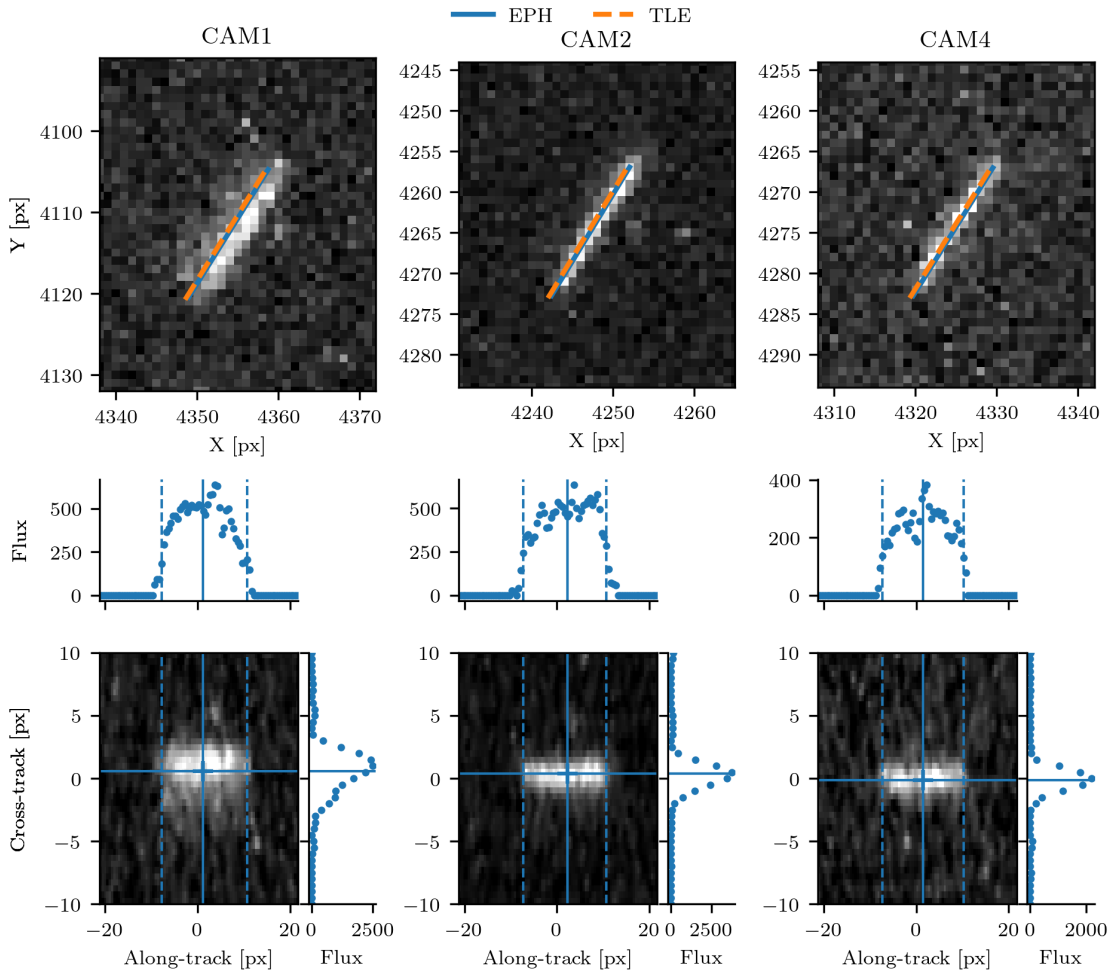

Figure 3.5: Predicted and measured morphology of a GPS satellite (NAVSTAR 81) streak observed in three cameras. **Top:** raw-frame zooms for CAM1, CAM2, and CAM4, centred on the identified streak, with the rolling-shutter-corrected GPS SP3 ephemeris prediction overlaid as a solid line and the TLE prediction overlaid as a dashed line. The SP3/EPH prediction uses iterative light-time correction to give the apparent satellite direction at reception, while the TLE prediction is obtained from pre-tabulated GCRS RA/Dec interpolation. **Middle:** summed along-track flux profiles from the rectified streak cutouts; the vertical lines mark the measured start, centre, and end positions. **Bottom:** background-subtracted rectified cutouts in a local along-track/cross-track coordinate system, with the corresponding cross-track flux profiles shown to the right. The vertical and horizontal guide lines indicate the measured along-track centre/extent and cross-track centroid used for comparison with the rolling-shutter-corrected predictions.

3.4.3 Position Extraction and Comparison

The position-extraction procedure is summarised as follows, and is illustrated for one streak in Figure 3.5:

1. The automated pipeline (described in Chote et al. (2025)) performed the early stages of the usual reduction steps. The `ObserveTLESidereal` action is reduced by grouping the raw frames by observed field and camera, combining the usable frames for each field into a stacked output image, and recorded the provenance of the contributing frames in the FITS headers.

In particular, the DATE-XXX header entries come from each input frame's DATE-OBS value, and the SEQN-XXX entries come from each input frame's GPS-SEQN value.

2. Candidate GPS streaks were supplied through a predefined list of identifiable stacked images and associated SEQN ranges. For each stacked FITS frame, the `DATE-XXX` and `SEQN-xxx` header cards were read, and only the entries whose SEQN values lay within the prescribed range for that stack were retained.

3. For every retained entry, the corresponding raw FITS frame was recovered from the timestamp encoded in the stacked frame `DATE-XXX` card together with the satellite and camera identifiers. The raw start time `DATE-OBS`, end time `DATE-END`, and row readout delay `ROWDELTA` were read from the raw FITS header, while the astrometric solution was taken from the associated `.head` file.

4. A single reference prediction of the satellite position was generated from the precise GPS SP3 ephemeris. SP3 position and velocity samples were parsed from the ephemeris files and interpolated in GPS time using cubic Hermite interpolation (see §3.4.1). For each requested time, a geometric light-time correction was solved iteratively, after which the satellite state was converted from ITRS to the topocentric GCRS frame at the observatory site. The satellite is not observed at its instantaneous position at the reception time t_{obs}, but at the earlier position from which the detected photons were reflected, at time t_{reflect}, where

$$t_{\mathrm{reflect}} = t_{\mathrm{obs}} - \tau, \tag{3.7}$$

and

$$\tau = \frac{\rho(t_{\mathrm{reflect}})}{c}. \tag{3.8}$$

Here, ρ is the satellite-observer range and c is the speed of light. Since the range depends on the satellite position at the unknown reflection time, the correction must be solved iteratively. In practice, the SP3 state is interpolated at a trial reflection time, the corresponding range is computed, and the reflection time is updated until convergence. The final apparent topocentric direction is then evaluated from this retarded-time state. The magnitude of this correction is not negligible. For the GPS satellites considered here, the light-travel time is typically of order

$$\tau \approx \frac{\rho}{c} \approx \frac{2.2 \times 10^4\,\mathrm{km}}{3.0 \times 10^5\,\mathrm{km\,s^{-1}}} \approx 7 \times 10^{-2}\,\mathrm{s}, \tag{3.9}$$

that is, about 70 ms. Over this interval, a GPS satellite moving at approximately $3.9\,\mathrm{km\,s^{-1}}$ travels

$$\Delta s \approx v\tau \approx (3.9\,\mathrm{km\,s^{-1}})(0.07\,\mathrm{s}) \approx 0.27\,\mathrm{km}, \tag{3.10}$$

which at a typical topocentric range of $\sim 2.2 \times 10^4$ km corresponds to an angular displacement on the sky of

$$\Delta\theta \approx \frac{\Delta s}{\rho} \approx \frac{0.27}{2.2 \times 10^4}\,\mathrm{rad} \approx 1.2 \times 10^{-5}\,\mathrm{rad} \approx 2.5\,\mathrm{arcsec}, \tag{3.11}$$

comparable to the astrometric residuals themselves. The light-time correction is therefore a necessary part of any arcsecond-level comparison with the SP3 reference orbit.

5. Because the detector is read out with a rolling shutter[1], the effective observation time depends on the image row. For a reference time t_{ref} and image row y, the effective time was taken as

$$t_{eff}(y) = t_{ref} + (y - y_{ref})\,\Delta t_{row}, \tag{3.12}$$

where Δt_{row} is the row delay from the FITS header and, in the present implementation, $y_{ref} = 0$. Start-, mid-, and end-of-exposure SP3 predictions were iteratively corrected for rolling shutter using the predicted row coordinate.

6. The rolling-shutter-corrected SP3 start, midpoint, and end predictions were used directly to define a guide streak in the raw image. A local rectified cutout was then formed in along-track and cross-track coordinates centred on the predicted streak midpoint by bilinear resampling of the raw image onto a regular (s, t) grid, and a local background level was estimated from the cutout border pixels and subtracted. The resampling slightly smooths the cutout and correlates neighbouring pixel noise, but, being linear in the data, it preserves the intensity-weighted first moment, so the centre-of-mass measurement of step 7 is not biased; only the diagnostic endpoint and width estimates are affected by the associated broadening.

7. The rectified streak was measured by a thresholded centre-of-mass estimator. A robust noise scale σ was derived from the background-subtracted cutout via the median absolute deviation, and pixels exceeding $n\sigma$ (with $n = 3$ by default) were retained as source pixels with weights equal to their background-subtracted values; all other pixels were masked. Intensity weighting is preferred to an unweighted centroid of the thresholded mask because it reduces sensitivity to the exact threshold: marginal edge pixels contribute in proportion to their flux rather than with full weight. The flux-weighted first moment also uses the signal-to-noise carried by each pixel rather than simply its presence above threshold. The streak midpoint (s_{com}, t_{com}) was then computed as the intensity-weighted centre of mass of the surviving pixels using `photutils.centroids`. Approximate streak endpoints were obtained from the along-track profile, formed by summing the thresholded weights over a narrow cross-track band centred on t_{com}, and taking the outermost samples exceeding a fixed fraction (default 25%) of the profile peak. This fraction lies well above the noise floor of the summed profile and below its plateau, placing the endpoints on the steep PSF roll-off at the streak ends and making them relatively insensitive to the precise choice of fraction. The endpoints serve only as a diagnostic: the half-length $L_{com} = \frac{1}{2}(s_{end} - s_{start})$ is compared against the predicted streak length, but does not enter the astrometric residuals, which use the centre of mass alone.

8. The measured midpoint was transformed to sky coordinates using the WCS, and the SP3 prediction was re-evaluated at the rolling-shutter-corrected time corresponding to the measured row. Residuals were then computed as the scalar angular separation, the tangent-plane components $(\Delta\alpha \cos\delta, \Delta\delta)$, and the projections of this vector onto the detector x/y axes and onto the local along-track and cross-track basis defined by the guide streak direction.

[1] Rolling shutter refers to an electronic readout method, not a physical/mechanical shutter.

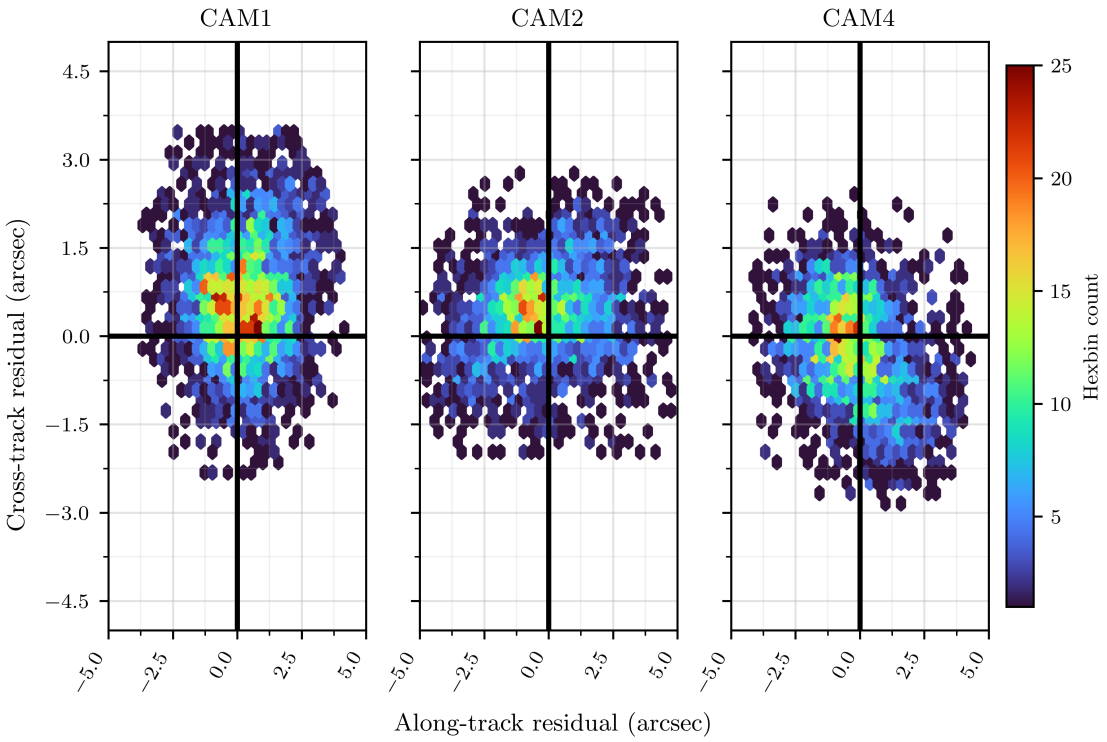

FIGURE 3.6: Hexbin density maps of the mid-exposure SP3 ephemeris residuals in local along-track and cross-track coordinates for CAM1, CAM2, and CAM4, shown from left to right. The black horizontal and vertical lines mark zero residual in each component, and the colour scale gives the number of measurements per hexagonal bin. All three distributions are centred close to the origin, showing no significant bias against the SP3 reference; the elongation along the horizontal axis reflects the larger along-track scatter of Section 3.4.3, and the slight upward displacement of the CAM1 distribution is the $+0.67''$ cross-track offset.

Camera	N_{used}	Retained (%)	Mean along-track offset $\pm\sigma$ [arcsec]	Mean cross-track offset $\pm\sigma$ [arcsec]
CAM1	3298	87.4	0.257 ± 1.326	0.674 ± 1.002
CAM2	2416	74.7	-0.224 ± 1.843	0.371 ± 0.810
CAM4	2257	92.4	-0.162 ± 1.544	-0.253 ± 0.907

TABLE 3.1: Camera-wise along-track and cross-track residual offsets after iterative 3σ clipping.

The procedure described above was applied to the full GPS dataset obtained with CAM1, CAM2, and CAM4. Measured midpoint positions were compared directly with the rolling-shutter-corrected SP3 reference prediction, and an iterative 3σ clip was applied to suppress extreme outliers.

The resulting EPH residual distributions are shown in Figure 3.6, and numerical summaries are given in Section 3.4.3. In all three cameras the distributions are approximately Gaussian in both directions, with skewness values of $|\gamma_1| \leqslant 0.24$ and excess kurtosis below 0.4. The along-track means are small ($\leqslant 0.26$ arcsec, $\leqslant 0.07$ px) and consistent with no systematic offset. Cross-track means are similarly small for CAM2 and CAM4, while CAM1 shows a modest offset of $+0.67$ arcsec (0.17 px) that likely reflects a minor camera-specific calibration effect. The scatter is dominated by the along-track direction, where standard deviations range from 1.33 to 1.84 arcsec

(0.33–0.46 px), compared with 0.81–1.00 arcsec (0.20–0.25 px) in the cross-track direction. This is expected: the centre of mass is tightly constrained across the narrow streak, while along the streak it responds to brightness variations along the trail and to any residual timing error, both of which act in the along-track direction only. CAM2 exhibits the broadest distributions and the lowest clipping retention rate (74.7%), consistent with it having the most asymmetric PSF. Overall, the pipeline delivers sub-pixel astrometric performance across all three cameras, with near-Gaussian residual distributions centred close to zero. Omitting the light-time correction displaces the along-track residual centroid by ≈ 3 arcsec, consistent with the estimate of Equation (3.11), confirming that the correction is essential.

3.5 SUMMARY

This chapter has characterised STING against the four observational requirements set out in §2.2.1, and each is now quantified rather than assumed.

The photometric requirement is met: a noise model built from the measured detector and site properties reproduces the lower envelope of the observed light curve scatter under both bright- and dark-sky conditions, with single-exposure precision of $\approx$1% for targets brighter than 9th magnitude to 10% at 14th magnitude. As anticipated, the coarse plate scale exacts its price only in the sky-limited faint regime, while the bright survey targets are scintillation limited. Because the four bands are recorded simultaneously, this precision applies directly to colour indices.

The spatial requirement is met with quantified limits: the empirical horizon and obstruction mapping shows the core GEO survey region to be unobstructed for all four cameras, with 28% of the catalogued GEO population accessible above the 20° working limit and camera-dependent losses confined to the extreme eastern limb. These limits are encoded as scheduling constraints, so survey completeness is set by geometry rather than by un-modelled instrumental effects.

The astrometric and timing requirements are met and verified against an independent truth: GPS satellites with centimetre-level SP3 ephemerides agree within a sub-arcsecond with one to two arcsecond scatter (0.2–0.5 px), an order of magnitude inside the TLE along-track errors that catalogue association must tolerate. Taken together, these results establish the performance envelope within which the science of the following chapters operates.

4 A Comprehensive Survey of the GEO-belt with STING

"The only thing that's for sure about
Sting is that nothing is for sure!"

Sting (Steve Borden), 1996

4.1 Motivations

The STING instrument was newly commissioned during my PhD such that it was a good opportunity to undertake a comprehensive survey of the GEO-belt with the instrument. This was the first large systematic colour survey of the GEO regime. A large portion of the content within this science chapter was taken from Airey et al. (2025a) with minor wording changes and the additions of §4.4.4.1 and §4.4.4.2. The technical background of the survey i.e., instrument control, scheduler and field geometry was undertaken by Paul Chote. The scheduler was run by myself and I performed all the analysis on the reduced light curves outlined later in this chapter.

Previous surveys of GEO have limitations. These limitations include works that implement observations of satellites across only one filter band (see e.g., Vrba et al., 2009; Wang et al., 2018; Cognion, 2013). Observations taken across a single filter will most likely be unable to probe features relating to materials or specific components given the dependency of reflectance on wavelength, whereas observations taken across multiple filters likely allow for the identification and characterisation of surface materials and mounted components (solar panels, antennas etc).

However, colour measurements of observed GEO satellites are affected by complex factors such as the illumination geometry, variations in the spectral energy density of reflected light near specular peaks, and the differing BRDFs (Bidirectional Reflectance Distribution Functions) (Hapke, 2012) of satellite surface materials (Zigo et al., 2023). Generally, it is found that using colour solely for satellite discrimination is extremely difficult given the variability of colour across a single night of observations (Jolley et al., 2015).

Despite these complexities, there have been numerous studies that have used colour measurements of GEO satellites to investigate and constrain properties such as ageing, bus configuration and manufacturing markers. Overall, the key findings from these studies are that dry masses and bus configurations seem to significantly impact colour measurements (Di Cecco et al., 2023) with evidence of clustering around certain colours within colour-colour diagrams and a null correlation between colour and time spent on-orbit (Schmitt, 2020). Multi-colour photometry can also be useful in characterising and classifying tumbling space debris given the relationships between rotational phase and colour indices (Silha et al., 2022).

There are common limitations of the studies described above. Often these surveys utilise telescopes which have a narrow FOV, which impacts upon the ability to conduct a large survey at GEO, i.e., observing as many GEO targets of interest as possible within a field of observation (Blake et al., 2023). Previous studies have also only focused on single object observations, which span over multiple months. Whilst this is indeed useful for characterising these objects individually and probing how seasonal and attitudinal variations impact light curve features, there are a very large number of satellites in the GEO regime and these vary in shape, size, design (usually in accordance with their bus configurations) and position on the sky. Therefore, it is difficult to apply the observations of a singular satellite across a range of satellites without many assumptions being made. In addition, multi-colour photometric surveys of GEO satellites have often made use of non-simultaneous colour measurements, which introduce systematic errors when the target's intrinsic brightness changes on a similar timescale to the observations. Another common limitation is that colour indices are often quoted at particular phase angles, neglecting the fact that the colour can change throughout the night as different components of the target become illuminated.

4.2 Observational Strategy

The observation campaign consisted of 17 full night observations between the 12th of April and the 18th of May 2023. Observations began at the start of nautical twilight (Sun altitude = -12 °), and continued until the end of nautical twilight the following morning. This corresponded to approximately 20:30 through 5:30 UTC during the observation campaign. Observations used a standard exposure time of 10s.

The GEO belt was tiled into 17 fields spanning hour angles between $\pm$ 70 °. A $\sim$ 2.5 ° overlap between adjacent fields ensures that the GEO targets of interest remain within the FOV of all four cameras for the entire night of observing (accounting for the small misalignment in the camera footprints, and the motion of targets throughout an observation series). The vertical FOV was cropped to 1 ° to reduce the data volume; this is sufficient for observing active objects that maintain low-inclination orbits, but sacrifices observation duration for debris and objects in inclined orbits (measured to have a duration at most around 170 minutes) as they pass through the limited vertical field. A single field was observed robotically each night, chosen to ensure that the moon did not approach within 30 ° of the field centre, and where possible to avoid the galactic plane (with very high background star density) entering the field during the observation. Due to visibility constraints from the walls of the enclosure, some fields were unable to be observed with some cameras, such was the case for the R_{RGB} camera observations in Field B. In this case, the observations taken by the camera were neglected. The observational fields and objects observed within are displayed in Figure 4.1.

4.3 Overview of Survey Demographics

The majority of geostationary satellites feature a 'box-wing' design, with large flat solar panels that have freedom to pivot around the north-south axis in order to track the Sun. Satellites of this type

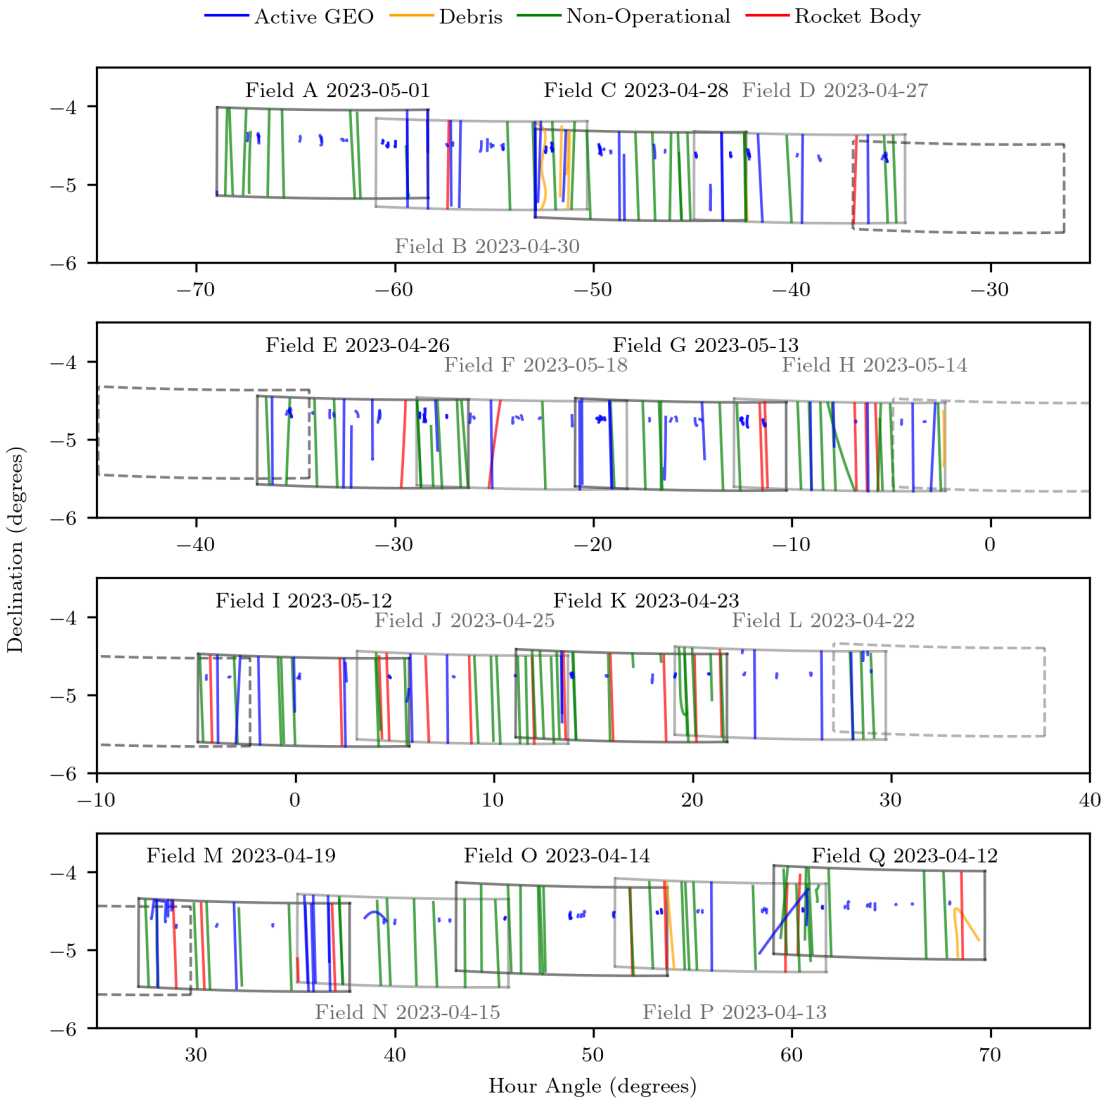

Figure 4.1: The observational fields (A-Q) used in this survey of the GEO belt in hour angle and declination space, with each field having overlap with the next. The tracks of each type of object observed in the survey are displayed by their respective markers in the figure legend. The borders of each field are alternated between black and grey in order to distinguish between them. Dashed field lines are drawn to display the fields on the previous/next row.

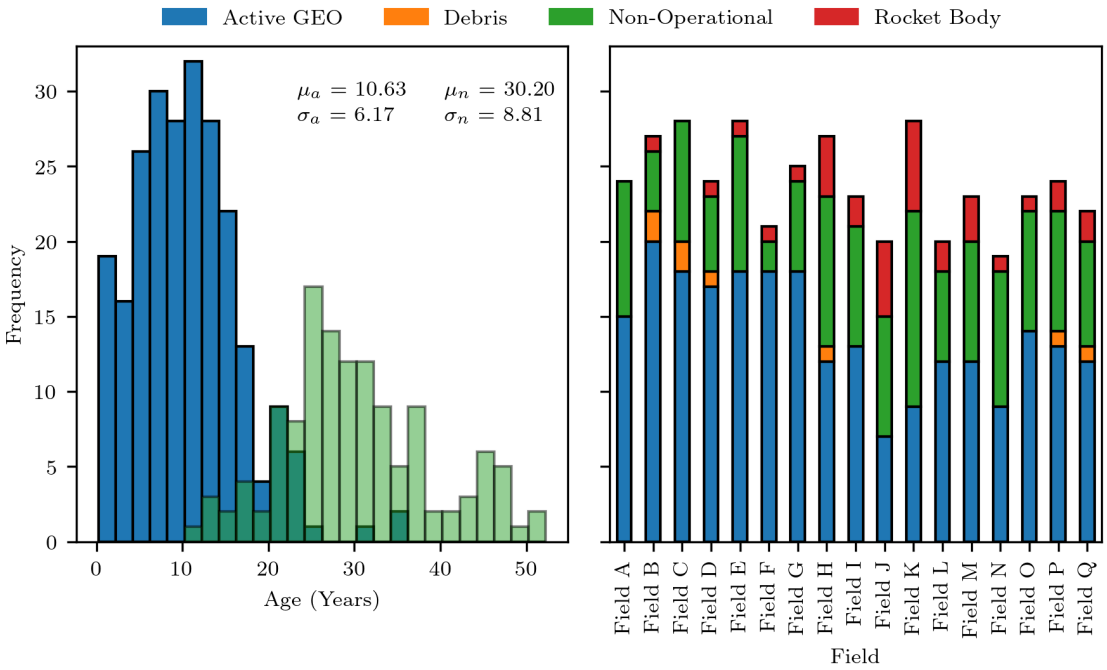

Figure 4.2: Demographics of observed objects upon conclusion of the survey. Left panel: Histogram displaying the overlapping age distributions of observed active and non-operational geostationary satellites in 2 year bins. Here μ_a, σ_a and μ_n, σ_n refer to the mean and standard deviation of the active and non-operational satellites respectively. A transparency is applied to the distribution of non-operational satellites. Right panel: Stacked bars indicating the type and number of objects observed within each field of the survey.

typically display strong specular features around solar equatorial phase angle (see §2.1.1.3) of 0° due to contributions of specular reflection from the solar panels (Payne et al., 2007).

The observations were carried out in the weeks following the March equinox. Around the equinoxes, the Sun-satellite-observer plane aligns with the rotation axis of the Earth, so the component of the phase angle perpendicular to the equator passes close to zero and the solar panels of box-wing satellites produce their strongest specular reflections – the so-called 'glint season' (Weisz et al., 2016). The survey period sits in the weeks immediately after this peak, where the geometry remains favourable for detecting specular reflections from the satellites without suffering from the extreme brightness changes that would otherwise saturate the detectors during the brightest glints.

The bus of a GEO satellite will typically have a fixed orientation relative to the Earth so as to maintain a consistent antenna footprint. Other light curve features can be possibly explained by reflections of the satellite bus or mounted components on said bus.

The analysis focuses on the active GEO satellite population, which spans a wide distribution of time on orbit, defined here as the elapsed time between a satellite's launch date, taken from Jonathan's space report (McDowell, 2024), and the survey epoch. This allows the possible relationship between time on orbit and reddening of satellites over time due to material degradation to be investigated. Figure 4.2 and Table 4.1 provide an overview of the full observational sample. 305 unique objects in total were observed, 168 of these were active satellites. Objects are identified by matching detections to TLE-predicted positions (§2.4.4), each carrying its NORAD catalogue identifier; uncatalogued objects are therefore not searched for. For objects at lower elevations

Field ID	ACT	DEF	RB	DEB	Total
A	15	9	0	0	24
B	20	4	1	2	27
C	17	9	0	2	28
D	17	5	1	1	24
E	18	10	1	0	29
F	19	2	1	0	22
G	17	7	1	0	25
H	11	11	4	1	27
I	11	9	2	0	22
J	7	8	5	0	20
K	9	13	6	0	28
L	11	7	2	0	20
M	12	8	3	0	23
N	9	9	1	0	19
O	14	8	1	0	23
P	13	8	2	1	24
Q	12	7	2	1	22
Total Observations	232	134	33	8	407
Unique Objects	168	109	22	6	305

Table 4.1: Summary of objects observed in each of the 17 survey fields. The different types of object can be broken down into active geostationary satellites (ACT), non-operational satellites (DEF), rocket bodies (RB) and debris (DEB). Objects may be observed in adjacent fields if they lie within the overlapping region, this means that the unique total of objects is less than the sum of observations.

(fields A, B, and Q), one or more of the cameras were obscured by the dome horizon and those specific colour bands are entirely absent. The 15 active satellites observed in Field A were missing two colour bands so were excluded. 37 active satellites on inclined orbits were excluded due to insufficient observation length and 4 satellites which were missing at least one colour band were excluded also, such that the final sample was 112 objects. Table 4.2 displays the different types of bus configuration associated with the final 112 objects. Figure 4.3 displays an example of a processed four colour light curve of the INTELSAT 901 (+ MEV-1) satellite observed on the 25th April 2023 with STING. This observation is taken after the docking of INTELSAT-901 with MEV-1 and therefore the colour signatures within could potentially be used to distinguish between INTELSAT-901 and MEV-1 components. There are notable colour features within this light curve which will relate to objects that were analysed in later sections, particularly §4.4.4. This data demonstrates the ability of the STING instrument to resolve these signatures.

Figure 4.4 and Figure 4.5 present a selection of G_{RGB} band light curves of satellites belonging to the SSL-1300 and Eurostar-3000 bus classes respectively, the two most populated bus classes in the sample. These illustrate the variety of features the survey light curves contain i.e., broad phase-angle structure, glints of differing sharpness and location, and object-to-object variation within a single bus class.

Bus Configuration	Count	Bus Configuration	Count
Eurostar-3000	26	GEOSTAR-3	2
SSL-1300	24	Eurostar-NEO	2
STAR-2	10	Spacebus-4000B3	2
BSS-702HP	5	LUXOR	2
Spacebus-4000B2	4	Ekspress-2000	2
BSS-702MP	4	Eurostar-2000+	2
Spacebus-4000C3	3	Spacebus-3000B3	2
Spacebus-Neo-200	3	Spacebus-3000B2	1
Spacebus-4000C4	3	Spacebus-Neo-100	1
A2100AXS	3	LM-2100	1
A2100A	1	SS-702	1
Eurostar-NEO	2	AMOS	1
Spacebus-4000B3	2	DS2000	1
LUXOR	2	Ekspress-1000N	1
Ekspress-2000	2	Eagle-Bus ESPA (ESPAStar-D)	1
Eurostar-2000+	2	BSS-702SP	1
Spacebus-3000B3	2	Ekspress-1000HTB	1
Spacebus-3000B2	1	ARSAT-3K	1
Spacebus-Neo-100	1	DFH-4	1
Grand Total	112		

Table 4.2: Count of bus configurations for the unique active GEO sample of satellites for which photometry was available in all four bands (B_{RGB}, G_{RGB}, R_{RGB} and i_{sdss}).

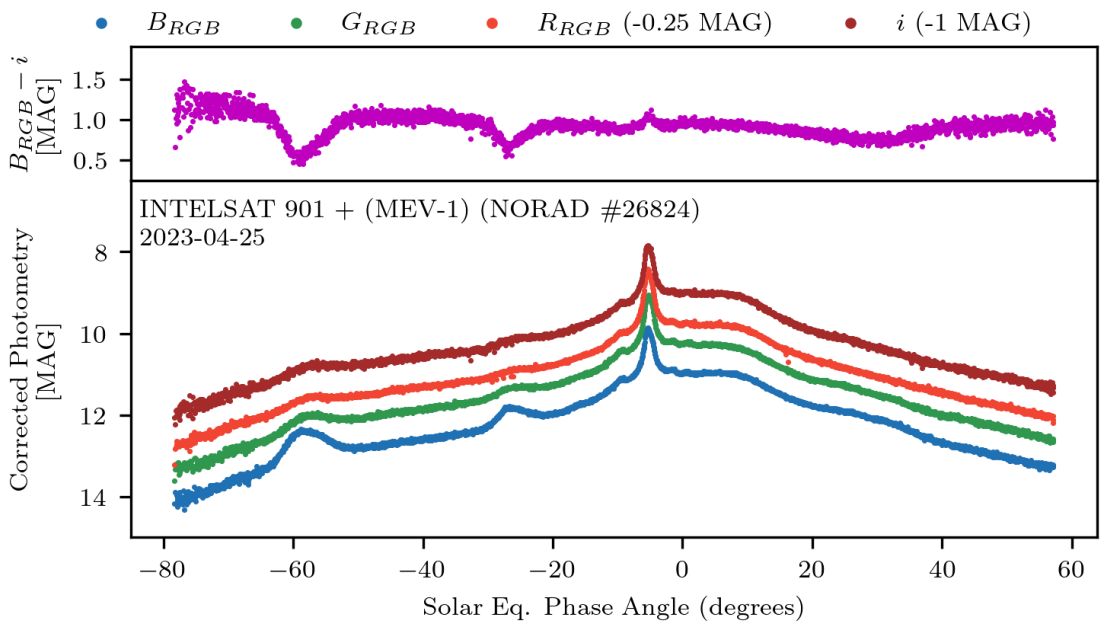

Figure 4.3: An example four colour light curve of INTELSAT 901 (+ MEV-1) on the 25th April 2023. The upper panel displays the $B_{RGB} - i$ colour index with the lower panel displaying the corrected photometric measurements in the four colour bands. A 0.25 and 1 mag offset is applied to the R_{RGB} and i band measurements respectively to aid readability.

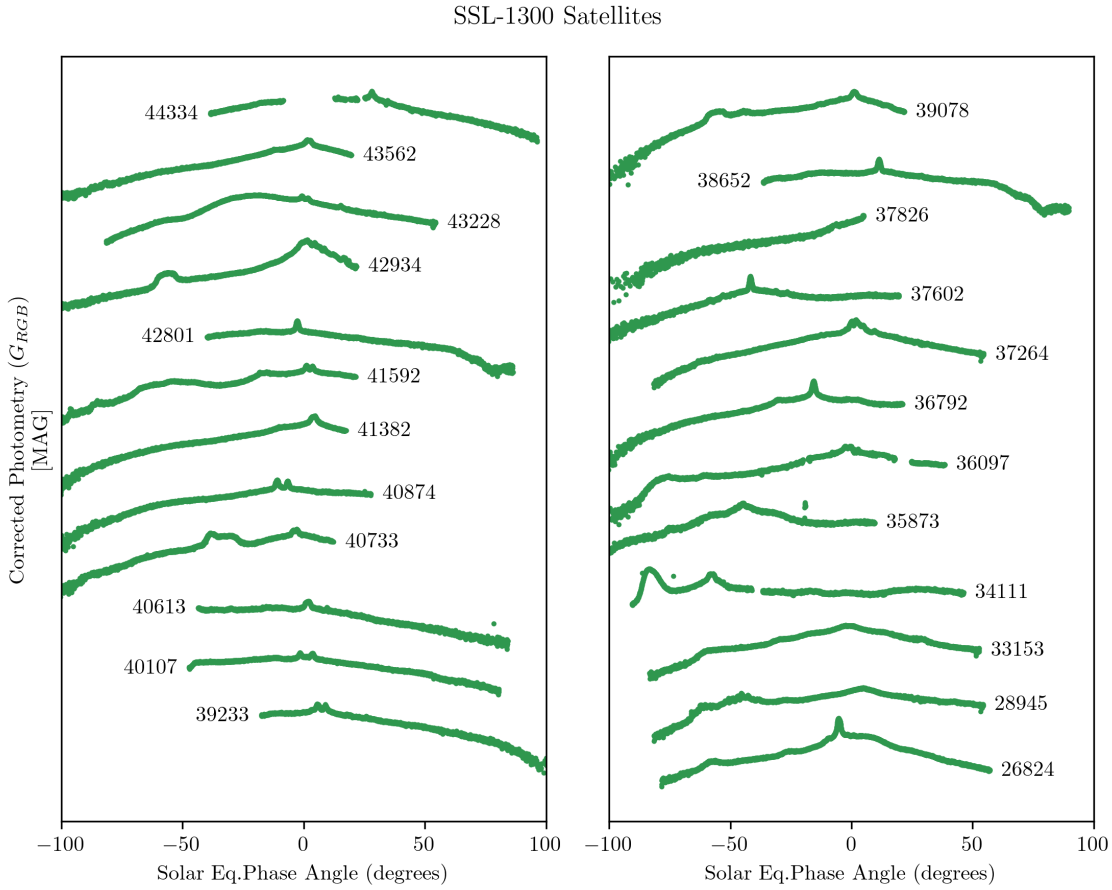

Figure 4.4: Illustrative example of G_{RGB} band light curves for SSL-1300 class satellites with an offset of 4 magnitudes applied to each subsequent light curve to avoid overlap. Gaps within a few of the light curves are due to filtering out bad zero-point measurements. The NORAD ID of each object is shown adjacent to its light curve.

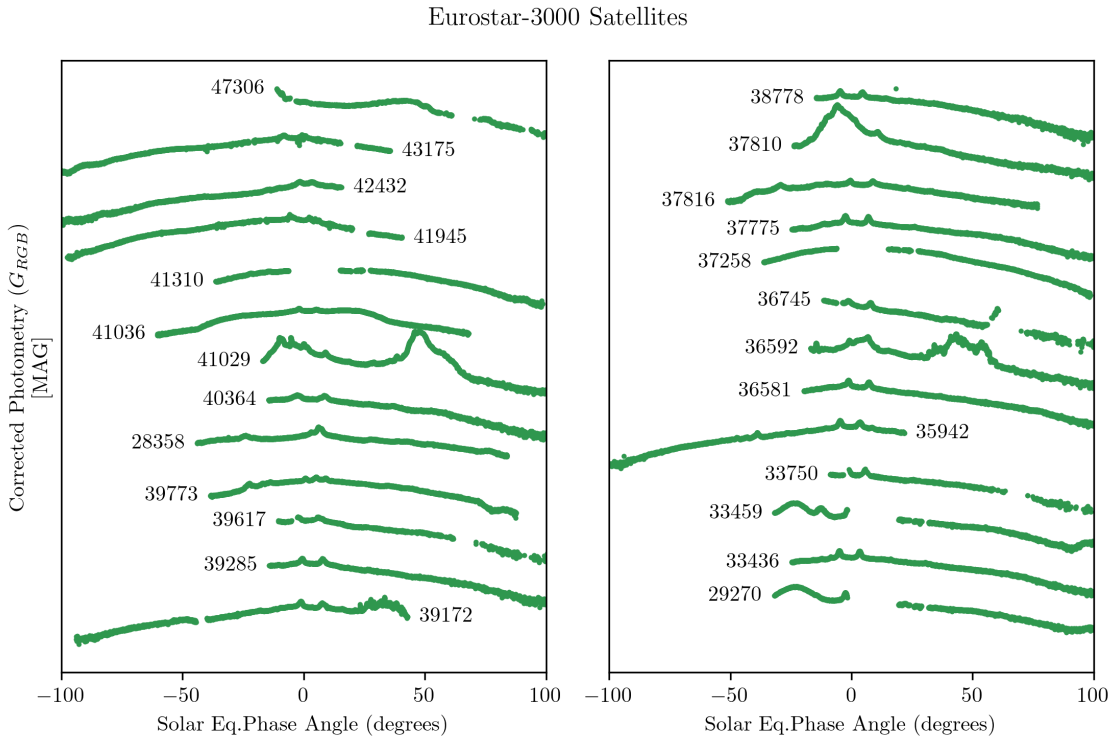

Figure 4.5: Illustrative example of G_{RGB} band light curves for Eurostar-3000 class satellites with an offset of 4 magnitudes applied to each subsequent light curve to avoid overlap. Gaps within a few of the light curves are due to filtering out bad zero-point measurements. The NORAD ID of each object is shown adjacent to its light curve.

4.4 Results and Analysis

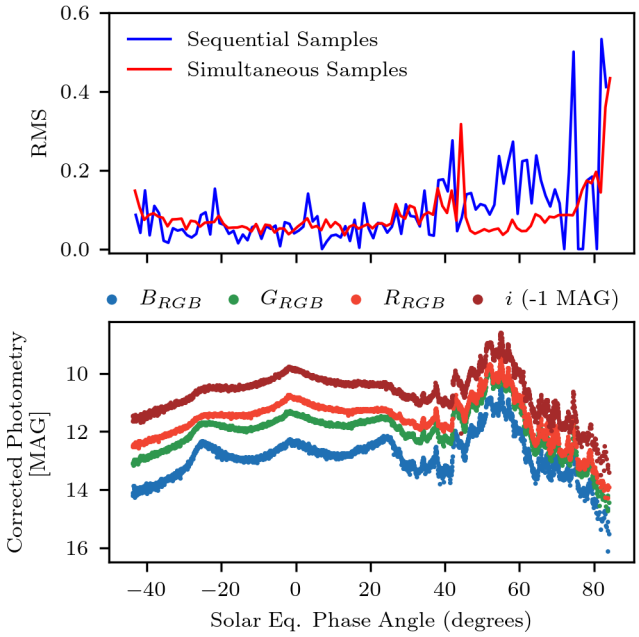

Figure 4.6: Upper: Binned RMS values calculated for the sequential and simultaneous measurements of $B_{RGB} - i$ within bins of width of 5 minutes. Lower: Original THOR-6 light curve with truth simultaneous measurements. The sequential measurements have a definitively larger RMS over the glinting region due to not being able to effectively resolve these short timescale features. A 1 mag offset is applied to the i band measurements.

4.4.1 Induced Colour Error

Sequential colour measurements can induce an error when the object brightness is rapidly changing with respect to the filter cycle, this is referred to as a induced colour index error (Castro et al., 2022).

It is appropriate to demonstrate the improvements in the resolution of light curve features achieved through simultaneous colour measurements, as compared to the sequential colour measurements predominantly used in previous studies. How and where simultaneous observations enhance the resolution of features within the light curves is analysed below.

Figure 4.6 is an example of the light curve of THOR-6 with a binned RMS measurement taken in bins with widths of 5 minutes. For the majority of the light curve, the two colour features have a negligible deviation difference when comparing the simultaneous and sequential observations. However, as the region between 40-75 ° is approached, where the brightness changes on short timescales, there is a larger associated RMS with the sequential data measurements, this highlights the importance of simultaneity such that simultaneous colour measurements are most useful across short timescale features and these features in particular are likely to have interesting colour behaviours to investigate.

Given there were numerous objects in the dataset which show regions of features occurring on short timescales, the induced colour index error associated with the sequential data measure-

NORAD ID	σ_{seq} (M)	σ_{sim} (M)	Δ (%)
42934	0.113	0.086	27.1
29273	0.103	0.093	10.6
41866	0.105	0.112	-6.0
39168	0.082	0.064	24.4
39172	0.211	0.103	68.5
28945	0.156	0.105	39.0
45027	0.123	0.124	-0.50
44457	0.161	0.102	44.2
37836	0.222	0.188	16.7
41029	0.255	0.188	30.4
36592	0.212	0.145	37.5
40258	0.136	0.072	61.9
36830	0.183	0.149	20.9
36033	0.307	0.229	29.2
36831	0.130	0.098	27.5
39163	0.443	0.367	18.7

Table 4.3: Standard deviation measurements for both sequential and simultaneous colour measurements ($B_{RGB} - i$) of short timescale glinting regions within the light curves of the named satellites. The standard deviation measurements are in units of magnitudes (M). The percentage difference between the standard deviations is denoted as Δ. The percentage difference is calculated as $\frac{a-b}{0.5(a+b)} \times 100$, where a and b are the standard deviation measurements of the sequential and simultaneous samples respectively.

ments was characterised and measured. Table 4.3 shows the measured standard deviation of the colour index ($B_{RGB} - i$) for various satellite glinting regions for both simultaneous and sequential observations. Generally, the standard deviation increases significantly between the simultaneous and sequential data measurements. The average percentage difference and increase between the simultaneous and sequential measurements are approximately 28% and 36% respectively. The GOES 16 (NORAD: 41866) and EUTE KONNECT (NORAD: 45027) satellites actually show a reduction in the error, however, this can be explained as the glinting seen within their light curves is occurring over slightly longer timescales than the others (see §4.4.4.5). It is evident that simultaneous colour measurements allow for a higher resolution of these short timescale glinting regions.

4.4.2 Colour Contrasts Between Bus Configurations

To further analyse the observed differences between satellites of different bus configurations, the approach was taken of investigating two colour indexes ($B_{RGB} - i$ and $G_{RGB} - R_{RGB}$) as a function of solar equatorial phase angle for each satellite. A consideration for this again is to exclude the $G_{RGB} - R_{RGB}$ measurements for satellites observed within field B (large variability in the R_{RGB} zero-point).

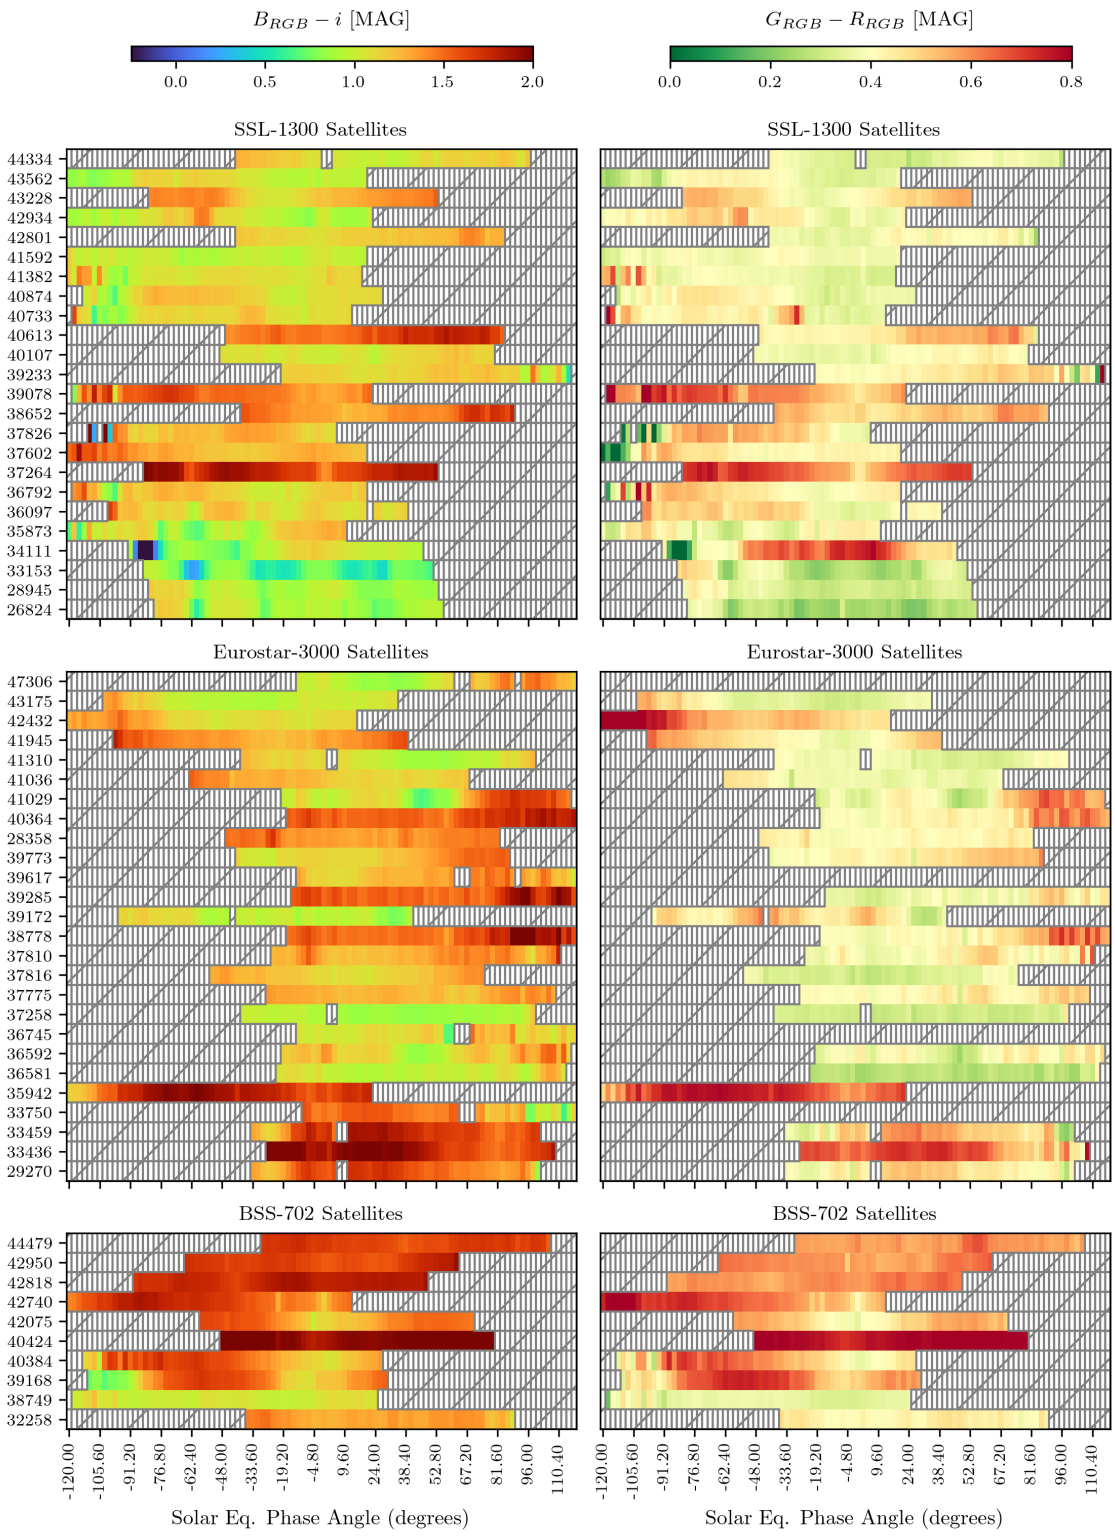

Figure 4.7: Colour light curves transformed into 2D image data for $B_{RGB} - i$ and $G_{RGB} - R_{RGB}$ colour indices. Each row represents the colour light curve of a satellite object identified by its given NORAD ID and each block of objects belong to a particular bus configuration (SSL-1300, Eurostar-3000 and BSS-702). Each colour light curve is split into 100 bins of equal solar equatorial phase angle width. The row order is determined by increasing time on orbit and entire missing rows correspond to objects observed in field B.

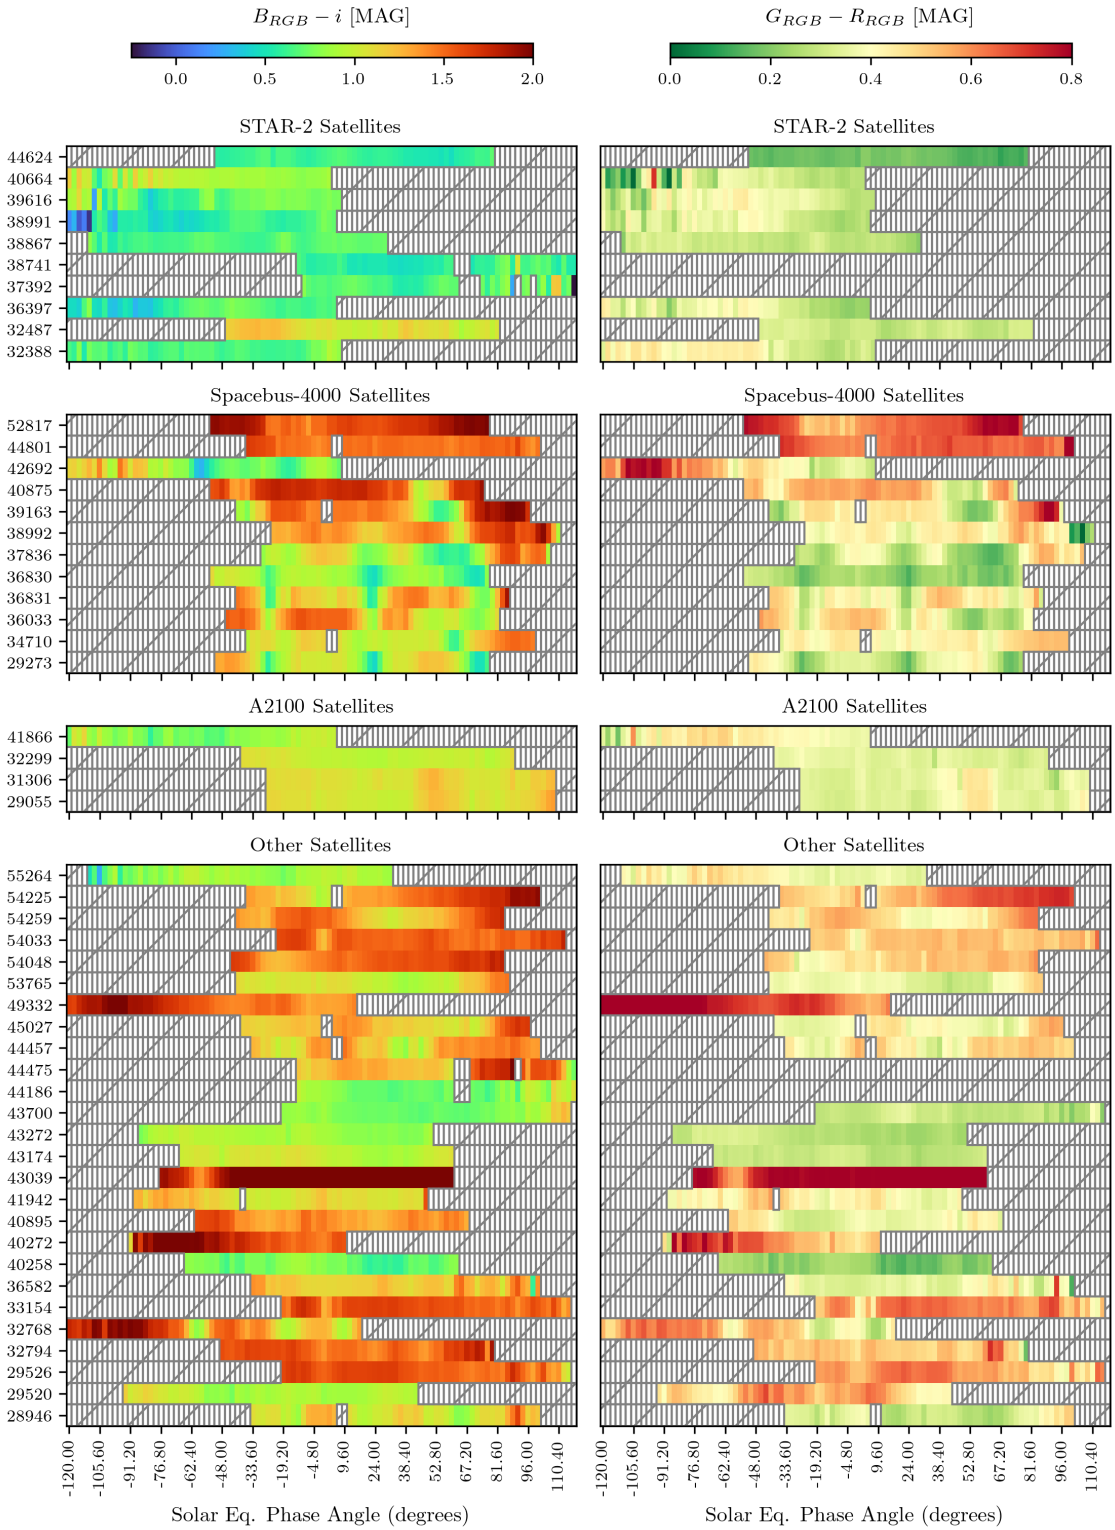

Figure 4.8: Colour light curves transformed into 2D image data for $B_{RGB} - i$ and $G_{RGB} - R_{RGB}$ colour indices. Each row represents the colour light curve of a satellite object identified by its given NORAD ID and each block of objects belong to a particular bus configuration (STAR-2, Spacebus-4000, A2100 and satellites belonging to other classes). Each colour light curve is split into 100 bins of equal solar equatorial phase angle width. The row order is determined by increasing time on orbit and entire missing rows correspond to objects observed in field B.

The phase range of -120 ° to 120 ° was chosen as this captured most if not all of the light curve datapoints of the sample targets. The colour light curves were split into 100 bins ($\simeq$ 10 minute bin width, small enough to not lose any resolution on significant features). A suitable colour map range for $B_{RGB} - i$ and $G_{RGB} - R_{RGB}$ respectively could be found, ensuring comparison between every bus configuration and produce the 2D image plots shown in Figure 4.7 and Figure 4.8. These colour light curve maps are particularly useful for displaying the variability of colour across an observation as well as making the identification of features shared between objects easier.

The colours of GEO satellites are expected to vary both across different bus configurations and throughout the night for individual objects. Starting with the SSL-1300 class satellites, a few interesting features can be commented upon. Generally the value of $B_{RGB} - i$ tends to be fairly consistent for individual satellites that were launched more recently and this value of colour index is comparable to that seen in Eurostar-3000 and A2100 satellites. For the objects which have spent more time on orbit, there tends to be 'bluer' features present and in general more variability in the $B_{RGB} - i$ index values. The value of the $G_{RGB} - R_{RGB}$ colour index shows greater variability across all SSL-1300 satellites, however the same cluster of objects with the intense $B_{RGB} - i$ excursions (excursion referring to an increase or a decrease in the colour index value) are 'greener' than the other SSL-1300 satellites.

The Eurostar-3000 satellites display an effect of the satellites which have spent the most time on orbit displaying much 'redder' colours with respect to satellites of the same bus class that were launched more recently. The Eurostar-3000 satellites launched after AMAZONAS-2 (NORAD: 35942) don't reach the same values of $B_{RGB} - i$ consistently. However, looking at the $G_{RGB} - R_{RGB}$ index, whilst Eurostar-3000 satellites are 'greener' with respect to the cluster of objects launched before ASTRA-3B (NORAD: 36581), the amount of colour difference between this group and the others is not as drastic as shown in the $B_{RGB} - i$ index.

BSS-702 satellites tend to be 'redder' across the entire phase range. This behaviour is also seen in the $G_{RGB} - R_{RGB}$ colour index. Contextually, their $B_{RGB} - i$ indices are comparable to that of the cluster of red Eurostar-3000 satellites. Conversely, STAR-2 satellites are consistently 'bluer' across the entire phase range than any of the other satellite platforms within the $B_{RGB} - i$ indices. The $G_{RGB} - R_{RGB}$ values tend to be only marginally smaller than satellites within other bus classes.

In the Spacebus-4000 class, a few satellites tend to show similar colour features at similar solar equatorial phase angles with some near-symmetry between features around $\pm$ 25 °. These features are apparent in both colour indices and are instances in which the satellite appears 'bluer'. Conversely, the sample size of A2100 satellites is small however, generally the colour of these satellites tends to be fairly consistent across the entire phase range. The two older satellites (NORAD IDs: 31306 and 29055) show a red feature around 60 ° of phase. The A2100 also display slightly 'greener' colours compared to other satellite bus configurations.

4.4.3 COLOUR-COLOUR DIAGRAM

Figure 4.9 displays the average colour measurements of satellites belonging to a specific bus class within a specific phase angle range of their light curves. The colour indices ($B_{RGB} - G_{RGB}$ and $R_{RGB} - i$) were chosen as they best demonstrate the behaviour seen between classes at the very

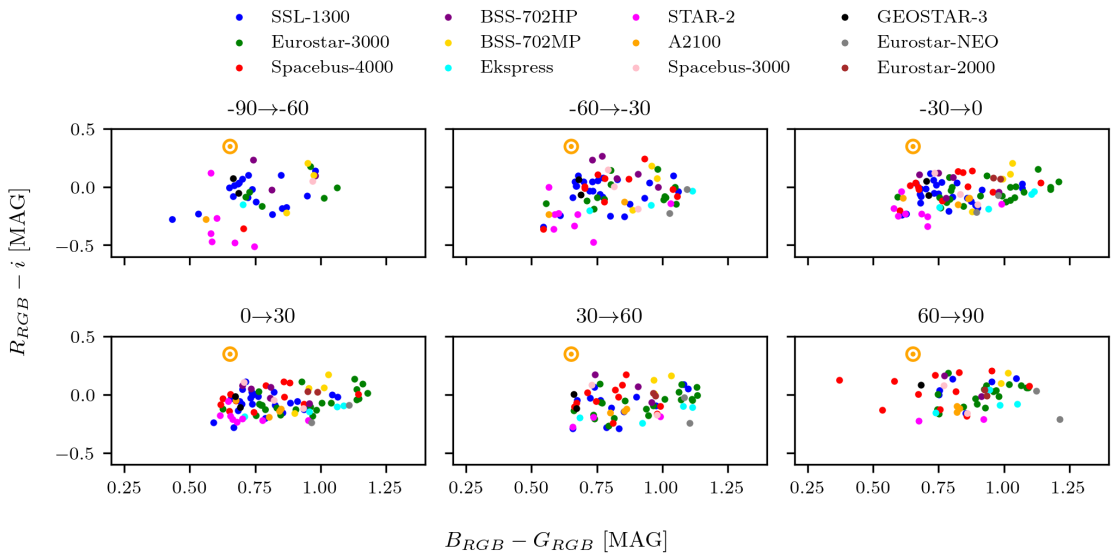

Figure 4.9: Colour-colour diagrams showing the average colour measurements of individual satellites belonging to a particular bus configuration class within given solar equatorial phase angle ranges from $\pm$ 90 °. The error-bars were significantly small on these measurements, so they are not displayed. The average colours of the Sun are denoted by the orange solar marker on each of the diagrams.

blue ($\sim$ 400 — 550 nm) and very red ($\sim$ 600 — 850 nm) regimes across each axis respectively. There is significant overlap in the colour measurements for many of the bus configurations, which indicates characterisation may be difficult based on the colour alone for identification of many of the bus classes. In particular, as the phase angle varies between ± 30 °, the colours of the satellites of varying bus configurations cluster closely together. This highlights the difficulty in using colour measurements to discriminate between bus configurations overall. This overlap is itself an informative result: it quantifies the qualitative impression of Figure 4.7 and Figure 4.8 – that colour, averaged within phase-angle ranges, is insufficient to separate most bus classes, because object-to-object colour offsets within a class are comparable to the differences between classes. A differential analysis, measuring each satellite's colour change across phase angle relative to its own quiescent baseline, would remove these offsets and may provide a stronger discriminant.

All of the satellite colour measurements overall have 'redder' colours than the solar spectrum, which is in agreement with the measurements from Schmitt (2020).

4.4.4 Light Curve Features

4.4.4.1 LDPE-3A Manoeuvres

LDPE-3A (ROOSTER-3) is a satellite designed with two experimental payloads for space-domain awareness[1]. The first is named Catcher, which is designed to demonstrate new miniaturised technology capable of diagnosing the adverse effects of radiation, charged particles, and other space weather events on spacecraft in orbit. The second is WASSAT, this is a wide-area sensor with four cameras that searches for and tracks other spacecraft and space debris in geosynchronous orbit, where communications, missile detection, intelligence-gathering and weather monitoring

[1] https://space.skyrocket.de/doc_sdat/ldpe-1.htm

satellites operate. It also has the most unique design of all satellites within the dataset given it is a ring-shaped spacecraft with a singular solar panel, LDPE-3A and GOES 16 are the only satellites in the sample with singular solar panel designs.

The phase and colour light curves of LDPE-3A are presented in Figure 4.10 and denoted are several interesting features. Features denoted by **(a)** and **(b)** are likely body manoeuvres of LDPE-3A in such a way that the area of LDPE-3A presented to STING has decreased such that there is a sharp decrease in brightness. Although the differing factors between both features is that there is a slight increase in the brightness before the sharp decrease in **(b)**, which is shown by the zoomed inset axis placed in Figure 4.10. An instrumental or atmospheric origin can be excluded: the photometry is calibrated against field stars in each individual frame (§2.4.3), the zero-point scatter is stable across these intervals, and no other objects in the same field show corresponding features.

Both of these features are associated with slight shifts towards 'redder' colours which could be reflections off the gold MLI covering of some of the instruments. Although **(b)** does show a shift towards 'bluer' colours whilst the brightness decreased. The inset axes which hold the unbinned datapoints of this portion of the light curve does reveal that during the decrease in brightness, there is a short lived glint where the brightness increases which suggests subtle changes in the attitude, i.e., more surface area is in the line of sight of STING. Following **(a)** there is a slight step up in the light curve which suggests again a slight attitude adjustment of LDPE-3A. It would also suggest that this is not a station-keeping manoeuvre such that rather than returning to the same orientation it now lies in a different orientation. Given that the LDPE series of satellites have instruments aboard for space domain awareness, it could be possible that it is now tracking a different object.

Overall there is not much variability in the colour and the colour could be comparable to that of the STAR-2 satellites. Given that these were designed by the same manufacturer (Orbital Science Corp, later Northrop Grumman), a shared surface-material heritage is therefore plausible, but cannot be confirmed from unresolved photometry alone, and the colour similarity should be read as consistent with, rather than evidence of, common materials.

Similar signatures to **(b)** can also be seen in the light curves of the MEV-2 satellites during proximity operations with Intelsat 10-02 (Chote et al., 2025). The features identified in this light curve could be useful for future identifications of manoeuvres in other light curves of GEO satellites.

4.4.4.2 Fitting Manoeuvres

A selected manoeuvre region in the light curve of LDPE-3A (as seen in Figure 4.10) is modelled using a facet-based spacecraft geometry (see Figure 4.11) and a parametric attitude history. The unknown attitude parameters are estimated by minimising the mismatch between observed and modelled brightness.

The measured flux in this segment is denoted by

$$y_i = F_{obs}(t_i), \qquad i = 1, \dots, T.$$

A pre-manoeuvre reference flux, F_{pre}, is estimated from data preceding the dip, and both observed

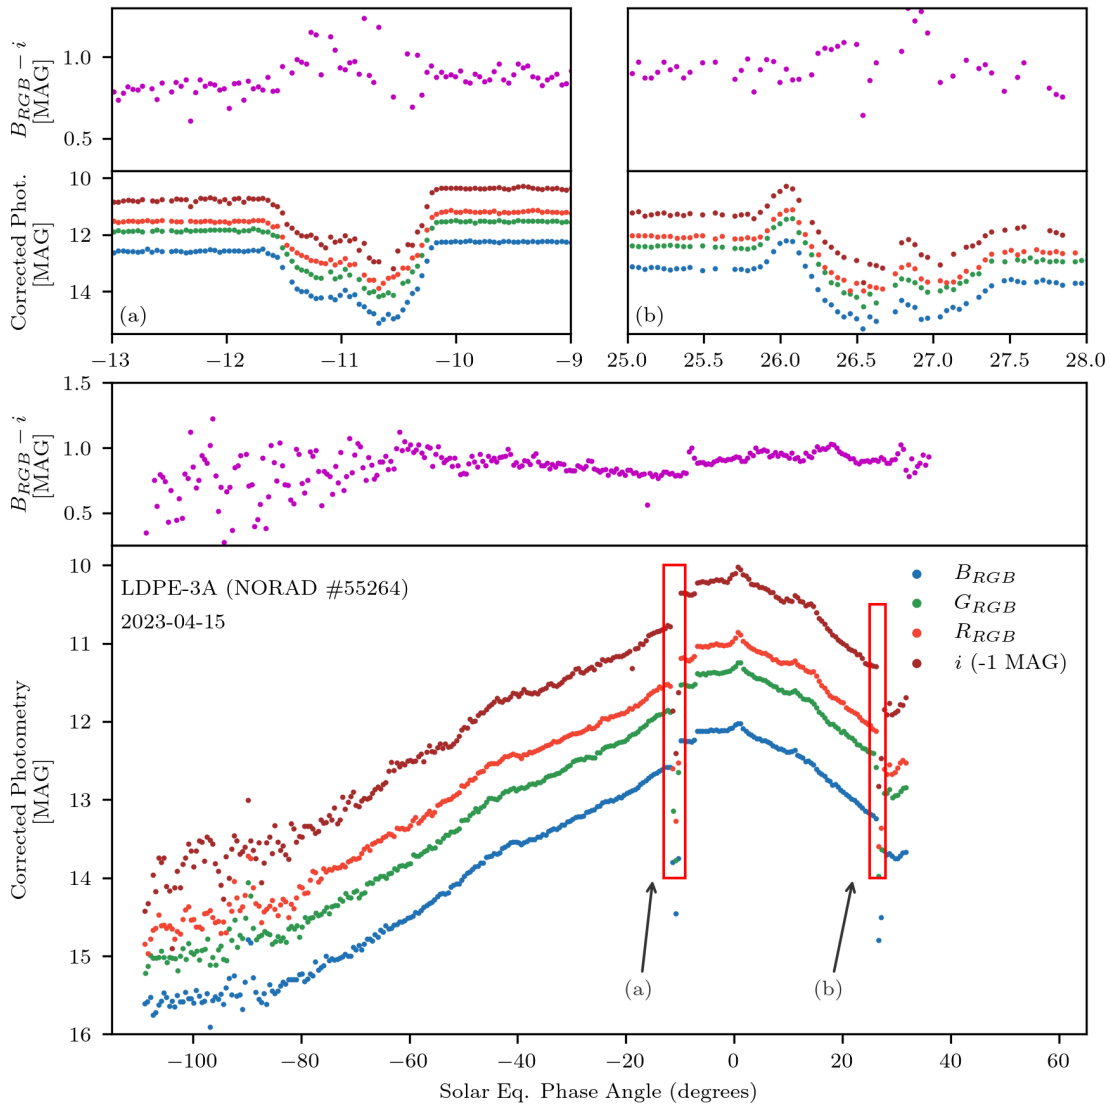

Figure 4.10: Binned (bin widths of 7.5 minutes) four colour phase light curves for satellite LDPE-3A (NORAD: 55264) observed on the 15th April 2023. At this point, LDPE-3A had only spent 90 days on orbit. Interesting features are labelled by **(a)** and **(b)**. It is suspected that each of these features are caused by the manoeuvring of LDPE-3A about its own axes. The upper two panels show the zoomed in and unbinned region of the light curve for features annotated by **(a)** and **(b)** respectively with these regions being shown by the boxes drawn in the lower panel. Slight stepped features are apparent across the phase light curve which further suggests slight manoeuvring. A 1 mag offset is applied to the i band measurements for clarity.

and modelled fluxes are compared in normalised form:

$$r_i^{\mathrm{obs}} = \frac{y_i}{F_{\mathrm{pre}}}.$$

For each time bin t_i, the Sun direction in ECI coordinates is computed, and an orthonormal reference frame is constructed whose z-direction is aligned with the Sun. The line-of-sight unit vector to the observer, obtained from right ascension and declination, is projected into this frame to give $\hat{\mathbf{o}}_{b0,i}$.

This defines the observer direction before the attitude perturbation is applied.

A triangular mesh of the spacecraft is loaded from an .OBJ file that I created myself. Facets are divided into two groups: array facets, identified by a specified RGBA colour, and bus facets, defined as the remaining facets. For computational efficiency, only the largest bus facets, ranked by area, are retained. For each subset, facet normals and areas are stored as

$$\{(\hat{\mathbf{n}}_k^{\mathrm{arr}}, A_k^{\mathrm{arr}})\}, \qquad \{(\hat{\mathbf{n}}_k^{\mathrm{bus}}, A_k^{\mathrm{bus}})\}.$$

The spacecraft attitude over the fitting interval is represented by a set of parameters consisting of a baseline roll motion, two transient tilt manoeuvres, a roll step, and a small bias tilt. In this model, the observer direction in the body frame is written as

$$\hat{\mathbf{o}}_b(t) = R_{\mathrm{tilt}}(t)^\top R_z[-\psi(t)]\, \hat{\mathbf{o}}_{b0}(t),$$

where $\psi(t)$ is the roll angle about the reference-frame z-direction aligned with the Sun, and $R_{\mathrm{tilt}}(t)$ is the composite tilt rotation.

For the two transient tilts, the rotation axis is defined in the xy-plane of this reference frame by an azimuth angle ϕ,

$$\hat{\mathbf{u}}(\phi) = \begin{bmatrix} \cos\phi \\ \sin\phi \\ 0 \end{bmatrix},$$

and the corresponding tilt angle is represented by a trapezoidal time profile,

$$\theta_j(t) = \theta_{\mathrm{max},j}\, \mathrm{trap}(t; t_j, t_{s,j}, t_{h,j}), \qquad j \in \{A, B\},$$

where t_j is the manoeuvre start time, $t_{s,j}$ is the ramp time, $t_{h,j}$ is the hold duration, and $\theta_{\mathrm{max},j}$ is the maximum tilt angle. Manoeuvre A and Manoeuvre B are therefore described by $(\phi_A, t_A, t_{sA}, t_{hA}, \theta_{\mathrm{max},A})$ and $(\phi_B, t_B, t_{sB}, t_{hB}, \theta_{\mathrm{max},B})$, respectively.

The roll motion is modelled as an initial roll angle together with a step-like roll change,

$$\psi(t) = \psi_0 + \Delta\psi\, \mathrm{ramp}(t; t_R, d_R),$$

where ψ_0 is the initial roll angle, $\Delta\psi$ is the magnitude of the roll step, t_R is the time at which the step begins, and d_R is the duration over which that step is applied.

A smaller bias tilt is also included,

$$\theta_{\mathrm{bias}}(t) = \theta_{\mathrm{bias}}\, \mathrm{ramp}(t; t_b, d_b),$$

about an axis

$$\hat{\mathbf{u}}(\phi_{\text{bias}}) = \begin{bmatrix} \cos\phi_{\text{bias}} \\ \sin\phi_{\text{bias}} \\ 0 \end{bmatrix},$$

where θ_{bias} is the bias tilt amplitude, ϕ_{bias} is its direction, t_b is its onset time, and d_b is its duration.

At each time step, five photometric features are computed:

$$\mathbf{x}_i = \begin{bmatrix} D_i^{\text{arr}} & S_i^{\text{arr}} & D_i^{\text{bus}} & S_i^{\text{bus}} & 1 \end{bmatrix},$$

where the terms correspond to diffuse and specular contributions from the array and bus together with a constant background term. The diffuse terms are Lambert-like,

$$D_i = \sum_k A_k \, \max(0, \hat{\mathbf{n}}_k \cdot \hat{\mathbf{s}}_{b,i}) \, \max(0, \hat{\mathbf{n}}_k \cdot \hat{\mathbf{o}}_{b,i}),$$

and the specular terms are computed using a Cook-Torrance microfacet model (see Appendix A) with a GGX normal distribution, Schlick Fresnel approximation, and Smith masking-shadowing. Stacking these feature vectors yields the design matrix

$$X \in \mathbb{R}^{T\times 5}.$$

Given a candidate attitude history, the brightness is modelled as

$$y'(t) = c_{D,\text{arr}} D_{\text{arr}}(t) + c_{S,\text{arr}} S_{\text{arr}}(t) + c_{D,\text{bus}} D_{\text{bus}}(t) + c_{S,\text{bus}} S_{\text{bus}}(t) + c_0,$$

where the coefficients are constrained to be non-negative. These coefficients are estimated by solving a weighted non-negative least-squares problem,

$$\min_{c\geqslant 0} \left\| W^{1/2}(Xc - \mathbf{y}) \right\|^2.$$

The non-linear attitude parameters are scored using the weighted sum of squared errors in normalised flux,

$$\text{SSE}_{\text{ratio}} = \sum_{i=1}^{T} w_i \left(\frac{y_i}{F_{\text{pre}}} - \frac{y'_i}{F_{\text{pre}}} \right)^2.$$

The weighting scheme emphasises the dip and the early recovery (i.e., the point at which the flux ratio begins to climb). The attitude parameters are estimated by randomized search. A coarse set of candidates is first sampled around a seed defined by the dip start and end times. The best candidate from the coarse stage is then refined by repeated perturbation at progressively smaller scales. For each candidate, the attitude series is generated, the feature matrix is computed, the non-negative photometric coefficients are estimated, and the weighted objective is evaluated. The candidate with minimum $\text{SSE}_{\text{ratio}}$ is retained as the final solution.

Section 4.4.4.2 summarises the fitted parameters for the two selected events, while Section 4.4.4.2 defines the fitted parameters and their roles within the attitude and photometric model. The fitted manoeuvre models against the observed light curves are shown in Figure 4.12.

Figure 4.13 displays the evolution of the attitude angles (yaw, roll, pitch and bias) over the two manoeuvres for the best-fitting models in both cases. For the first event, shown in the upper panel,

the fitted model provides a close representation of the observed brightness evolution. In particular, it reproduces the transition from the pre-dip level at points A and B into the principal decline, the low-flux plateau around C, the deeper minimum near D, and the subsequent rapid recovery through E to the post-step level at F. The agreement between model and observation is consistent with the comparatively small weighted error reported in Section 4.4.4.2, where $\mathrm{SSE}_{\mathrm{ratio}} = 0.686$.

For the second event, shown in the lower panel, the fitted model captures the broad structure of the dip, including the onset near B, the depressed flux interval spanning C to D, and the subsequent recovery towards E and F. However, the agreement is noticeably poorer than for the first event. Most notably, the pronounced pre-dip brightening immediately before the drop is not reproduced by the model, and larger residual differences remain during the recovery and post-step phases. This behaviour is reflected in the larger error value reported in Section 4.4.4.2, where $\mathrm{SSE}_{\mathrm{ratio}} = 6.017$. Taken together, these results indicate that the adopted manoeuvre parameterisation is sufficient to describe the principal features of the first dip, but is less able to account for the full photometric complexity of the second event (particularly the large peak before the brightness decrease), which may contain additional attitude dynamics or reflectance effects not represented in the present low-fidelity model. It should be emphasised that these fits are not unique: as discussed in §2.5.6, different combinations of attitude history and surface reflectance can reproduce similar photometric signatures, so the fitted parameters represent a plausible manoeuvre consistent with the observations rather than a definitive reconstruction, and alternative or additional manoeuvre components cannot be excluded. Referring back to §2.2.1, this section demonstrates the capability of high-precision photometry to not only be able to resolve manoeuvres but, to be able to fit for a plausible solution for said manoeuvres.

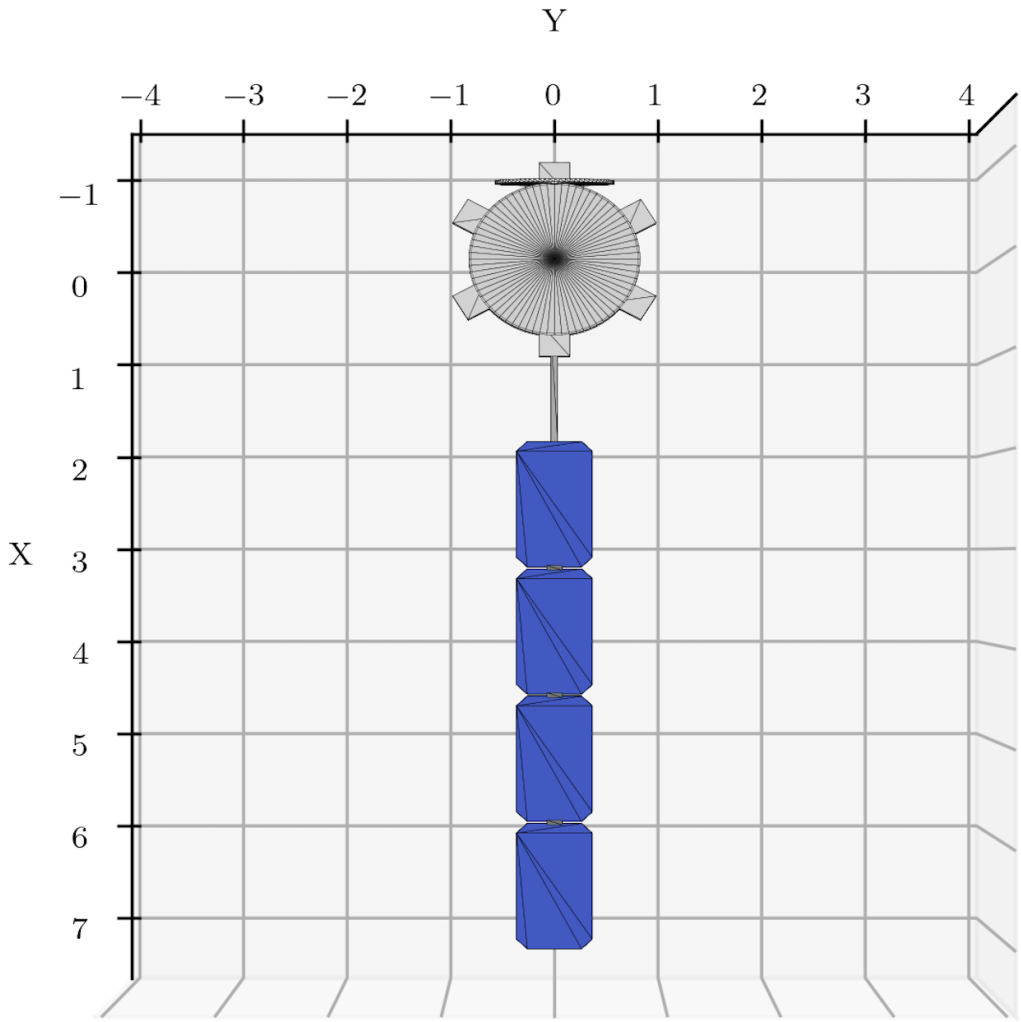

Figure 4.11: Top-down view of the simple 3D model of the LDPE-3A satellite which is used in the BRDF model to best reconstruct manoeuvres that are seen in the light curves.

Parameter	Manoeuvre 1	Manoeuvre 2
Dip interval (UTC)	2023-04-16 03:00:50 to 03:05:30	2023-04-16 05:30:40 to 05:34:00
Initial roll angle	-205.9°	-193.3°
Manoeuvre A direction, ϕ_A	49.6°	28.1°
Manoeuvre A maximum angle, $\theta_{max,A}$	50.5°	62.7°
Manoeuvre A start time, t_A	562.1 s	566.8 s
Manoeuvre A ramp time, t_{sA}	87.8 s	68.5 s
Manoeuvre A hold duration, t_{hA}	174.0 s	166.0 s
Manoeuvre B direction, ϕ_B	43.3°	166.4°
Manoeuvre B maximum angle, $\theta_{max,B}$	18.4°	23.3°
Manoeuvre B start time, t_B	751.6 s	701.3 s
Manoeuvre B ramp time, t_{sB}	21.0 s	16.2 s
Manoeuvre B hold duration, t_{hB}	98.2 s	142.3 s
Roll step, $\Delta\psi$	-54.2°	$+54.9^\circ$
Roll-step time, t_R	874.6 s	806.4 s
Roll-step duration, d_R	39.9 s	13.1 s
Bias tilt angle, θ_{bias}	-7.2°	$+7.1^\circ$
Bias tilt direction, ϕ_{bias}	27.4°	36.1°
Bias start time, t_b	839.5 s	828.2 s
Bias duration, d_b	49.2 s	29.6 s
Fit error, SSE_{ratio}	0.686	6.017

Table 4.4: Summary of fitted manoeuvre parameters for the two LDPE-3A events. Angles are given in degrees, and times are in seconds relative to the start of the fitted segment.

Term	Meaning	Role
Initial roll angle, ψ_0	Starting rotation about the reference-frame z-direction aligned with the Sun	Baseline roll angle in $\psi(t)$; sets the initial orientation before transient manoeuvres are applied.
Manoeuvre A direction, ϕ_A	Direction of the first tilt axis in the reference xy-plane	Defines the tilt axis $\hat{\mathbf{u}}(\phi_A) = [\cos\phi_A, \sin\phi_A, 0]^\top$.
Manoeuvre A maximum angle, $\theta_{max,A}$	Largest amplitude reached by Manoeuvre A	Sets the peak of the first trapezoidal tilt profile $\theta_A(t)$.
Manoeuvre A start time, t_A	Time at which Manoeuvre A begins	Marks the onset of the first trapezoidal manoeuvre.
Manoeuvre A ramp time, t_{sA}	Time taken for Manoeuvre A to rise to full amplitude	Controls the rising edge of $\theta_A(t)$.
Manoeuvre A hold duration, t_{hA}	Time for which Manoeuvre A remains near its maximum	Controls the plateau of $\theta_A(t)$.
Manoeuvre B direction, ϕ_B	Direction of the second tilt axis	Defines the axis $\hat{\mathbf{u}}(\phi_B)$ for the second transient tilt.
Manoeuvre B maximum angle, $\theta_{max,B}$	Largest amplitude of Manoeuvre B	Sets the peak of the second trapezoidal tilt profile $\theta_B(t)$.
Manoeuvre B start time, t_B	Time at which Manoeuvre B begins	Marks the onset of the second trapezoidal manoeuvre.
Manoeuvre B ramp time, t_{sB}	Time taken for Manoeuvre B to reach full amplitude	Controls the rise of $\theta_B(t)$.
Manoeuvre B hold duration, t_{hB}	Time for which Manoeuvre B remains near its maximum	Controls the plateau of $\theta_B(t)$.
Roll step, $\Delta\psi$	Sudden additional change in roll angle	Appears in $\psi(t) = \psi_0 + \Delta\psi\,\mathrm{ramp}(t; t_R, d_R)$. Positive and negative values indicate opposite roll directions.
Roll-step time, t_R	Time at which the roll step starts	Sets the onset of the smooth roll change.
Roll-step duration, d_R	Time over which the roll step is applied	Controls how abruptly or gradually the roll step is introduced.
Bias tilt angle, θ_{bias}	Small additional tilt offset	Amplitude of a secondary bias rotation added to the main manoeuvre.
Bias tilt direction, ϕ_{bias}	Direction of the bias tilt axis	Defines the axis $\hat{\mathbf{u}}(\phi_{bias})$.
Bias start time, t_b	Time at which the bias tilt begins	Sets the onset of the bias tilt term.
Bias duration, d_b	Time over which the bias tilt is applied	Controls the smooth growth of the bias tilt.
Fit error, SSE_{ratio}	Overall mismatch between model and normalised observations	Computed as $\sum_i w_i \left(y_i/F_{pre} - \hat{y}_i/F_{pre}\right)^2$; smaller values indicate a better fit.

Table 4.5: Interpretation of the fitted manoeuvre parameters used in Section 4.4.4.2.

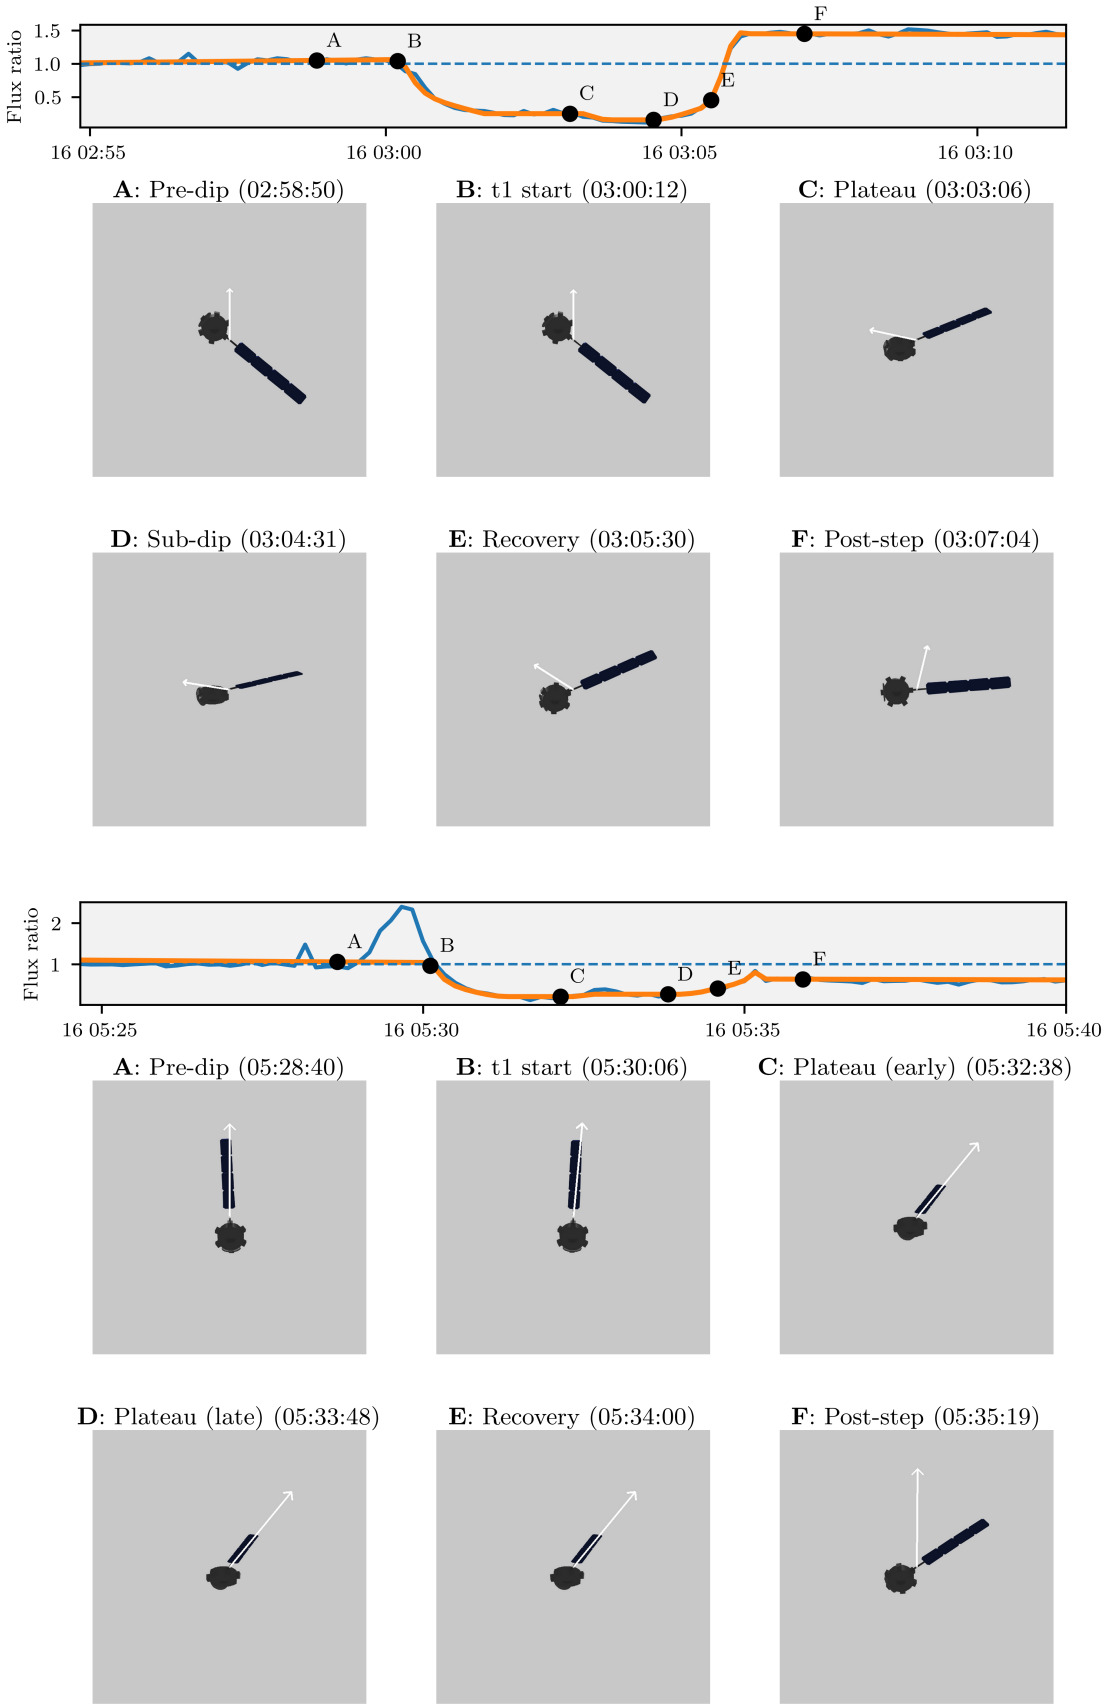

FIGURE 4.12: A demonstration of fitting the simulated light curve to the observed 'dips' caused by the spacecraft manoeuvring. The light curve is simulated using a set of specular Cook-Torrance BRDF terms and diffuse Lambertian terms and a simple 3D model representing the LDPE-3A satellite. The lower panels below the light curves display the object in key frames across the best fit manoeuvre.

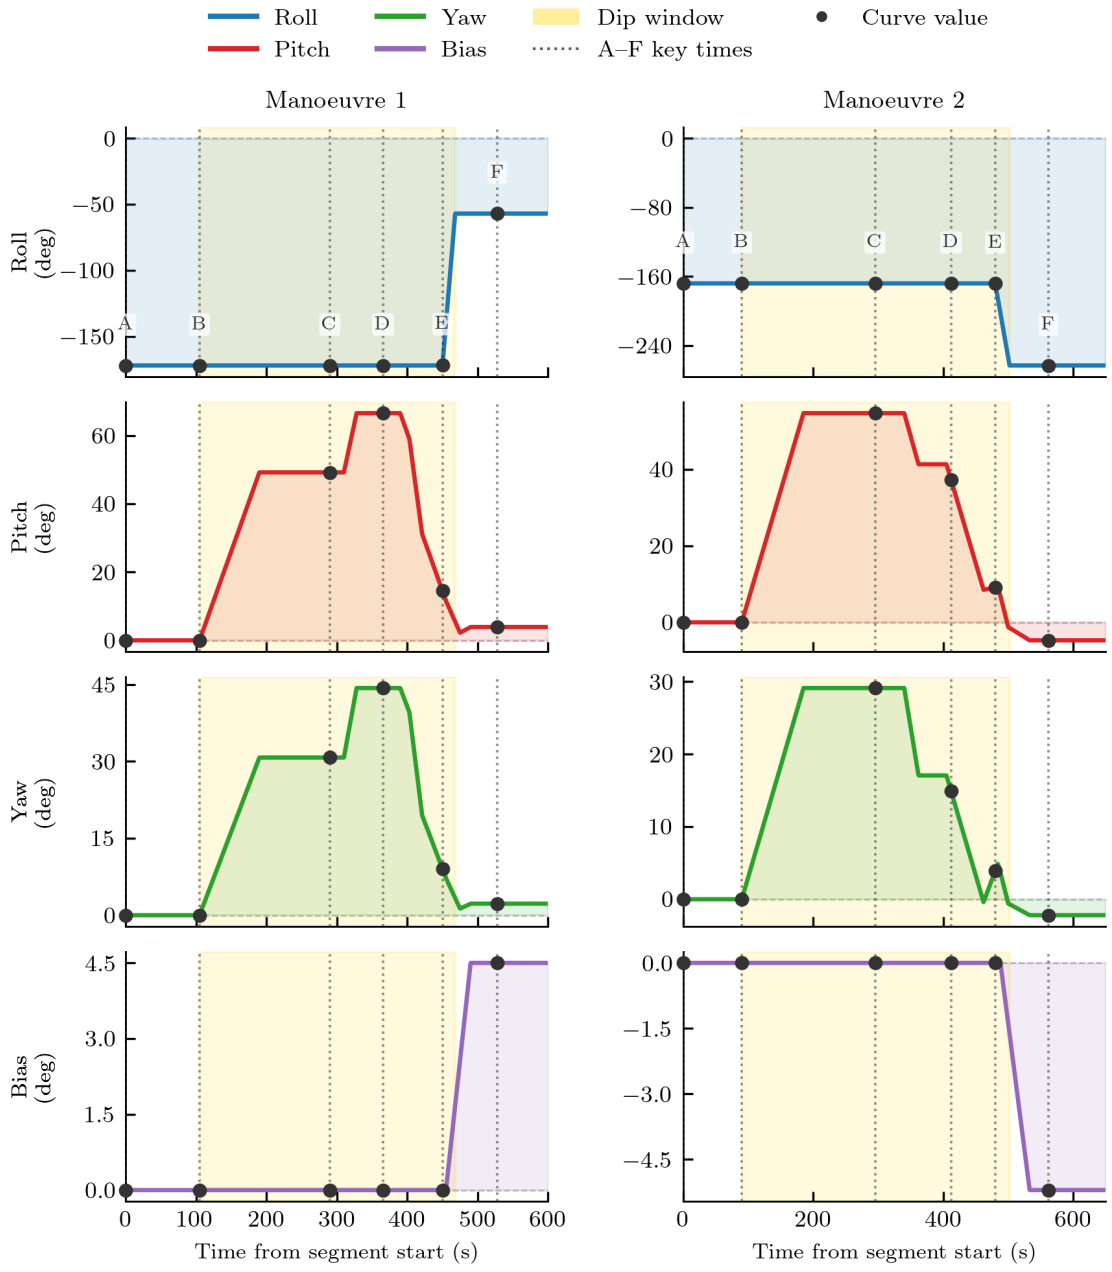

Figure 4.13: Reconstructed attitude angles for two LDPE-3A manoeuvres derived from the fitted model. Each column shows, from top to bottom, the roll, pitch, yaw, and bias angles (all in degrees) as a function of time from the start of the manoeuvre segment. Pitch and yaw are the cross-track and along-track projections of the net tilt, respectively; bias is the isolated contribution of the $\Delta\theta_{\text{bias}}$ ramp-step component. The gold shading spans the manoeuvre window (manoeuvre onset t_1 to roll-step completion $t_R + \Delta t_R$). Dotted vertical lines and filled circles mark six key epochs shown in Figure 4.12.

4.4.4.3 Blue Colour Features

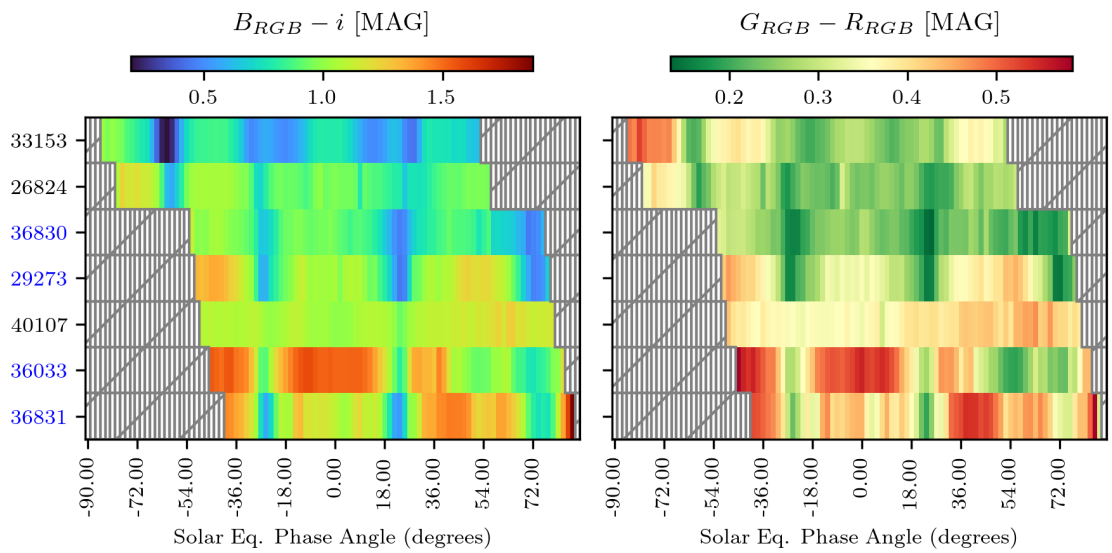

Figure 4.14: Colour light curves transformed into 2D image data for $B_{RGB} - i$ and $G_{RGB} - R_{RGB}$ colour indices with respect to the solar equatorial phase angle. Each row represents the colour light curve of a satellite object identified by its given NORAD ID. Blue ID labels indicate Spacebus-4000 satellites; black ID labels denote SSL-1300 satellites. Each colour light curve is split into 100 bins of equal phase width. All objects exhibit common structures within their colour light curves, characterised by three distinct peaks, including two blue features occurring at $\pm 25\,^{\circ}$. The objects are ordered by increasing longitude relative to La Palma.

Night Observed	NORAD ID	c_1 (M)	δ_1 (M)	c_2 (M)	δ_2 (M)
2023-05-13	36033	-0.64	0.03	-0.59	0.03
2023-05-13	36830	-0.3	0.02	-0.4	0.01
2023-05-18	36831	-0.66	0.02	-0.74	0.03
2023-05-13	29273	-0.57	0.04	-0.62	0.03
2023-05-13	40107	-0.06	0.01	-0.14	0.02
2023-04-23	33153	-0.21	0.01	-0.19	0.02
2023-04-25	26824	-0.28	0.04	-0.14	0.01

Table 4.6: Measurements and standard errors for the near-symmetrical colour changes measured in units of magnitudes (M). These features occurred for 7 satellites, each denoted here by their allocated NORAD ID. Blue ID labels indicate Spacebus-4000 satellites; black ID labels denote SSL-1300 satellites. The notation of c refers to the measurements of the colour index change of $B_{RGB} - i$ and the respective errors on the first and second colour change are denoted by $\delta_{1,2}$.

The appearance of blue colour features occurring for 7 satellites near symmetrically about $0\,^{\circ}$ of solar equatorial phase angle are reported. Four of these satellites belong to the Spacebus-4000 class and 3 belong to SSL-1300 class.

As shown in Figure 4.14, there are noticeable changes in the colour index as the night progressed for all 7 objects. In particular, the colour excursions are more visible in the $B_{RGB} - i$ colour space

as opposed to the $G_{RGB} - R_{RGB}$ colour space where the excursions are not as large. There is often a strong colour excursion either at the start of the night or at the end of the night, this feature may relate to the findings of Kerr et al. (2021) in which they see a strong brightness increase late in the observation window. They believed that this was caused by reflections off the side panel of the satellite which are very reflective.

Two of these changes of interest occur near-symmetrically at $\approx \pm 25^\circ$ solar equatorial phase angle. These colour excursions are occurring at approximately the same equatorial phase angle for each object and there is a shift towards 'bluer' colours. To constrain these features further, fitted measurements and standard errors on these 14 $B_{RGB} - i$ colour excursions were calculated.

To perform this calculation, fits were made across three sub-regions of the colour excursion of interest. This fitting consisted of two linear fits either side of the feature as a proxy for the baseline and a 4th order polynomial fitted across the colour change itself. The value for the colour change is approximated by taking the averages of the midpoint values of the two linear fits away from the minima of the polynomial fit, this was the most appropriate way of approximating an amplitude measurement from the baseline. With this in mind, it was deemed appropriate to apply a bootstrapping technique (Ader et al., 2008). The data are re-sampled with replacement and the fitting process and colour excursions are measured with the methods described above for over 3000 iterations. The mean and standard deviation of this distribution of measured colour changes can be found in Table 4.6. Whilst indeed it is very important to note that the sample size is limited, it does appear that on average the colour excursions seen with the Spacebus-4000 class are larger than the measurements for the SSL-1300 satellites. The absolute difference between the two excursions is minimal for most satellites with Intelsat 901 (NORAD: 26824) having the largest difference between the two colour changes. However, these results should be taken with caution given the small sample size.

Overall, this relates to findings of Dunsmore et al. (2017) where they find that a considerable portion of the solar spectrum greater than 600 nm is not reflected and from that they conclude that the surfaces responsible for these glints are reflecting primarily at shorter wavelengths and absorbing photons at longer wavelengths, which is typical of older satellite solar panels. Although they are working with slitless spectroscopy and this data in optical photometry, the same conclusions can be reached with the findings above, where the near-symmetrical colour changes are strongly reflecting at blue wavelengths.

4.4.4.4 Triple Peak Features

There were some shared features that were found within certain light curves. Most notably the appearance of triple peaks, which appears to be from dual offset solar panels plus an additional specular reflection from a component. Figure 4.15 displays clear examples of these features.

Interestingly, these satellites share a common feature found in their designs and that is a 'tab' component mounted upon the solar panel: the Eurostar attitude-control (ADCS) flap[2], a small solar-sail surface at the array used to manage solar-radiation-pressure torques, similar to the artistic impression of BADR-4 in Figure 4.16. The additional reflection may be caused by this. In terms of

[2] https://resilience.esa.int/archives/projects/e3000-second-generation-solar-array

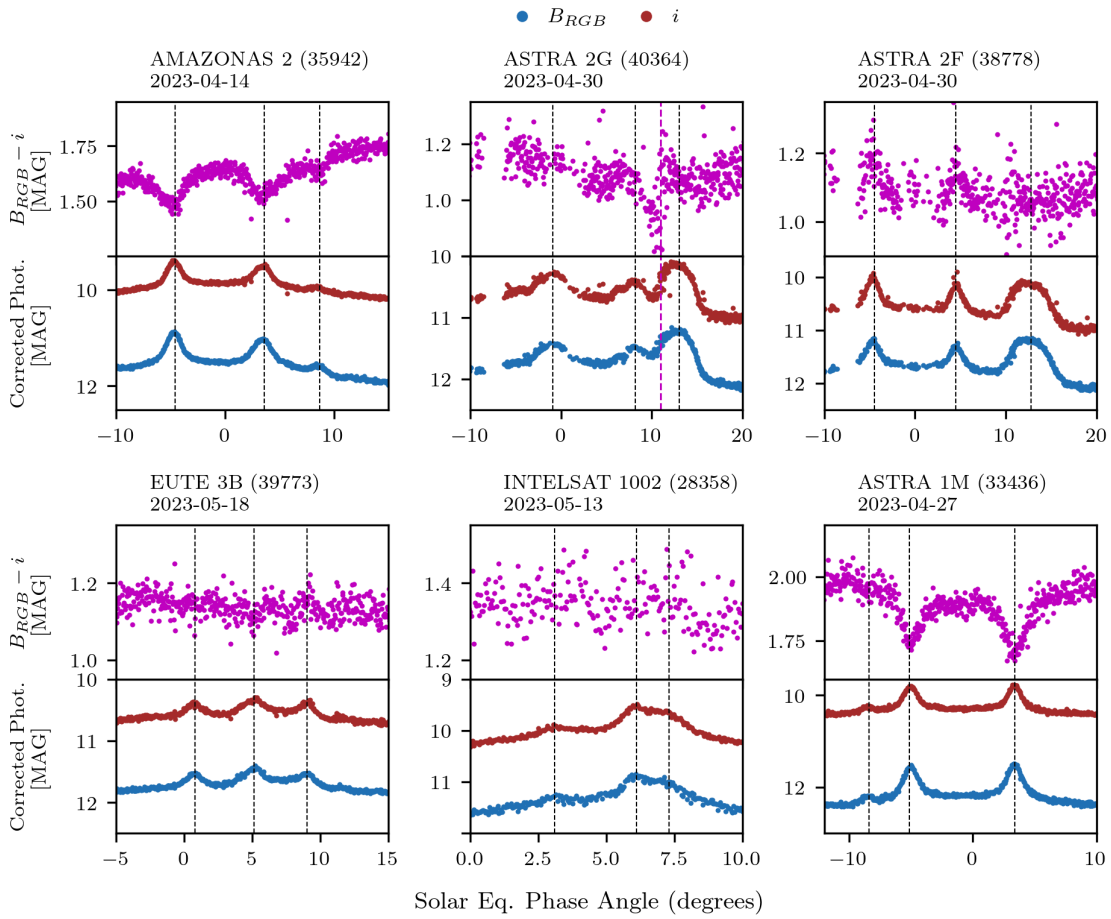

FIGURE 4.15: Triple peak features found in the light curves of 6 example satellites in the survey sample. The $B_{RGB} - i$ measurements are shown above the B_{RGB} band measurements for each satellite. All satellites belong to the Eurostar-3000 bus configuration class. The larger plateau feature in the INTELSAT 10-02 single band light curve is likely due to the addition of the MEV-2 (Mission Extension Vehicle) solar panels. Black dashed lines are given at each of the glint locations and the magenta dashed line is given where there is a significant 'bluer' feature for ASTRA 2G.

the colour response in $B_{RGB} - i$, they tend to vary across the satellites; AMAZONAS 2 shows a shift towards 'bluer' colours across all three peaks, ASTRA 1M does not show any colour change across the smaller additional peak but does across the others, EUTE 3B and Intelsat 10-02 do not show any colour changes and ASTRA 2G shows a stepped increase in the colour index across the 2nd and 3rd peaks. A unique colour feature is also seen within the light curve of the BADR-6 satellite. The bus configuration of BADR-6 is EUROSTAR-2000+. Figure 4.16 displays the artistic impression of BADR-4, which is also built upon the Eurostar-2000 design[3]. Figure 4.17 displays the portion of the BADR-6 light curve where the colour behaviour occurs. As the brightness is increasing, there is first a shift towards 'bluer' colours and then a shift towards redder colours at the end of this feature. The signature in the B_{RGB} and i bands is similar to the third peak found in the ASTRA 2F and ASTRA 2G light curves in terms of profile. The colour features are similar to that seen in the colour light curve of ASTRA 2G which is denoted by the dashed magenta line in Figure 4.15. It could be that the mounted solar panel component is responsible for this colour behaviour alongside

[3] https://space.skyrocket.de/doc_sdat/arabsat-4a.htm



Figure 4.16: Artistic impression of the BADR-4 satellite. It is noted that there is a 'tab' like component mounted upon the solar panel as notable feature. This feature is noted to be within the designs of many Eurostar-3000 satellites as well as Eurostar-2000.

favourable illumination geometry. Further observations would be able to display how this colour feature changes seasonally and constrain its geometry.

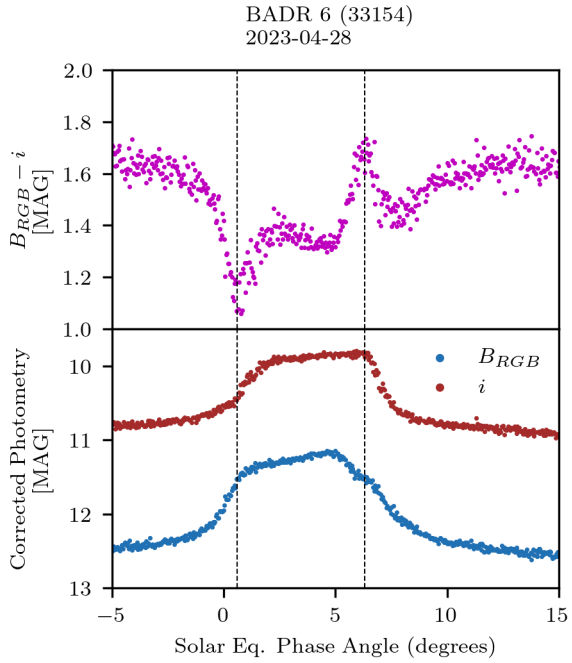

Figure 4.17: Portion of the BADR-6 light curve in the B_{RGB} and i bands. This portion contains both a very strong blue and then red colour response. Black dashed lines indicate the positions of the glints seen in the colour measurements.

4.4.4.5 Short Timescale Glints

Given the observations were carried out in the weeks during the 'glint season', it was expected that there would be features with variable brightness on short timescales due to the favourable

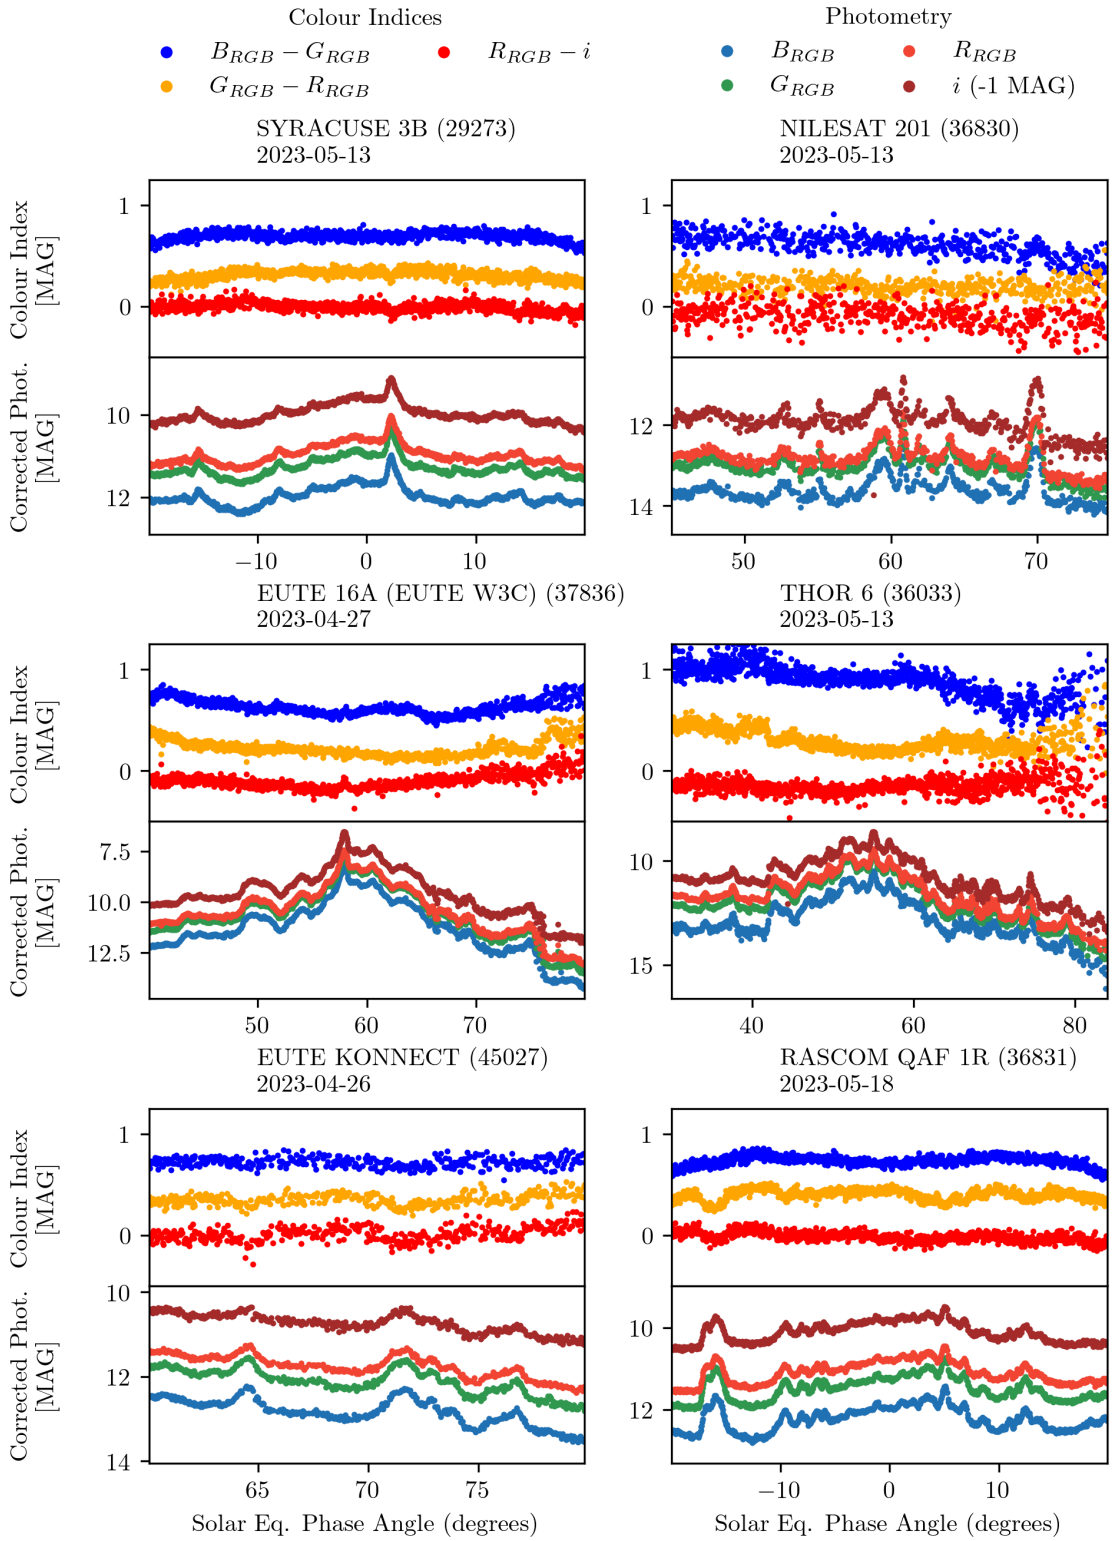

FIGURE 4.18: Examples of the glinting regions seen within light curves of satellites in the dataset. 1st row: corrected photometric measurements as a function of solar equatorial phase angle in the four colour bands. 2nd row: binned $B_{RGB} - i$ colour index as a function of solar equatorial phase angle. 3rd row: binned $G_{RGB} - R_{RGB}$ colour index as a function of solar equatorial phase angle. A 1 mag offset is applied to the i band measurements.

Sun-observer-satellite geometry.

Figure 4.18 displays 6 illustrative examples of the glinting seen within 16 satellites in the sample and how these glinting features appear in colour space. The three colour indices ($B_{RGB} - G_{RGB}$, $G_{RGB} - R_{RGB}$ and $R_{RGB} - i$) were chosen to display the general colour variation seen across regions of short timescale glinting across the different wavelength bands. These 6 satellites were chosen as examples as they best highlight the uniqueness in profile. It is found that the amplitudes of these glints tend to only increase by a few tenths of a magnitude and each individual glint tends to last only a fraction of a degree of phase (< 4 minutes), whilst overall the glinting regions tend to last 10's of degrees of phase. In the $B_{RGB} - G_{RGB}$, $G_{RGB} - R_{RGB}$ and $R_{RGB} - i$ colour spaces, the colour indices tend to stay fairly consistent across the glinting regions with only faint colour responses. In particular, there is not necessarily a strong colour response associated with a glint or set of glints (e.g., THOR-6 (NORAD: 36033), RASCOM QAF-1R (NORAD: 36831), EUTE 16A (NORAD: 37836)). Over the four colour bands, structurally the glinting regions have unique profiles which is expected given the uniqueness of every satellite-Sun-observer geometry and potentially the individual pointing of the satellites themselves.

From the measurements displayed in Figure 4.19 and Table 4.7, the mean glinting time is ≈ 124 minutes, with the start and end of each region marked by eye as the first and last discernible glints of the sequence. Figure 4.20 displays the distributions of brightness in all four colour bands and the two colour indexes considering all objects that had regions of glinting within their light curves. In terms of brightness, they tend to reach around 10th to 12th magnitude on average but can reach up to 8th in the B_{RGB}, G_{RGB}, R_{RGB} and i bands. The distributions in colour index space have small deviations especially within the $G_{RGB} - R_{RGB}$ distribution ($\sigma = 0.11$) which informs us that generally these glints tend to be achromatic and do not generate a significant colour response which suggests that the materials of the mounted components which are responsible for these features have spectrally flat profiles. For RASCOM QAF-1R (NORAD: 36831) and SYRACUSE-3B (NORAD: 29273), which are two satellites which display the blue colour features, the glinting region is near symmetrical about 0 ° of solar equatorial phase angle and occurs between the two colour changes discussed in §4.4.4 which suggests that the glinting is caused by off-angle reflections off the same component or one located close by. The colour of the glints on average is 'bluer' (reflecting at shorter wavelengths) than before and after the glinting window.

4.4.5 Average Colour Measurements

Figure 4.21 displays the average colour measurements ($B_{RGB} - i$ and $G_{RGB} - R_{RGB}$) versus the dry masses of satellites and the time spent on orbit (denoted as age in the figure). This is highlighted for satellites belonging to the most populous bus configurations.

Firstly, there is no strong overall correlation between the colour and time spent on orbit, which is in agreement with the findings of Schmitt (2020). The Pearson r correlation coefficient was measured between colour ($B_{RGB} - i$, $G_{RGB} - R_{RGB}$) and the time spent on orbit as -0.15 and -0.21 respectively. These values suggest very weak negative correlations. The corresponding p-values are 0.13 for $B_{RGB} - i$ and 0.03 for $G_{RGB} - R_{RGB}$. These results suggest that while the correlation between $B_{RGB} - i$ and time on orbit is not statistically significant, the weak negative correlation

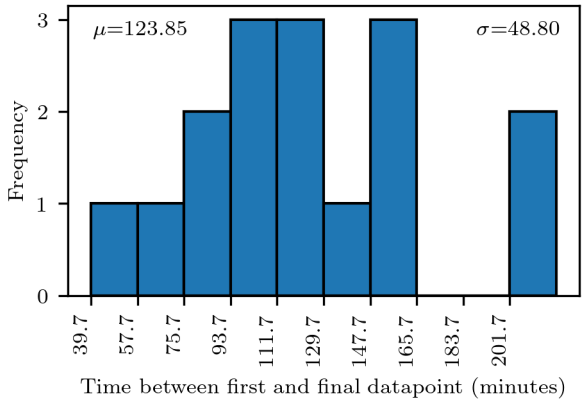

FIGURE 4.19: Distribution of the duration of glinting regions in minutes for all satellites that had glinting regions within their light curves. The mean and standard deviation of the distribution are given on the axes as μ and σ respectively.

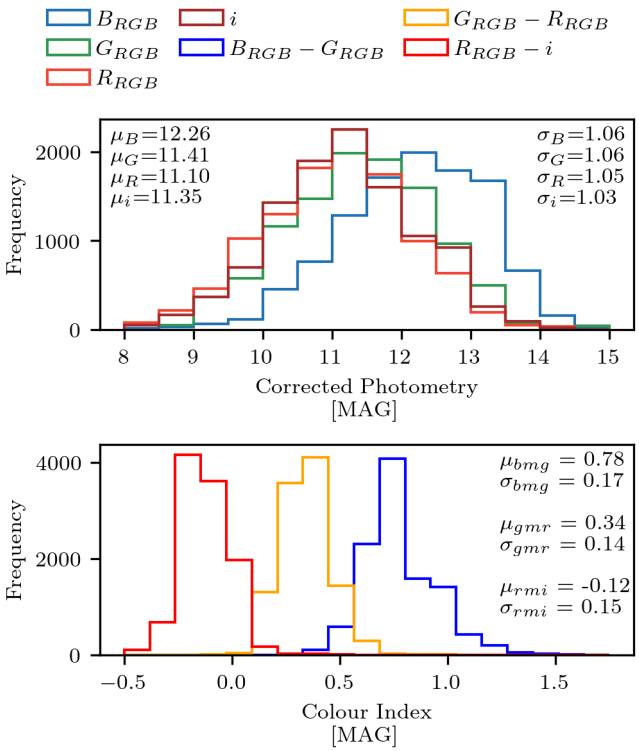

FIGURE 4.20: Photometric distributions of glinting regions within all four colour bands and the three colour index spaces ($B_{RGB} - G_{RGB}$, $G_{RGB} - R_{RGB}$ and $R_{RGB} - i$) respectively. The mean value and standard deviation are given for each distribution as μ and σ respectively.

NORAD ID	Start Time (UTC)	Duration (Minutes)
42934	2023-04-15 03:47:17	119.9
29273	2023-05-13 22:49:20	160.1
41866	2023-04-13 00:48:58	119.9
39168	2023-04-16 02:29:21	79.9
39172	2023-04-23 04:13:52	59.9
28945	2023-04-23 21:25:32	99.9
45027	2023-04-27 03:14:23	79.4
44457	2023-04-26 20:43:05	109.5
37836	2023-04-28 01:14:50	159.9
41029	2023-04-28 20:50:36	39.7
36592	2023-04-28 23:30:19	99.9
40258	2023-05-12 21:48:48	139.7
36830	2023-05-14 03:17:46	119.7
36033	2023-05-14 01:49:44	214.6
36831	2023-05-18 22:12:28	160.1
39163	2023-05-19 01:14:12	219.7

Table 4.7: Time measurements of the glinting regions for satellites. Displayed are the measurements of the starting time (UTC) and the duration (minutes) of the glinting regions. All glinting regions were carefully inspected to make sure these were the result of real reflections as opposed to bad zero-point measurements.

between $G_{RGB} - R_{RGB}$ and time on orbit does reach statistical significance ($p < 0.05$). This implies that although both relationships are weak, there is stronger evidence for a slight tendency of $G_{RGB} - R_{RGB}$ to decrease with increased time in orbit. Likewise, the Eurostar-3000 satellites follow a trend of increasing colour with time spent on orbit, this feature is also supported by the colour light curve maps of the Eurostar-3000 satellites (refer back to Figure 4.7 and Figure 4.8).

Furthermore, the STAR-2 satellites and to some extent BSS-702 satellites are isolated in colour. They have smaller and larger values of $G_{RGB} - R_{RGB}$ index respectively than satellites belonging to other bus configurations. With respect to the $B_{RGB} - i$ colour index, it is found that there are comparable colours to both the STAR-2 and BSS-702 satellites and so they no longer become isolated in colour. Again, this finding is supported by the colour light curve maps (refer back to Figure 4.7 and Figure 4.8). In terms of characterisation, it would be easier to discriminate between the STAR-2 and BSS-702 satellites.

It should be stated however, that given these correlations are calculated from a limited and unbalanced dataset and with uncertainties in dry mass data, there should be a notion to take these findings as preliminary requiring further evidence from repeated monitoring of the objects.

4.4.6 Colour Features in Relation with Physical Parameters

It was also of interest to identify whether the aforementioned blue colour features in $B_{RGB} - i$ correlate with certain physical properties such as the longitude relative to La Palma, the orbital inclination, the dry mass and the time spent on orbit. The measured Pearson r correlation coefficients

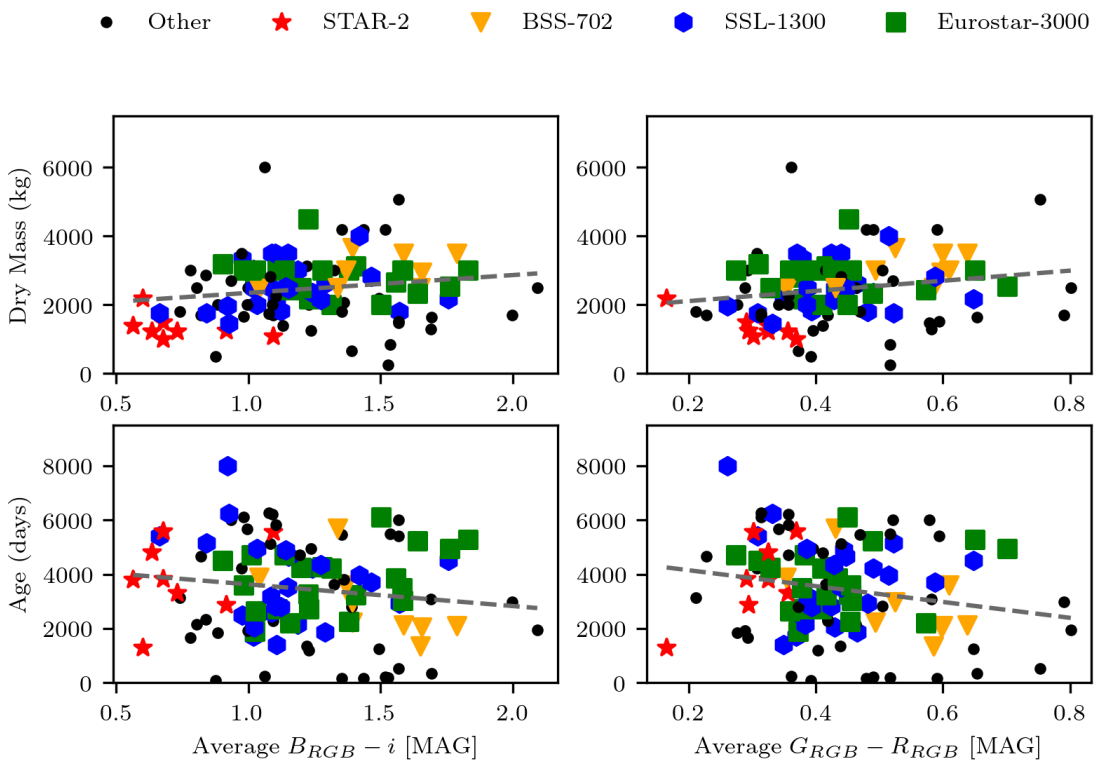

FIGURE 4.21: Time spent on orbit and dry mass for satellites with a variety of different bus configurations as a function of the average colour ($B_{RGB} - i$ and $G_{RGB} - R_{RGB}$) across their entire light curve. The grey dashed lines are linear fits applied across each population of datapoints.

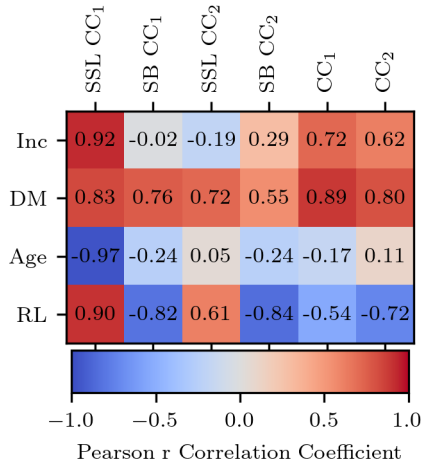

FIGURE 4.22: The measured Pearson r correlation coefficients arranged into a correlation matrix. Inc, DM and RL refer to the physical parameters of orbital inclination, dry mass and longitude relative to La Palma. SSL $CC_{1,2}$ and SB $CC_{1,2}$ refer to the 1st and 2nd colour change measurements for SSL-1300 (SSL) and Spacebus (SB) satellites respectively. $CC_{1,2}$ refers to considering the 1st and 2nd colour changes of all satellites.

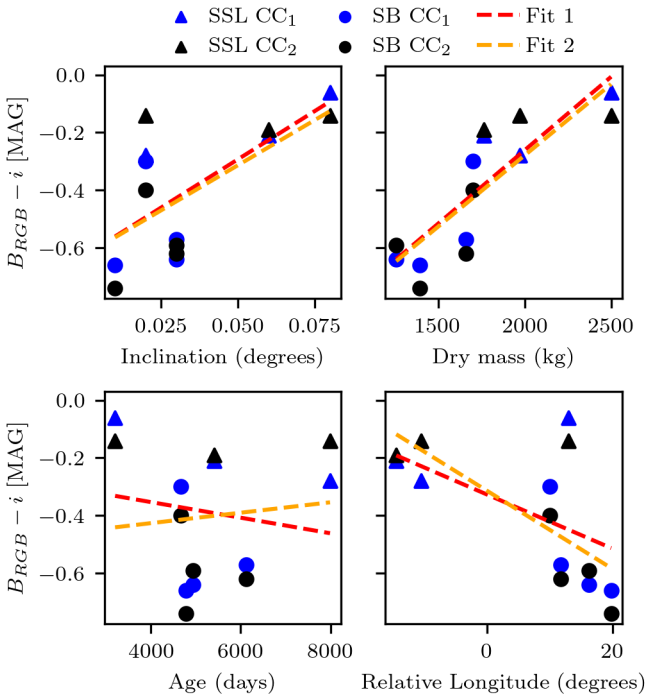

Figure 4.23: The measured $B_{RGB} - i$ values of the near-symmetrical colour changes as a function of different physical properties. The points in blue correspond to the first colour change (CC_1) and the points in black correspond to the second colour change (CC_2). The triangle and circle markers denote SSL-1300 (SSL) and Spacebus-4000 class satellites (SB) respectively. Linear fits are made between the 1st and 2nd colour change and the physical property, these are referred to by 'Fit 1' and 'Fit 2'. The errors on the colour measurements were significantly small so are not shown here.

being outlined within Figure 4.22, with the colour changes as functions of the physical parameters are displayed in Figure 4.23. The overall (not discriminating between SSL-1300 and Spacebus satellites) measured Pearson r correlation coefficients between each of the colour changes and orbital inclination are 0.72 and 0.62 respectively. These values suggest a strong positive correlation. However, when considering each of the bus classes separately, the correlation coefficients for the 1st and 2nd colour change were measured to be (0.92, -0.19) and (-0.02, 0.29) for the SSL-1300 and Spacebus satellites respectively. This contradicts the overall correlation that was measured and is contradictory between the same bus class. Similarly, there is a very strong positive correlation between each of the colour changes and dry mass, with measured Pearson r coefficients of 0.89 and 0.80 respectively. Considering the r correlation coefficients for the colour changes of each bus class with the dry mass, there is agreement that there is a strong positive correlation between them. There are strong negative correlations between the colour and longitude relative to La Palma with the measured Pearson coefficients being -0.54 and -0.72 respectively. This suggests that the further east the satellites are from La Palma, the 'bluer' the colour changes are. Considering each bus class, the Spacebus satellites follow the overall correlation and the SSL-1300 satellites follow the inverse, becoming 'bluer' the further west they are from La Palma. However, the SSL-1300 satellites are biased towards further western longitudes than La Palma.

However, there are very weak correlations between each of the colour changes and age with measured Pearson coefficients of -0.17 and 0.11 respectively. The correlation coefficients for the

Spacebus satellites and the second colour change of the SSL-1300 satellites are in agreement with this overall correlation. However, the first colour change of the SSL-1300 satellites are found to have a very strong negative correlation with age (r = -0.97). The fitting in Figure 4.23 between the colour changes and age is also poor, which further suggests that there is a lack of an overall relationship. There are only 14 datapoints (2 colour change measurements for each satellite) overall with fewer when considering either the Spacebus-4000 or SSL-1300 satellites individually, which is a small sample. Similarly, there is also a lack of datapoints across parameter space e.g., no datapoints between -5 ° and 5 ° of relative longitude. These per-class Pearson coefficients are computed on small subsamples and are sensitive to individual objects, so they should be regarded as preliminary until confirmed with larger samples. More measurements and longer term measurements would be needed to probe whether a relationship does exist between the size of these colour changes and each of the physical parameters outlined.

4.4.7 Solar Panel Offsets

Solar panel offsets can be determined directly by measuring the offset of specular glints. As argued in §2.6.2, a glint near 0° phase angle can only originate from a surface held near the Sun direction, which on a stabilised GEO satellite means the solar arrays; the m-shaped double signatures seen for 26 satellites, matching two arrays at slightly different offsets, support this identification. The offset of the specular glints are measured from 0° phase angle and take half of this value to yield the solar panel offset (Payne et al., 2006). Using this methodology, solar panel offsets for 54 satellites were measured, 26 of these satellites have both panels which are offset at slightly different angles from eachother. This presents itself as an m-shaped signature (see Figure 4.24) in the light curve. The distribution of these measurements are shown in Figure 4.25. It is found that most of these offsets are between -10 and 10 °, which is typically what is expected for geostationary satellites (Payne et al., 2006).

There are three outliers within this distribution. Firstly, satellite Telstar 14R (NORAD: 37602) which has an offset of approximately -41.5 °. It is understood from Rivera and Stewart (2021) that the northern solar array failed to fully deploy diminishing the amount of power available for transponders and thus reduced the life expectancy of the satellite.

Secondly, ALCOMSAT-1 (NORAD: 43039) has two solar panels offset at very large angles to what is normally expected, approximately -57.8 and -54.9 ° respectively. The satellite has spent almost six years on orbit so it is unlikely that this large offset is due to management of the solar arrays which is typically done in the early time spent in orbit. There is a possible link that these large offsets are implemented due to the bus on which it is configured. ALCOMSAT-1 is configured with the Chinese DFH-4 bus in which it provides high capability in output power and communication capacity. The approximate dry mass of ALCOMSAT-1 is 2500kg (McDowell, 2024). Considering these details, it is possible that the large offsets are to balance the impact of any torque generated by solar radiation pressure impacting upon the satellite. Finally, Echostar-15 (NORAD: 36792) which lies closer to the main distribution than the previous outliers, has an offset of about -15.7 °.

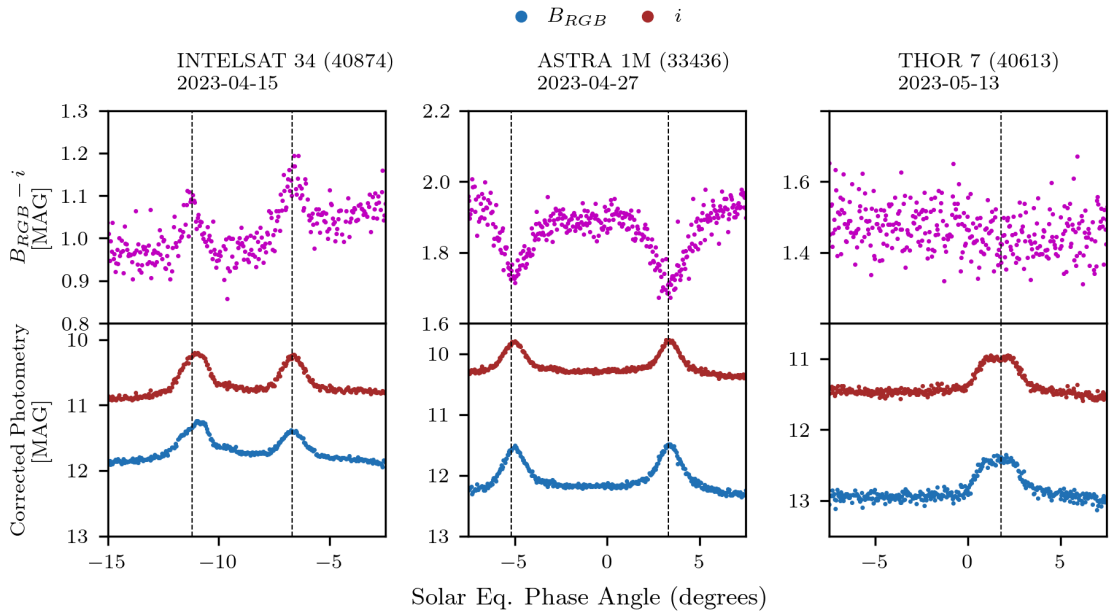

Figure 4.24: Three examples of the colour response across glints caused from specular reflections off of the solar panels. Upper: Colour measurements in the $B_{RGB} - i$ band. Lower: measurements made in the B_{RGB} (blue) and i (brown) bands. From left to right, there is a shift towards redder colours, towards 'bluer' colours and no significant colour change. Black dashed lines are used to highlight the locations of the solar panel glint(s) in both the single band and colour measurements.

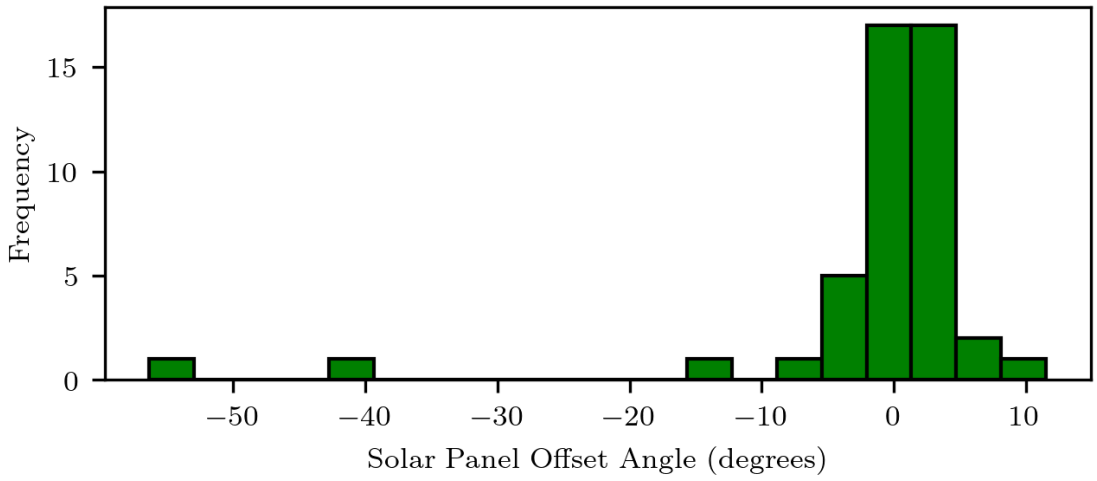

Figure 4.25: The distribution of solar panel offsets for 54 satellites in the sample. For satellites that have solar panels which are offset from each other, the average is taken between the two offsets. The three satellites that lie away from the main distribution are annotated by their NORAD IDs.

4.4.7.1 Solar Panel Glint Colours

It is highlighted in the works of Jolley et al. (2015) and Bedard and Wade (2014) that the illumination-observational geometry plays a significant role in the observed colour response from solar panel glints. The spectral energy density (SED) of reflected light off of a surface such as the solar panels is not invariant but varies and it is this variation of the overall geometry that leads to differences in the observed colour index of the solar panels. For example, at times around the summer and winter solstices, the angle at which the Sun's light illuminates a geostationary satellite solar array is the greatest. The specular reflection is thus directed away from Earth's equator; to the north during the northern hemisphere winter, and to the south during the northern hemisphere summer. For example, since the illumination angle is large on these nights close to the solstices and if the observer is more northern than where the specular reflection is directed, the extra contribution made by the solar panel glint to the measured flux, above the level measured without the specular reflection, consists of more longer wavelength light than shorter wavelength light. Thus, the overall colour shifts towards 'redder' colours.

There are three types of colour responses exhibited by the solar panel glints; stable behaviour, shifts toward redder colours and 'bluer' colours respectively. Table 4.8 shows the qualitative measurements of the solar panel glint colours in $B_{RGB} - i$. The majority of the colour responses are that there is no significant colour change across the solar panel glints.

Figure 4.24 highlights an example of each of these colour responses. There are no clear correlations of the type of colour response with the orbital slot longitude, age or bus configuration. The variation of the colour of the solar panel glints could possibly be due to the illumination angle at the time at which they were observed, seasonal changes in solar declination will result in variable solar array illumination angles. That is to say that if each of these solar panel glints was observed at different points in the year there would be variable colour responses. Further observations of these satellites would be needed to verify this.

Finally, 52 % of the solar panels that had dual offsets belonged to the Eurostar-3000 bus configuration class, which suggests this is a common operating mode.

Type	B	N	R	Total
Single	8	11	9	28
Dual	6	13	7	26
Total	14	24	16	54

Table 4.8: Colour responses of the measured solar panel offsets in the sample by the type of offset. B, N and R refer to the colour response of shifting towards 'bluer' colours, no observed colour response and shifting towards redder colours respectively. Single and Dual refer to where the solar panels were both offset at the same angle or at slightly different angles respectively.

4.5 Discussion

The colour light curve maps in this chapter serve as an extremely useful way to compare and contrast satellites belonging to different bus configurations. The key findings from these measurements are

as follows:

- Older Eurostar-3000 satellites are ‘redder’ than their successors, which contradicts the trend seen within the other bus classes
- STAR-2 satellites as a group are much ‘bluer’ than any other configuration and have colours comparable to the Sun
- BSS-702 satellites as a group are ‘redder’ than any other configuration

The colour evolution within the Eurostar-3000 bus configuration class could point towards manufacturing changes during the time when ASTRA-3B (NORAD: 36581) was being constructed. Further observations of more Eurostar-3000 satellites are needed to confirm whether this is a definitive trend. If it is, then this would be a powerful tool for identifying specific Eurostar-3000 satellites. It is also a possibility that the similarity in colour between the five red satellites is linked with their operator’s requirements at the time, for example three of these satellites are operated by EUTELSAT and are identical in design, these being Eute 33E (NORAD: 33750), Hotbird 13B (NORAD: 29270) and Hotbird 13C (NORAD: 33459). There is also evidence to support that this might be the result of manufacturing changes given the project documentation for which the Eurostar-3000 mechanical platform would be improved upon in European Space Agency (2014) and it is mentioned that solar panel manufacturing received two new design processes. The status update of this document was given in 2014, which would be before the majority of the satellites in this Eurostar-3000 sample were launched. There is not any evidence to suggest that the evolution of colour is related to either the satellites on-sky position or pointing given multiple satellites share similar communication footprints and have different colours. There is also no correlation with the solar declination angle.

Furthermore, STAR-2 satellites have a typical design with large flat antenna reflectors as shown in Figure 4.26, these are typically made from aluminium which as a material has a greater spectral reflectance value at shorter wavelengths (Quazi et al., 2015). This could explain why these satellites are much ‘bluer’ than any other configuration overall. It could also be a possibility that the STAR-2 satellites have less gold MLI materials either within their designs than the BSS-702 satellites or in the line of sight of STING.

With regards to the blue colour features seen within the SSL-1300 and Spacebus-4000B class satellites, these features could be caused by reflections of inclined components mounted upon the solar panels given their near-symmetrical nature and ‘blueness’ compared to the rest of the colour light curve. There is some near-symmetry with regards to the glinting features within RASCOM QAF-1R (NORAD: 36831) and SYRACUSE-3B (NORAD: 29273) which would further support this argument. The older SSL-1300 satellites show these near-symmetrical colour changes, which suggests that the newer satellites within this class have moved away from designs with the component which is causing this.

The short timescale glinting features found within this GEO-satellite sample are unique to each satellite and are likely a result of reflections off flat surfaces such as solar panels.

The glinting features tend to not generate a large colour response over an average duration of two hours. However, it is found that the average colour of the glinting regions is ‘bluer’ compared

to before and after these regions. If these unique features do appear seasonally, this would be a very powerful characterisation tool for identifying individual objects.



FIGURE 4.26: An artistic impression of the STAR-2 bus configuration design, the most notable feature of this design are the large flat antenna reflectors (Orbital ATK, 2017).

4.6 SUMMARY

Full night multi-colour light curves for over 100 GEO satellites were obtained using the STING instrument following observations over five weeks. This was the first ever large systematic survey of the GEO regime from which the analysis of full night light curves from the observations was presented. The error reduction between simultaneous and sequential observations was quantified.

The manoeuvres in the light curve of the LDPE-3A satellite were resolved and detected. Colour measurements were transformed into colour light curve maps which gave a relative comparison between satellites belonging to different bus configuration families. From this and further colour analysis, the STAR-2 and BSS-702 satellites are the most easily identified classes due to their colours and that there is tentative evidence showing that older Eurostar-3000 satellites are redder.

Regions of short timescale glints across multiple satellites were investigated and it was found that they are unique in profile to each satellite and imply possible reflections of inclined flat surfaces such as solar panels given their relative blueness with respect to the rest of the colour light curve, this uniqueness could serve as a very powerful characterisation tool for identifying individual satellites, especially if these profiles are seasonal. The changes in colour of the eastern and western illuminated faces of satellites were investigated and it was found the Spacebus-4000 satellites display the largest changes in colour collectively than any other class.

Finally, the solar panel offsets of 54 satellites in the sample were measured, and the colour response of their panel glints investigated. There is variation in the type of colour response, but the majority show no colour change across the glints, as opposed to shifting towards ‘redder’ or ‘bluer’ colours.

The findings from this chapter are drawn from a single five-week season; the correlations and class separations identified here are correspondingly preliminary, and repeated survey seasons are the natural route to confirming them.

5 Mapping the Spin-state of the Ajisai Satellite

"No, Frank! To try again, Frank... To try."

Matt Murdock (Daredevil) -
Daredevil TV Series, Episode 3:
"New York's Finest"

5.1 Motivations

The population of non-cooperative objects in the near-Earth environment has continued to grow in recent years (Blake, 2022; Serrano et al., 2024), such that active debris removal (ADR) techniques to achieve a more sustainable space and reduce collision risk are being proactively proposed and discussed (see e.g., Endo et al., 2022; Boonrath et al., 2024; Rossi et al., 2024; Sampath and Feng, 2024). In particular, most ADR techniques rely on a service vehicle to attach itself to the debris target; therefore, it is critical that the rotational state of the debris object be known in advance: the capture mechanism, approach trajectory and any de-tumbling strategy are designed around the target's spin-state well before rendezvous, with on-board sensing serving only to refine it during final approach (Shan et al., 2016).

In 2019, observations of the Experimental Geodetic Satellite (EGS), also known as Ajisai, were obtained using the SuperWASP (§2.2.2) instrument as part of a wider survey of low-Earth-orbit satellites (Chote et al., 2019). Previous work by (Shrive, 2024) used these observations as a case study for modelling Ajisai's specular glints.

Although this approach was sufficient enough such that it could reproduce the qualitative structure of the light curves, it had difficulties in obtaining consistent periods and was unable to quote uncertainties on spin-state parameters. The model remained partly empirical: it relied on digitised mirror coordinates, simplified glint profiles, and residual fitting to match the observed light curves, rather than predicting the reflected signal from the underlying optical and geometric properties of the satellite. The Ajisai problem was therefore revisited from first principles, comparing brightness in the body frame from spin-state parameter sets with the fixed modelled mirror positions. This new approach resolves the aforementioned issues and could be applied more generally to other targets.

As an overview, this chapter aimed to determine the attitude state of the Ajisai satellite purely from TLE-derived vectors and high-cadence optical photometry, expanding on the previous work by Hall et al. (2006) (Epoch-method) and, more recently, by Koshkin et al. (2024) (photometric-

patterns method). It is important to note that Ajisai's spin-state has been previously investigated using different approaches. Notably, Calatroni et al. (2025) employed high-frequency optical photometry (5–10 kHz) to determine the full rotation state, including the spin axis, period, and initial rotation phase, with high single-pass accuracy. Kucharski et al. (2016) achieved long-term monitoring of Ajisai's attitude evolution using 2 kHz Satellite Laser Ranging (SLR) data, benefiting from much broader temporal coverage while maintaining high-precision. Earlier optical studies such as Koshkin et al. (2017) demonstrated that attitude determination can also be performed with modest ground-based instrumentation, and therefore this work is particularly relevant as a precursor to the present study. The reflectivity of the Ajisai mirrors has also been studied in Kucharski et al. (2019b) and Kucharski et al. (2020), where the focus is purely on mirror reflectivity mapping rather than attitude determination.

Observational cadences on the order of milliseconds are needed to ascertain the spin-states of the fastest spinning objects (Kucharski et al., 2019a), which places strict limits on observing cadence with traditional CCD observations due to their relatively slow readout times and inefficiency at capturing rapid exposures without introducing significant dead time. Here a technique has been presented that does not rely on specialised kHz-level instrumentation, making it more readily scalable to a wider range of targets and observatories.

With the exception of §5.5, in which Ajisai was observed with the STING instrument in 2025, material in this chapter was presented in Airey et al. (2025b). Nikolay Koshkin of the University of Odessa kindly provided the modelled measurements of the Ajisai mirror positions in the body frame of the satellite.

5.2 Ajisai Background

The Ajisai satellite (NORAD: 16908, COSPAR: 1986-061A) was launched in 1986 as part of a geodetic mission sponsored by the NASDA (National Space Development Agency of Japan), with the goal of improving the accuracy of domestic geodetic triangulation with respect to other networks and accurately measuring the position of the remote Japanese islands (McDowell, 2002). The satellite is in Low Earth Orbit (LEO) approximately 1500 km above the Earth's surface with an orbit inclined by approximately 50°. Ajisai is a 685 kg hollow sphere with a diameter of 2.15 metres, and the surface is covered with 318 mirrors and 1436 corner-cube retroreflectors (CCRs) (Sasaki and Hashimoto, 1987) (see Figure 5.1). Work done in Kirchner et al. (2007) determined that Ajisai is spinning in the clockwise direction. The mirrors of Ajisai are curved at a radius of 8.55 ± 0.15 m and have a surface area of up to 393 cm^2 each (Kucharski et al., 2019b). There are two modes of operation that enable precise geodetic measurements: one involves sending a ground-based laser beam to the satellite's CCRs and measuring the round-trip travel time (laser ranging); the other involves imaging the illuminated satellite from ground stations against background stars to determine its position (eoPortal – Earth Observation Portal, 2012).

The mirrors covering the surface are organised into 12 primary latitudinal rings, along with two polar caps. Each primary ring is composed of several sub-rings, with each sub-ring containing three mirrors positioned at the same latitude and inclined in the same latitudinal direction. This

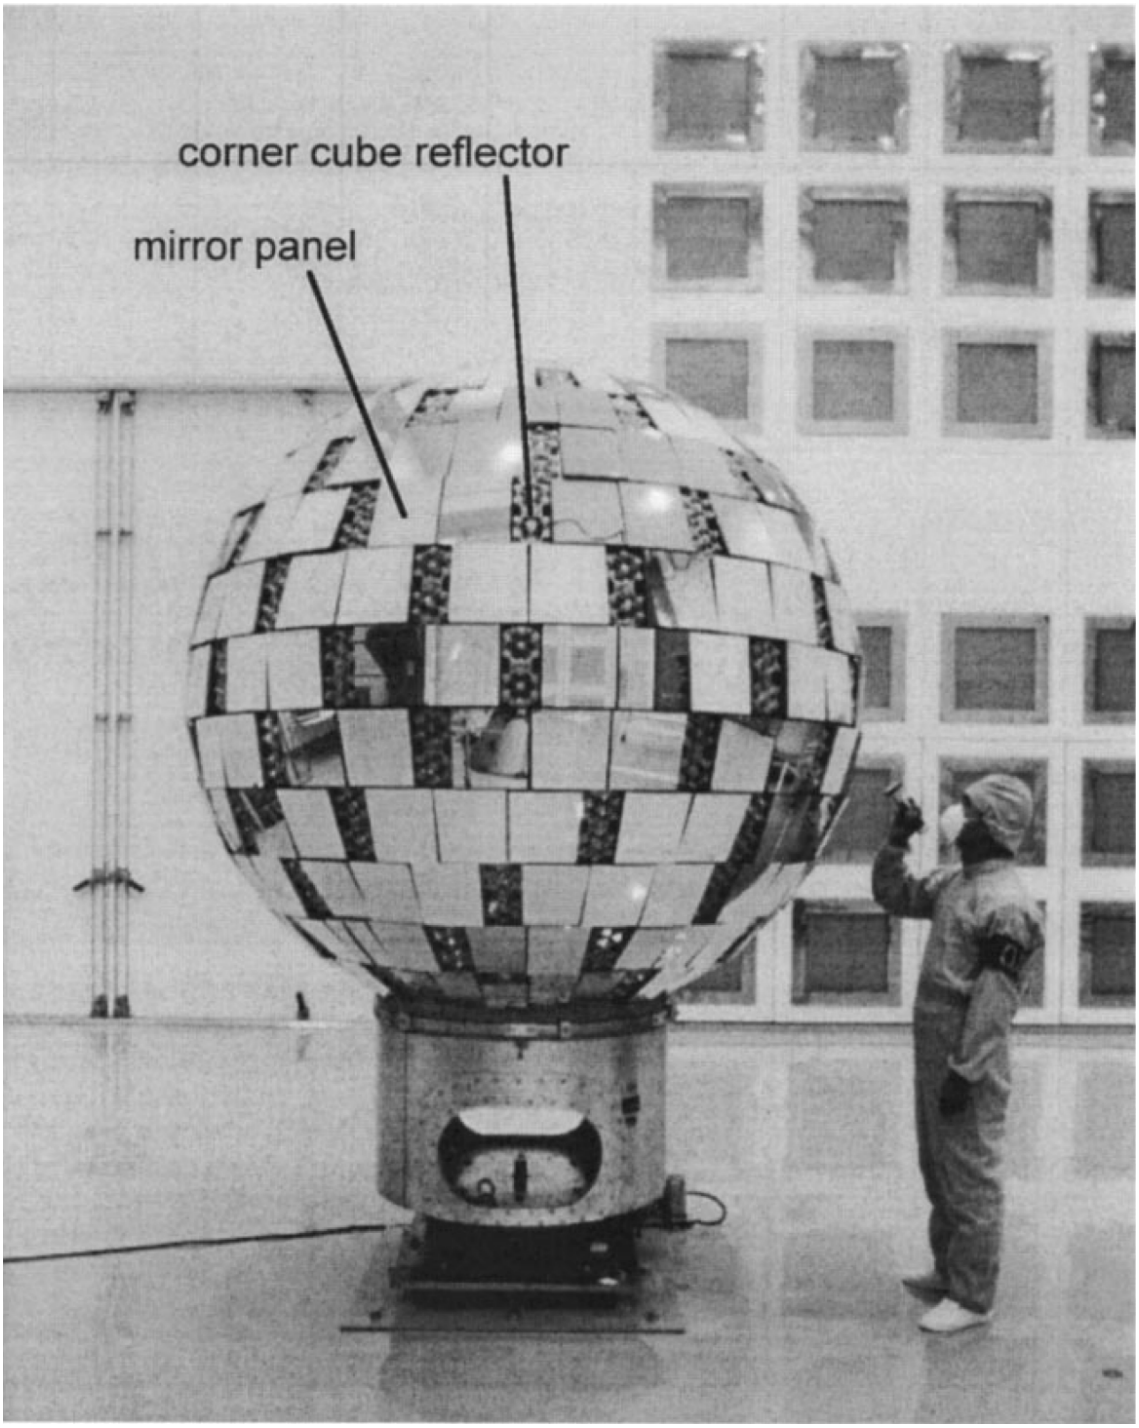

Figure 5.1: The Ajisai satellite with mirror panels and CCRs (corner cube reflectors) covering its spherical surface labelled on the figure. This figure is taken from Otsubo et al. (2000).

means that within a given revolution of the object around its spin-axis, a sub-ring of mirrors should always be able to reflect light towards an observer.

Ajisai was spin-stabilised, with an initial spin vector aligned with the Earth's rotation and a period of 1.5 s (Sasaki and Hashimoto, 1987). The spin-axis in the International Celestial Reference Frame (ICRF) and period of Ajisai at any time of observation can be approximated using the empirical models from (Kucharski et al., 2016). These are a set of time-dependent equations which describe the expected evolution of the precession and nutation cones of the spin-axis, and the period as an exponential trend function. The expected evolution of the spin-state of Ajisai in accordance with the empirical models is shown in Figure 5.2. The equations to calculate the spin-axis vector and period are given in Appendix C.

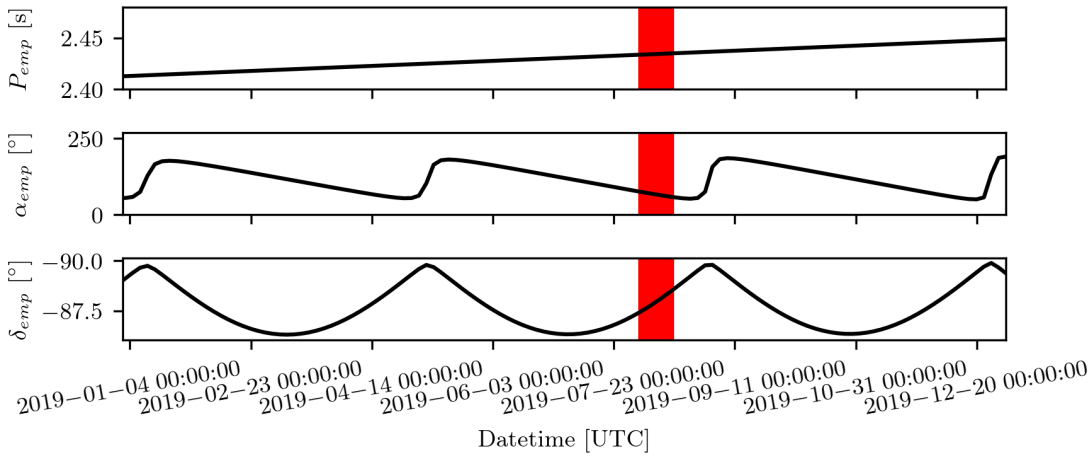

Figure 5.2: Propagation of the empirical model (Kucharski et al., 2016) for the orientation and period of Ajisai's spin axis. The red segment corresponds to the extent of the dates in which Ajisai was observed with SuperWASP. The empirical model defines δ_{emp} as being measured relative to the south celestial pole, with a clockwise spin relative to this axis.

5.3 Methodology

The procedure for extracting the spin-axis and period of Ajisai from the photometric observations can be broadly broken down into the following steps:

1. Defining a function that calculates (from an input Two Line Element ephemeris) the Phase Angle Bisector (PAB) vector in the ICRF reference frame as a function of input time.

2. Defining a function that transforms the PAB unit vectors, for a chosen set of rotation parameters (spin-pole coordinates, period and a rotational phase offset), into the RSO body-fixed frame.

3. Using these two functions to map each measurement in the light curve, for a chosen set of rotation parameters, to a latitude and longitude on the surface of Ajisai.

4. Comparing the resulting two-dimensional glint positions against the known mirror positions and using an MCMC (Markov Chain Monte Carlo) search to find the parameters (and associated uncertainties) that produce the best alignment.

The result of this procedure is a two-dimensional brightness map (discussed in §5.3.2) that visualises the surface reflectivity of Ajisai. The position of each mirror (strictly speaking, the orientation of its normal vector) is clearly visible as a localised spike in brightness centred on its position.

5.3.1 REFERENCE FRAME TRANSFORMATION AND VECTOR CALCULATIONS

As described in §2.1.1.3, the PAB, $\mathfrak{p}$, is defined as the unit vector halfway between the satellite-to-Sun ($V_{sat\to Sun}$) and satellite-to-observer directions ($V_{sat\to obs}$):

$$\hat{\mathfrak{p}} = \frac{\vec{V}_{sat\to obs}(t) + \vec{V}_{sat\to Sun}(t)}{\|\vec{V}_{sat\to obs}(t) + \vec{V}_{sat\to Sun}(t)\|} \tag{5.1}$$

It is a useful quantity in determining whether a reflection has occurred off a surface, as the specular condition holds that a reflection will occur when the PAB vector and the surface normal $\vec{n}$ to a flat facet nearly coincide (Hall and Kervin, 2013). The geometry of the PAB and angles related to the frame transform are illustrated in Figure 5.3. In order to compare the observations with the model, the PAB vectors must be transformed from the ICRF into the body-fixed frame. This transform depends on the spin-state of the satellite, such that finding the optimum spin parameters for this transformation will determine how well the surface reflectivity of the satellite in the body-fixed frame is mapped. The satellite-to-Sun and satellite-to-observer vectors are transformed from the ICRF frame to the body frame using a transformation matrix T, which follows standard right-hand rotation matrices around the z (R_3) and y (R_2) axes respectively (see e.g., Arfken and Mullin, 1967; Goldstein, 1980; Seeber, 2003):

$$R_3(\gamma) = \begin{bmatrix} \cos\gamma & \sin\gamma & 0 \\ -\sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{5.2}$$

$$R_2(\delta) = \begin{bmatrix} \cos(\frac{\pi}{2} - \delta) & 0 & \sin(\frac{\pi}{2} - \delta) \\ 0 & 1 & 0 \\ \sin(\frac{\pi}{2} - \delta) & 0 & \cos(\frac{\pi}{2} - \delta) \end{bmatrix} \tag{5.3}$$

$$R_3(\alpha) = \begin{bmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{5.4}$$

$$T = R_3(\gamma) \cdot R_2(\delta) \cdot R_3(\alpha) \tag{5.5}$$

where γ is the spin-angle, and α, δ are the right ascension and declination of the spin-pole in the ICRF respectively. The first rotation $R_3(\alpha)$ aligns the x-axis with the spin-pole's projection onto the equatorial plane. The second rotation tilts the new coordinate system such that the new z-axis points towards the spin-pole. Finally, the third rotation sets the orientation of the satellite about its

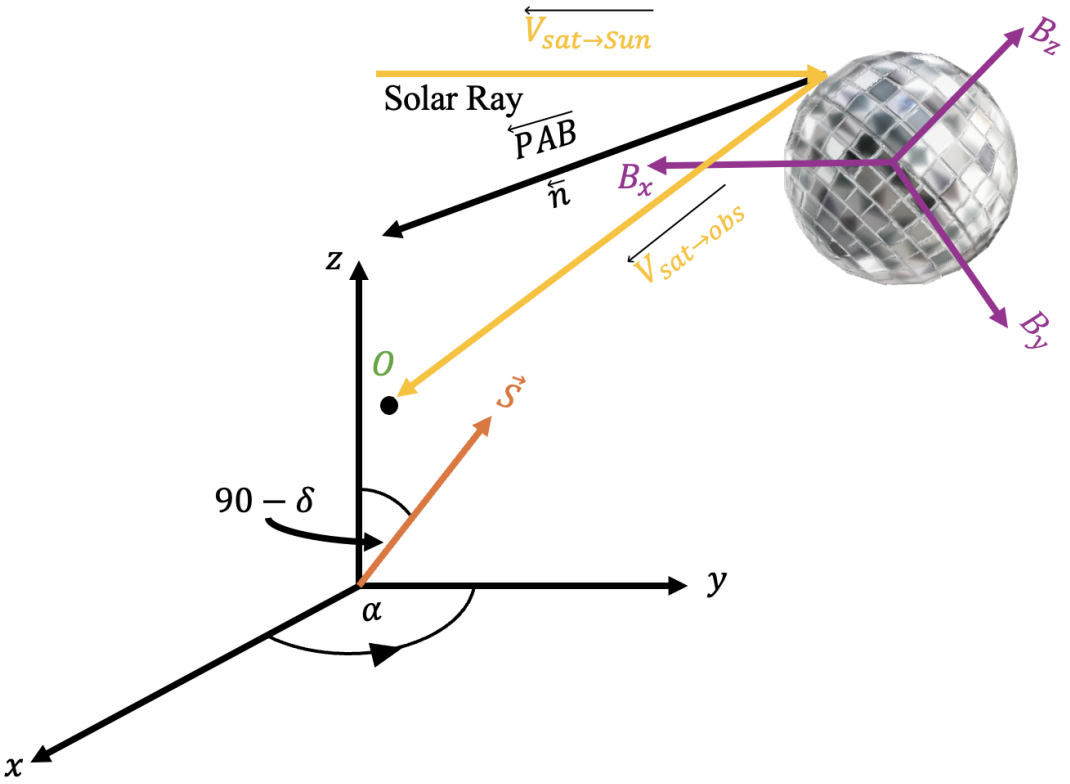

Figure 5.3: Diagram illustrating the defined angles, α and δ of the spin-axis vector, $\vec{S}$, with respect to the axes of the ICRF. The PAB is also defined with respect to the vector directions toward the Sun and observer O from the satellite with body frame axes B.

spin-axis at a given time-step, essentially defining the rotation around the newly aligned z-axis. Figure 5.4 illustrates the sequence of rotations from ICRF to body frame.

The satellite-Sun, satellite-observer and consequently the PAB vectors were calculated (see Equation 5.1) in Python using `Skyfield` (Rhodes, 2019), which effectively propagates a TLE into an orbital position from an observer location with respect to the desired reference frame, in this case the ICRF. The vectors were calculated from the TLEs at a 2 millisecond cadence between the start and end of each observation window. As a performance optimisation, these calculations were performed once and fitted with an interpolating spline, which was subsequently evaluated at the observation times.

The PAB unit vectors were then transformed into the body-fixed frame through the application of the transformation matrix T shown in Equation 5.5, and the latitudinal (ψ) and longitudinal (θ) components were extracted from the vector form via:

$$\psi = \arcsin \nu_z \tag{5.6}$$

$$\theta = \arctan 2(\nu_y, \nu_x) + \pi \tag{5.7}$$

where ν is the unit PAB vector in the body frame. This approach allows the flux measurements to be associated with a coordinate pair and thus a brightness map of the Ajisai surface can be created.

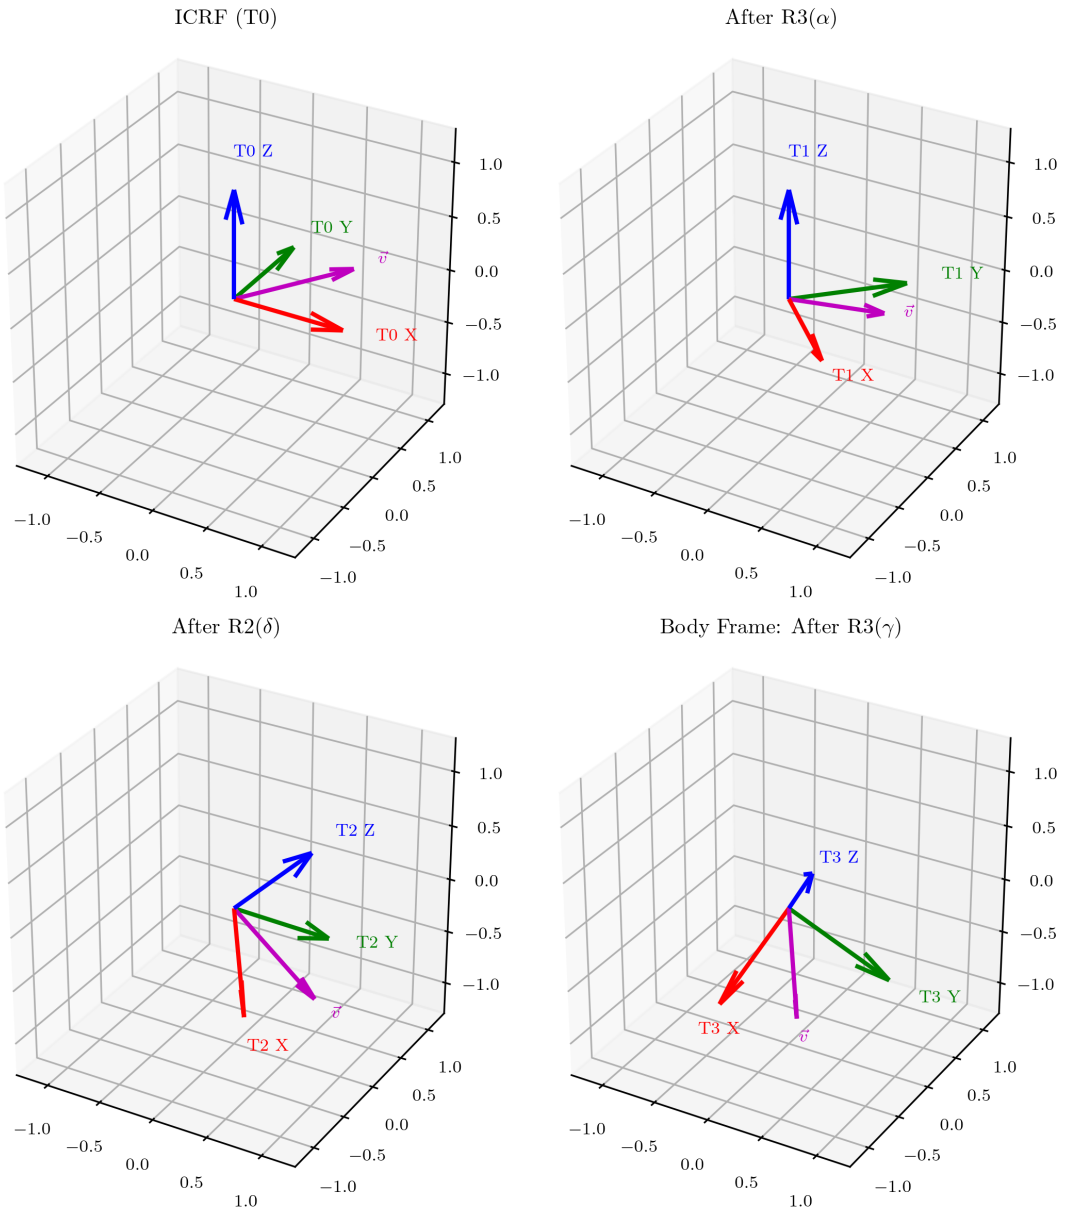

Figure 5.4: Illustration of the rotation sequence to transform a vector $\vec{v}$ from the ICRF to body frame. In this example, both the axes (T0) and vector are transformed by α, δ, and γ of 45°.

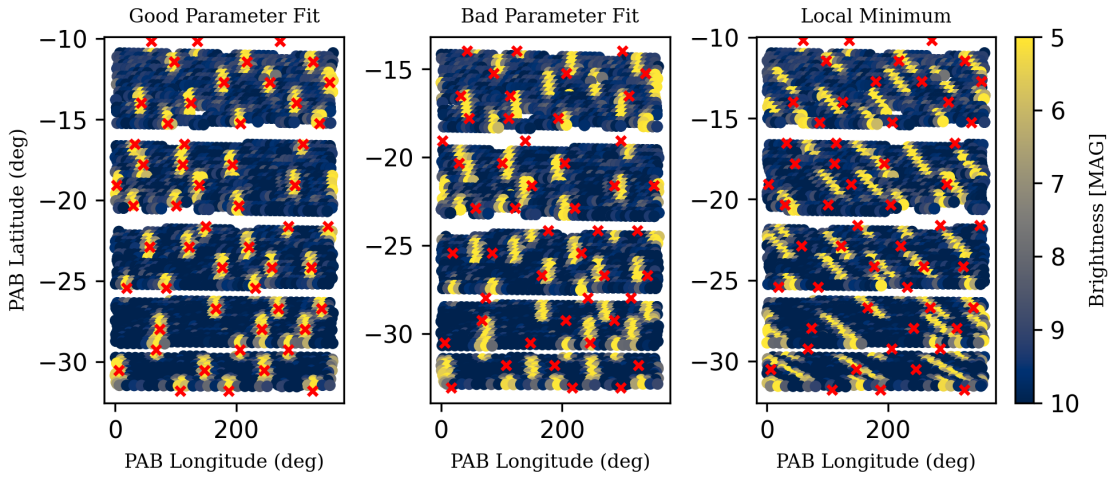

Figure 5.5: Examples of brightness maps from the 2019-08-01 observation with three different parameter sets compared to the Koshkin et al. (2017) mirror positions (red crosses); for clarity only the observed mirrors are overlaid here, as the comparison concerns the alignment of glints with mirrors. Otherwise, black crosses would be visible outside of the chosen latitude ranges. **Left:** The best-fit parameters transform all of the photometric glints onto an expected mirror position. **Centre:** A poor fit produced by offsetting the δ of the best fitted solution by 5°. **Right:** Local minima found in areas of the parameter space that project some of the photometric glints onto an expected mirror position, but project others in the gaps between mirrors.]

5.3.2 Interpreting Brightness Maps

Brightness maps represent a novel visualisation of the surface reflectivity of the satellite as a function of the derived spin-state solution. Each map is plotted in latitude–longitude body-frame PAB coordinates, where each point is associated with a measured brightness value (magnitude) indicated by the accompanying colour scale. Mirror positions from Koshkin et al. (2017) are overlaid as crosses:

- Black crosses define mirrors that were not observed.
- Red crosses define mirrors that were observed.

Better spin-state solutions produce glint patterns that align more closely with these reference mirror positions (see Figure 5.5). The effect of sub-optimal solutions is also illustrated: a solution trapped in a local minimum may still produce partial alignment, with some glints coinciding with reference mirrors, while most remain offset. Such a solution scores better than a poor (misaligned) fit but remains inferior to a fully aligned, optimal spin-state solution.

5.3.3 MCMC Fitting

The lat-lon coordinates can be compared to the modelled mirror points from Koshkin et al. (2017) in terms of spherical distance, which accounts for the periodic boundaries imposed by the longitudinal component.

Four parameters are required to transform the PAB vectors from the ICRF to the body-fixed frame: the spin-pole right ascension (α), declination (δ), rotation period (P), and phase offset (ϕ).

Parameter	Description	Prior Distribution	Support	Units
α	Spin-pole RA	$\mathcal{N}(\alpha_{emp}, 14^2)$	$[-140, -75]$	$^\circ$
δ	Spin-pole Dec	$\mathcal{N}(\delta_{emp}, 0.5^2)$	$[80, 90]$	$^\circ$
P	Rotation Period	$\mathcal{N}(P_{emp}, (0.000412)^2)$	$[-2.4355, -2.4322]$	s
ϕ	Phase Offset	$\mathcal{N}(180, 10^2)$	$[0, 360]$	$^\circ$

Parameter	Justification
α	Empirical RA + residuals
δ	Empirical Dec + residuals
P	Empirical period + RMS
ϕ	Full phase range centred

TABLE 5.1: Table of priors corresponding to the observation on the night of 2019-08-07 (see Figure E.3). The prior distributions are given in standard probability theory notation (Gelman et al., 2013). The initial phase offset distribution was set to search the full bounds but was then iterated upon after initial exploration as there is no prior knowledge of the phase offset. The convention adopted here assumes that Ajisai's spin pole lies towards the north celestial pole, implying a counter-clockwise spin direction and hence a negative rotation period.

The rotation period determines the angular velocity applied at each time step (see Equation 5.2), while the phase offset accounts for the satellite's initial orientation at the start of the observation.

To assess how well a given parameter set localises glint flux around known mirror orientations, a scoring metric is defined (Equation 5.8) that minimises the flux-weighted squared angular distance (see Appendix D for details on the angular separation calculation) between interpolated points and the closest mirror point. This penalises solutions that produce significant flux away from mirror locations. Well-aligned fits produce sharp, localised spikes, whereas poor fits smear them out. This behaviour could, in principle, be used to constrain the positions of unknown mirrors without relying on a priori knowledge of the mirror geometry (the spin-state priors of Table 5.1 concern only the rotation parameters, not the mirror positions).

$$\min_{(\alpha, \delta, P, \phi)} \sum_{i=1}^{n} f_i \cdot \min_j \left\{ \text{sep}(\mathbf{x}_i, \mathbf{y}_j) \right\}^2 \tag{5.8}$$

Subject to:

$$\mathbf{x}_i = (\text{ilon}_i, \text{ilat}_i) \quad \forall i \in \{1, \ldots, n\} \tag{5.9}$$

$$\mathbf{y}_j = (\text{K-Lon}_j, \text{K-Lat}_j) \quad \forall j \in \{1, \ldots, m\} \tag{5.10}$$

where n is the number of interpolated points (ilon, ilat), m is the number of mirror points in the dataset (K-Lon, K-Lat), f_i is the flux value for the i-th interpolated point, and $\text{sep}(\mathbf{x}_i, \mathbf{y}_j)$ is the angular separation between coordinates $\mathbf{x}_i$ and $\mathbf{y}_j$ in degrees.

An MCMC sampler was implemented using emcee (Foreman-Mackey et al., 2013) to explore the posterior distribution of these parameters. An example set of prior distributions is listed in Table 5.1, with the empirical estimates and uncertainties from Kucharski et al. (2016) used to inform the search band for each observation. Specifically, the empirical uncertainties are approximately $\alpha_\sigma \sim 14^\circ$, $\delta_\sigma \sim 0.5^\circ$, and $P_\sigma \sim 0.412 \times 10^{-3}$ s. Each MCMC production run consisted of 5000 steps, with the first 1000 steps discarded as burn-in. The best-fit parameters and their uncertainties were estimated using the mode and full width at half maximum (FWHM) of the posterior distributions.

Individual corner plots are generated using `corner` (Foreman-Mackey, 2016) and can be found in Chapter E (see Figure E.1 - Figure E.4).

5.4 SuperWASP Observations of Ajisai (2019)

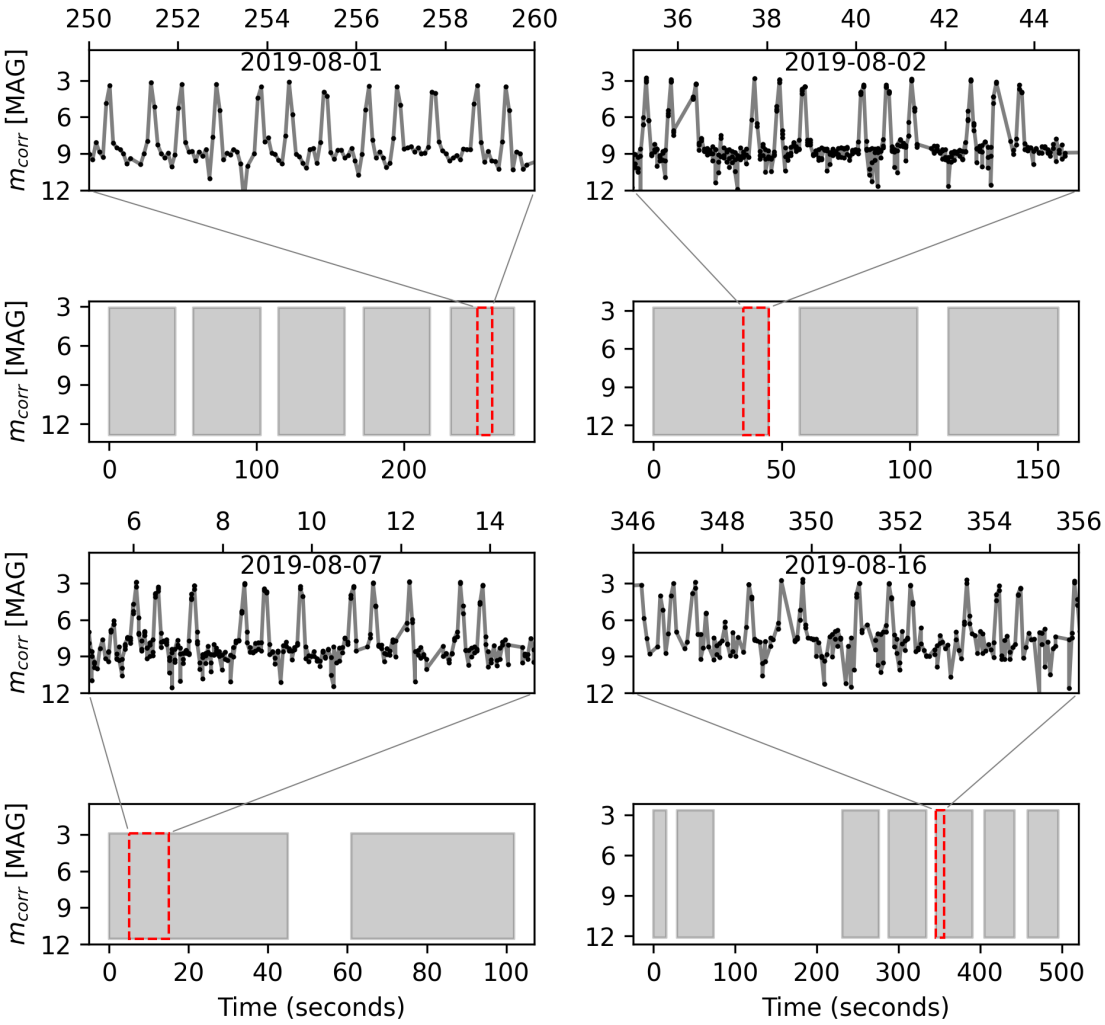

Figure 5.6: Summary of the SuperWASP observations of the Ajisai satellite. The upper panel of each observation shows a small section of the light curve. Because the spin period ~ 2.5 s is significantly shorter than the observation lengths (tens to hundreds of seconds), a plot of the full light curves cannot resolve the individual glints; for visual simplicity, the full light curves are represented in the lower panels as shaded regions indicating the time ranges of each field and the inter-field gaps. The dashed red rectangular boxes display the bounds of the zoomed-in region. The triplet features due to mirror reflections are clearly demonstrated within the zoomed-in regions. All brightness measurements are range-corrected.

5.4.1 Instrument and Observations

Ajisai was observed four times between the nights of 1st and 16th of August 2019. Figure 5.6 and Table 5.2 provide a summary of these data. The rate of motion of Ajisai compared to the background stars changes during each pass, producing a time resolution that is continuously changing. Figure 5.7 illustrates that the typical time resolution is 100–150 ms: an order of magnitude larger than the ~ 10 ms wide glint features (Calatroni et al., 2025). The photometric map methodology is sensitive primarily to the spacing between glints, so it is not necessary to resolve the individual glints.

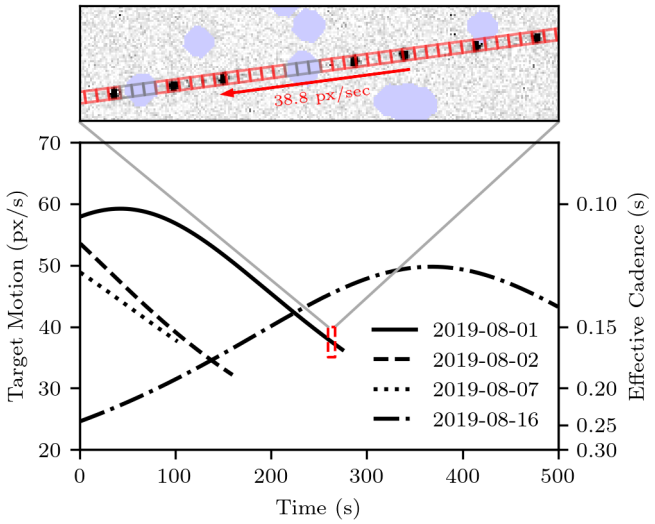

Figure 5.7: Bottom: The target motion and resulting extraction cadence plotted for the four Ajisai observations. **Top:** An example difference image segment with apertures visualised. Ajisai is visible as a faint continuous streak interspersed with point-source-like glints. Masked pixels containing difference-image artefacts are tinted blue, and apertures that intersect these regions are excluded. Photometry was extracted from the target streaks using a fixed aperture size of 6×6 pixels, chosen to avoid over-sampling the spatial resolution of the image. This figure was produced by Paul Chote.

T_{start} [UTC]	T_{end} [UTC]	D_{obs} [Days]	N_{data}	N_{fields}
2019-08-02 3:25:03	2019-08-02 3:29:38	12042.28	6787	5
2019-08-03 2:33:25	2019-08-03 2:35:18	12043.25	3665	3
2019-08-08 2:06:06	2019-08-08 2:07:55	12048.23	1597	2
2019-08-16 23:56:23	2019-08-17 00:05:42	12057.14	6224	7

Table 5.2: Summary of the 2019 SuperWASP observations of Ajisai. D_{obs}, N_{data}, and N_{fields} refer to the number of days since launch ($D_{launch} = 46654.86$ [MJD]), the number of observed datapoints, and the number of sidereal fields the satellite was observed in, respectively.

5.4.2 Results

Figure 5.8 shows the posterior distributions of the four spin-state parameters for each of the four observations. Figure 5.9 displays the brightness maps created from the best-fit MCMC parameters. The arrangement of mirrors into triplets is clearly seen across varying latitudes in the body-fixed

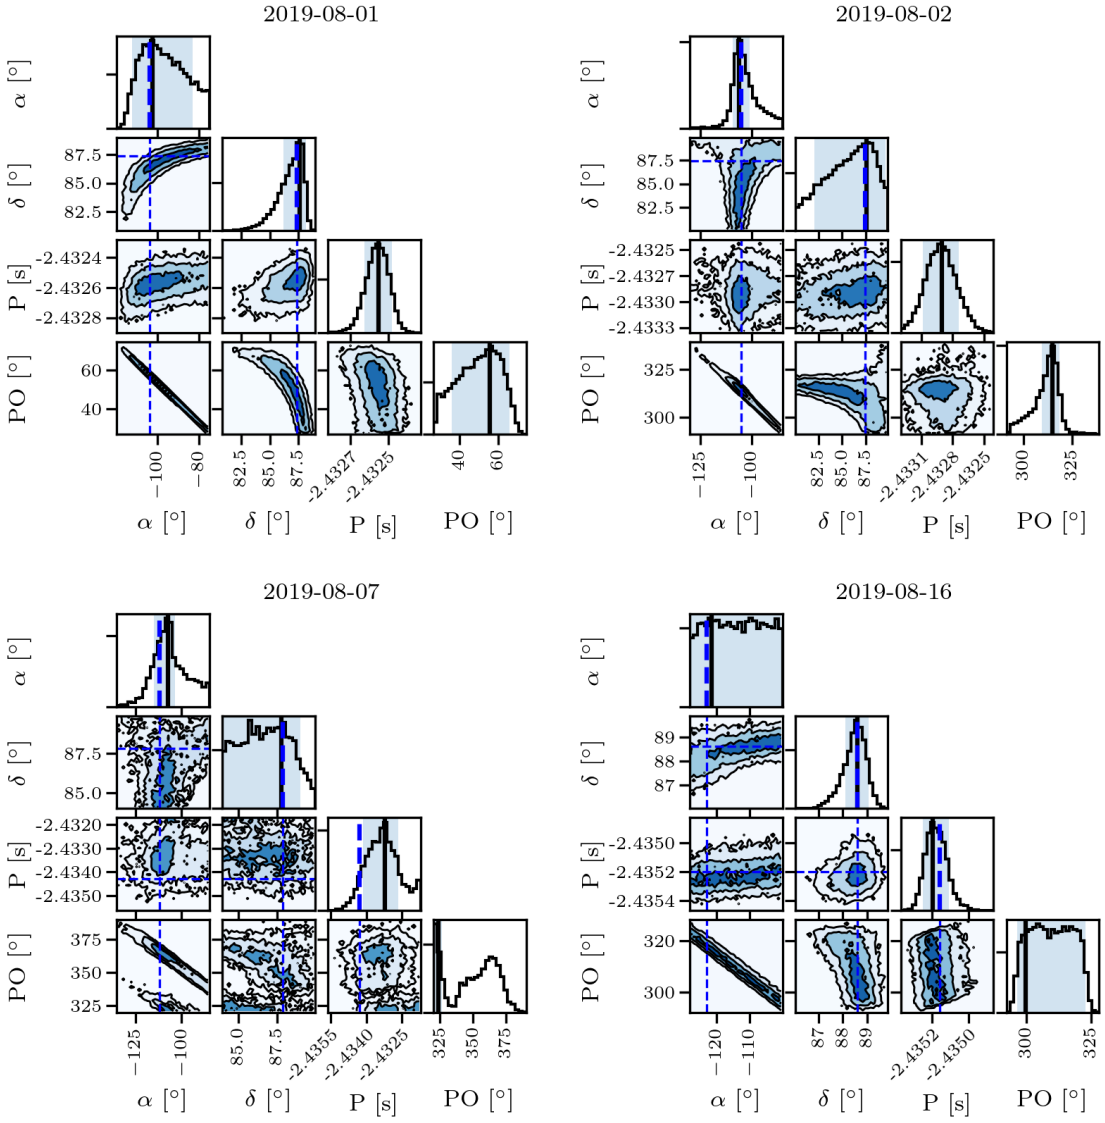

Figure 5.8: Corner plots from the MCMC runs on the data from each of the four observation nights (2019-08-01, 2019-08-02, 2019-08-07, and 2019-08-16, shown clockwise from top left). The blue dashed lines indicate the empirical value estimates from Kucharski et al. (2016), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimates (e.g., rotational period) do not coincide with a posterior maximum across all nights, which may be indicative of variations such as seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density, with contour levels corresponding to approximately 39.3%, 86.5%, and 98.9% credible regions. The rotational period and phase offset have been denoted as P and PO respectively for clarity in this figure.

Night Observed	Estimate	α [°]	δ [°]	P [s]
2019-08-01	MCMC	$-102.4^{+21.8}_{-10.0}$	$87.65^{+0.53}_{-1.49}$	2.43260 ± 0.00009
	Empirical	-103.7	87.39	2.4337
2019-08-02	MCMC	$-105.9^{+5.4}_{-3.4}$	$87.6^{+1.9}_{-5.4}$	2.43290 ± 0.00017
	Empirical	-105.0	87.46	2.4338
2019-08-07	MCMC	$-111.2^{+21.4}_{-7.8}$	$87.5^{+1.6}_{-2.9}$	$2.4333^{+0.0007}_{-0.0010}$
	Empirical	-111.8	87.84	2.4343
2019-08-16	MCMC	$-120.9^{+20.8}_{-6.8}$	$88.60^{+0.46}_{-0.51}$	$2.43520^{+0.00011}_{-0.00007}$
	Empirical	-122.9	88.60	2.4352

Table 5.3: Best estimates of parameters and their uncertainties derived from post burn-in MCMC samples for the 2019 SuperWASP observations. The estimate and asymmetric errors are derived from taking the mode and FWHM of the sample distributions. The α and δ of the empirical spin-pole (ICRF frame) and period calculated at time of observation are given; the spin-axis $\hat{\Psi}_{\text{emp}}(\alpha_{\text{emp}}, \delta_{\text{emp}})$ is inverted as this forms a right-hand coordinate system consistent with Koshkin et al. (2017). In this convention, the spin-pole of Ajisai points toward the north celestial pole and the spin direction is counter-clockwise.

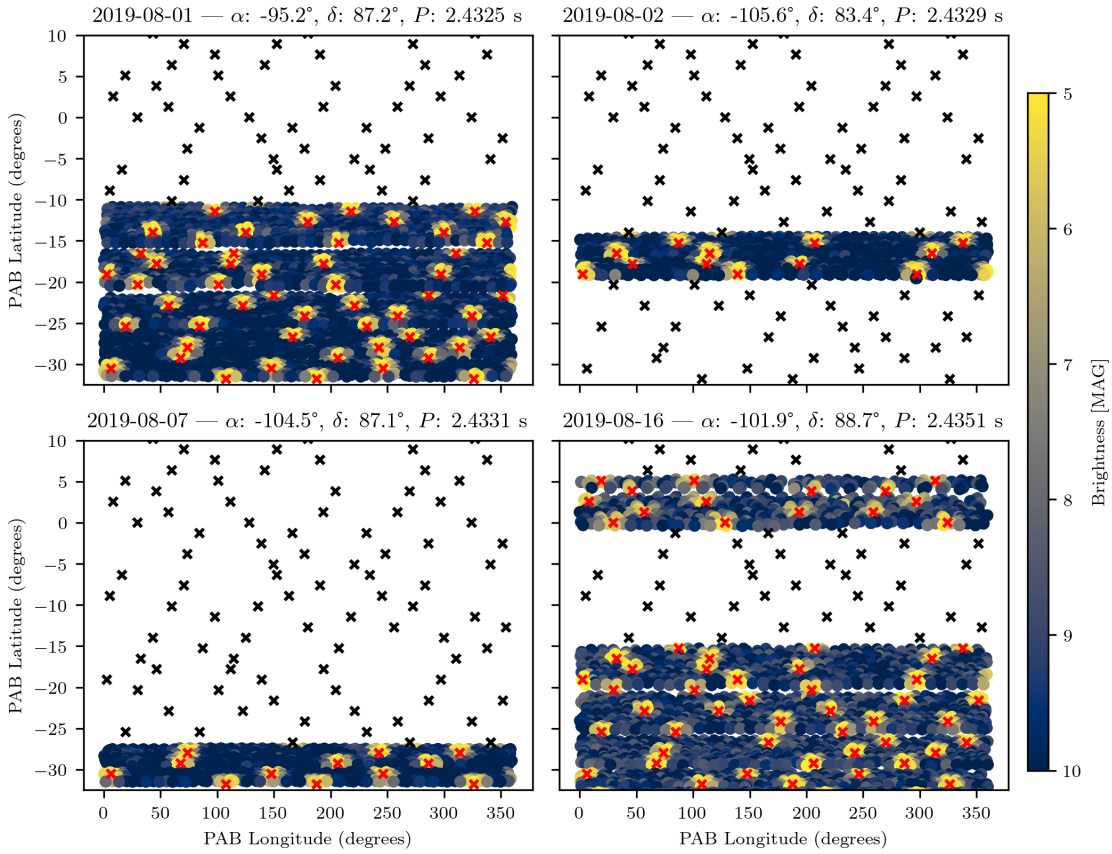

Figure 5.9: Brightness maps generated using the MCMC-derived parameter sets (shown in the subplot titles) with the maximum log-probability parameter sets for each observation, given above each brightness map. Red crosses denote the reference mirror points, overlapping with coloured points that indicate interpolated body-fixed lat-lon coordinates and associated measured magnitude values. Mirror position data was taken from Koshkin et al. (2017), used with permission.

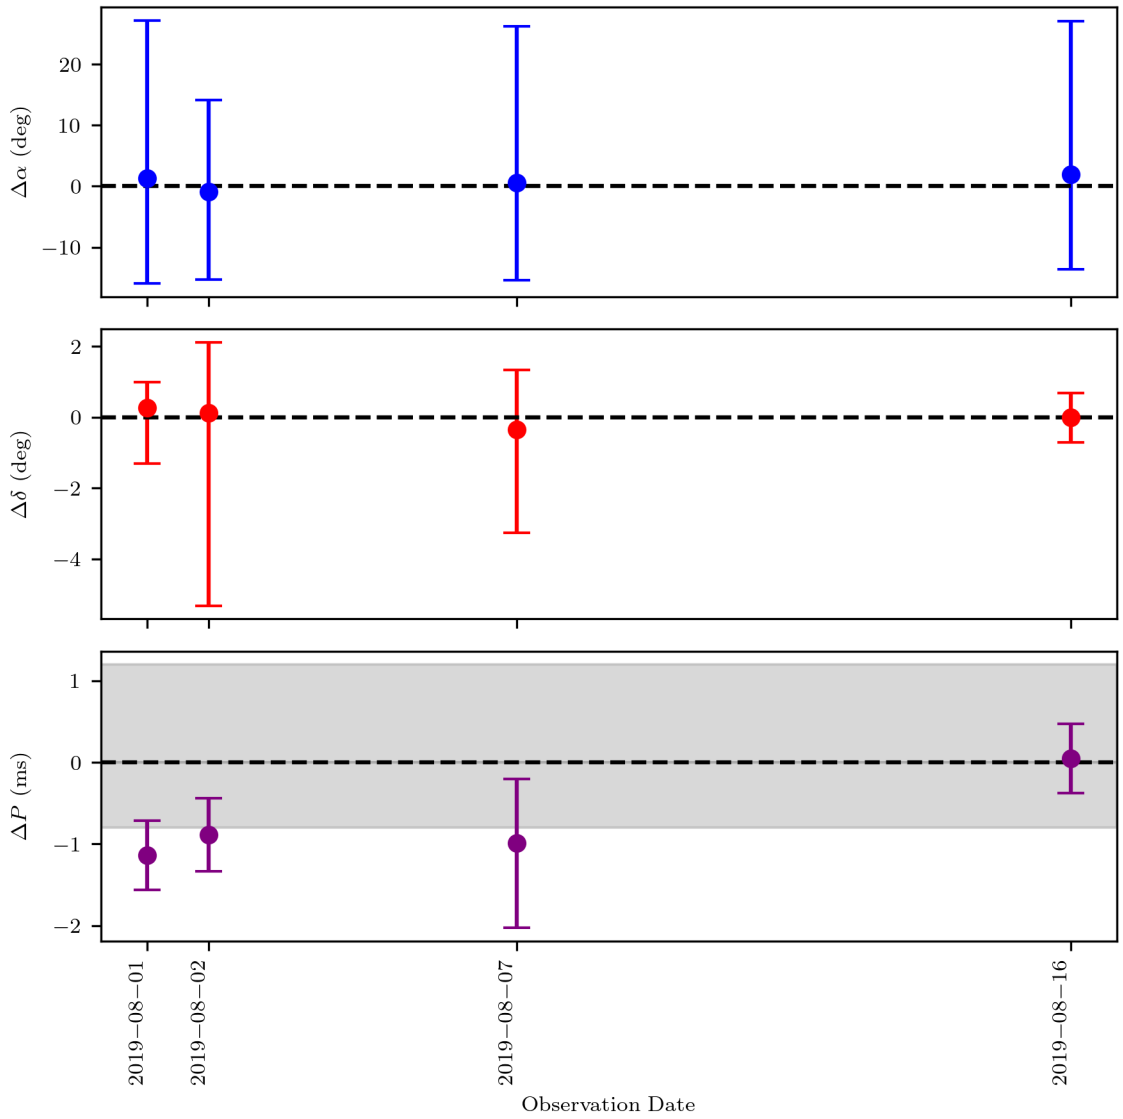

Figure 5.10: Residual measurements of parameters (estimate − empirical) for each SuperWASP observation of Ajisai. Error bars are calculated in quadrature from MCMC estimates and quoted residual/RMS measurements from the empirical models. The grey shaded bounds illustrate the maximum and minimum variation around the empirical period seen due to changes in solar illumination duration reported in Koshkin et al. (2017).

frame and the brightest regions of the surface reflectivity align well with the expected mirror positions. Table 5.3 displays the best estimates and errors.

Figure 5.10 displays the residual measurements between the MCMC estimate and the empirically predicted values with the errors combined in quadrature. The spin-pole estimations are close to the empirical value predictions; the first and fourth observations are particularly well constrained. The spin-pole δ residual for the second observation is not well constrained, as expected given this observation covers fewer mirrors.

The empirical model predicts only the long-term, secular evolution of the spin period and does not account for short-term variations, which could be driven, at least in part, by factors such as changes in solar illumination (Koshkin et al., 2017) and asymmetry in the surface reflectivity

introduced by the launch adapter ring and polar cap (Hattori and Otsubo, 2019). It is found that only the 2019-08-16 observation agrees with the empirical model, but all four lie within the interval (−0.8 to 1.2 ms) reported by Koshkin et al. (2017).

It is also worth noting that the relative motion between the observer and the target affects the measured apparent spin period regardless of the reference frame used, whether Earth-Centred Inertial (e.g., ICRF), Earth-Centred Earth-Fixed (e.g., ITRF), or the satellite's body frame. The rotation period is determined by forward-modelling the ICRF phase-angle bisector into the body frame and maximising its alignment with a static body-fixed mirror map; the fitted P is therefore the sidereal spin period, with observer/Sun relative motion handled inside the model rather than corrected afterwards. In contrast, Calatroni et al. (2025) first measure an apparent (synodic) period from inter-flash spacings within a triplet ("four-flash method"), and then apply a geometric correction to obtain the sidereal period.

It should also be noted that for the purpose of creating a right-handed coordinate system, the spin-axis results would be inverted to represent the true orientation of Ajisai, i.e., pointing towards the south celestial pole as opposed to the north.

5.5 STING Observation of Ajisai (2025)

5.5.1 Instrument and Observations

An observation of Ajisai was acquired with the STING telescope on the night of 29th August 2025. Additional observations were attempted on following nights, however the passes were too short to be able to give meaningful results in this approach.

5.5.2 Observation Planning

Observations were planned using the same field-hopping strategy as the GPS observations in §3.4, in which the telescope steps through a sequence of fixed sidereal pointings rather than tracking the target continuously. Within each pointing, a series of 2 s exposures are taken while the satellite drifts across the stationary field; these frames are subsequently stacked into a single image per field.

The sequence of pointings was generated by propagating the target TLE using `Skyfield` (Rhodes, 2019) at 1-second intervals across the observing window. Samples below an altitude of 30° or a declination of $\delta = -30°$ were removed to respect the dome horizon and mount limits respectively. For each remaining sample, the target position was computed both at the sample time and 2 s later so that the apparent on-sky trail vector during a single exposure could be estimated, with right ascension unwrapped to avoid discontinuities at the 0°/360° boundary.

The track was then divided into fields, the field was expanded one second at a time until the projected angular extent of the target's motion — cosine-corrected in right ascension — would exceed the camera field of view ($10.5° \times 7.5°$). Each field centre was set to the midpoint of the accepted coordinate range, producing a contiguous sequence of pointings that together covered the full visible pass. These were scheduled using the `ObserveTLESidereal` action. Any frames in which Ajisai fell within the Earth's shadow were discarded during the later extraction steps.

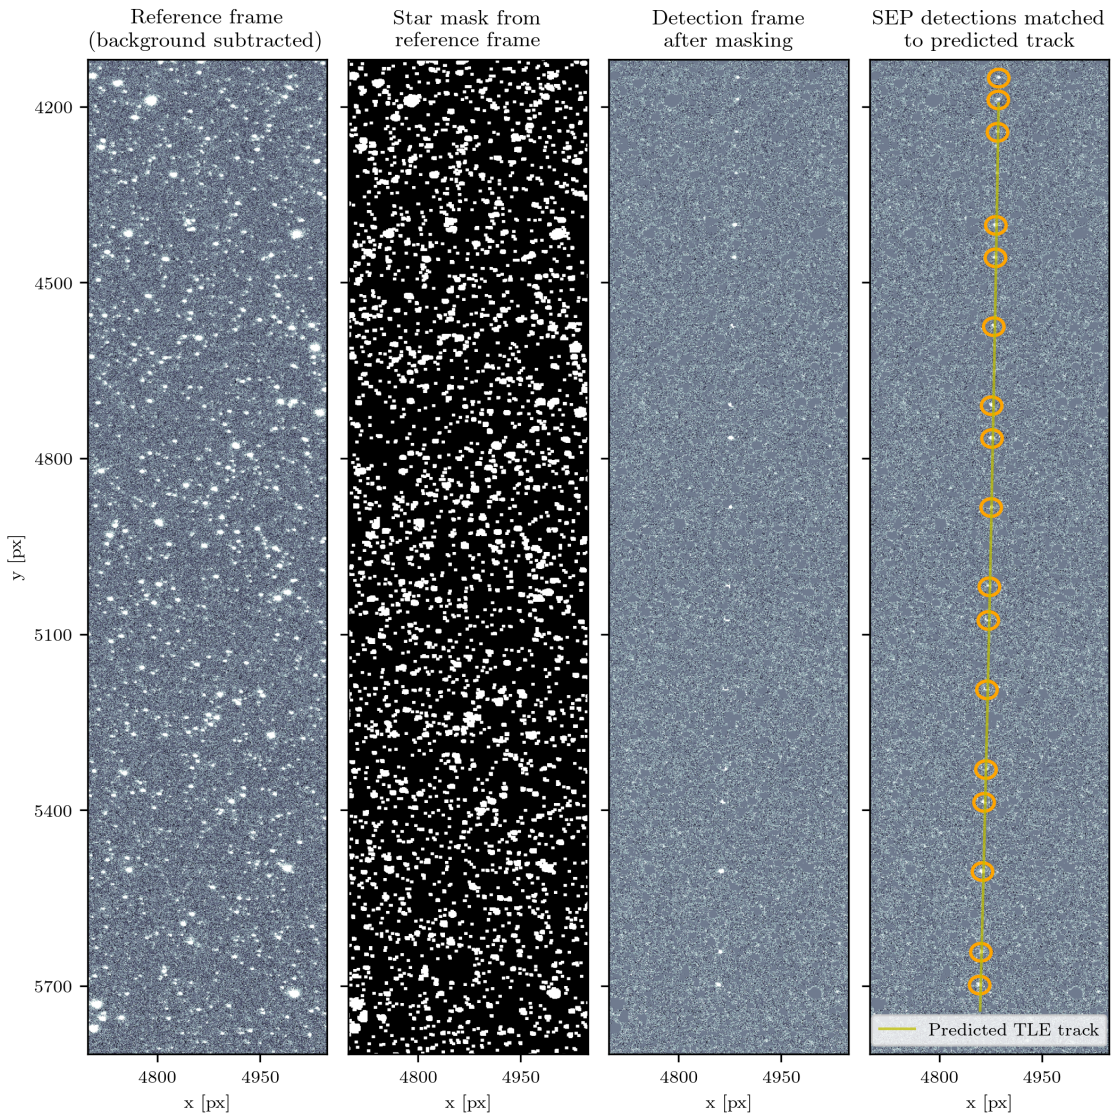

Figure 5.11: Illustration of the detection pipeline for Ajisai in STING science frames. From **left** to **Right**: reference frame after background subtraction. Stellar mask constructed from the reference frame. Detection frame after application of the image and stellar masks. SEP detections overlaid on the masked detection frame, with the predicted TLE track shown in yellow and matched detections circled in orange. All panels are cropped around the predicted satellite track for clarity.

5.5.3 Detection Pipeline

The STING observation of Ajisai was reduced from stacked FITS frames, with one stacked image produced for each action–field–camera combination. Each stack represents the high-speed raw frames accumulated during 30 exposures of 2 seconds each. In the timing analysis, the satellite track is sampled at 1 ms intervals using the exposure time recorded in the FITS header. The TLE track is evaluated using

$$N_{\text{sample}} = \text{round}(1000\, t_{\text{exp}} + 253),$$

where the additional 253 ms accounts for a full readout of the frame. This term does not affect the background subtraction or the initial SEP source extraction directly. However, it sets the temporal extent over which the predicted TLE track is evaluated, and therefore can affect the subsequent matching and timing interpolation of detections close to the beginning or end of the sampled track.

1. Each STING frame is first calibrated using the corresponding master dark and master flat for the camera. A reference image is used to estimate the background and to identify stars. The background is estimated using SEP (Bertin and Arnouts, 1996; Barbary, 2016) which calculates the per-pixel RMS noise. Stellar sources are masked by applying a threshold of 3 times the background reference to the reference frame. This threshold is chosen to balance sensitivity to faint sources while avoiding excessive false positives. Small gaps in the mask are filled to ensure completeness, and the masked regions are slightly expanded to fully exclude fixed celestial sources from the satellite search.

2. A predicted Ajisai track is generated from the TLE stored in the FITS header. The TLE positions are projected into the image plane using the astrometric WCS solution. Only portions of the predicted track for which Ajisai is sunlit are retained, since these are the parts of the pass capable of producing optical glints.

3. Two timing refinements are applied to the predicted track: a light-travel-time correction and a rolling-shutter correction. These corrections are the same as those applied in §3.4.

4. Candidate sources are extracted from the calibrated, star-masked detection image using SEP (Bertin and Arnouts, 1996; Barbary, 2016). Each extracted centroid is compared with the corrected TLE track, and detections more than 10 pixels from the predicted track are rejected. For accepted detections, the centroid is projected onto the local linear segment of the predicted track, allowing the detection time to be interpolated to sub-millisecond precision.

5. These detections are then combined across cameras, and the MCMC sampling and brightness-map fitting follow the method described in §5.3, with two modifications. First, extracted detections are treated as binary measurements, with a value of 1 assigned to a detection and 0 elsewhere. Second, a new set of empirical models from Shakun et al. (2026) is adopted as priors for this 2025 dataset; see Chapter C for details of the differences between the two model sets.

Figure 5.11 demonstrates the process of extracting the position of Ajisai from the stacked STING frames. The yellow line shows the TLE-predicted path after applying the light-time and

rolling-shutter corrections. Detections lying close to the track are accepted as Ajisai glints, while sources farther from the track are rejected. The accepted centroids provide the measured glint times used in the subsequent phase-angle-bisector and mirror-orientation analysis.

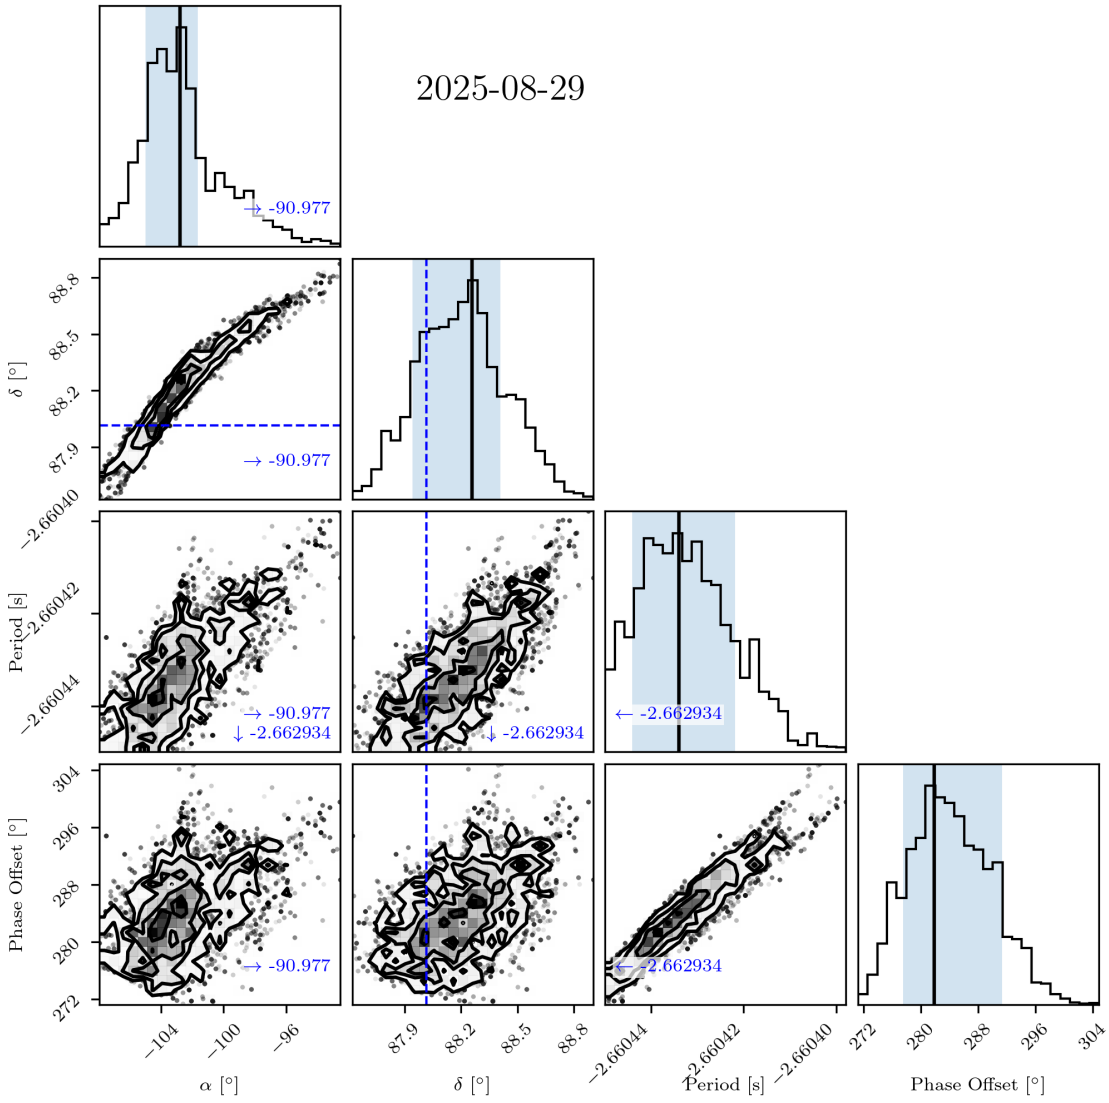

Figure 5.12: Corner plot from the MCMC run on the data from the night of 2025-08-29. The blue dashed lines indicate the empirical value estimate from Shakun et al. (2026), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimate (e.g. rotational period) do not pass through a minimum, which may be indicative of variations such as seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density, with contour levels corresponding to approximately 39.3%, 86.5%, and 98.9% credible regions.

5.5.4 Results

A similar frame-mapping procedure to that used for the four SuperWASP observations of Ajisai was applied here, but because these data were binary detections rather than calibrated brightness measurements, the mapping represents detections rather than brightness. Table 5.4 displays the best-fit parameter estimates and their associated uncertainties derived from the MCMC analysis of

the 2025-08-29 observation, alongside the empirical reference values from Shakun et al. (2026) for comparison. Figure 5.12 displays the corresponding corner plot showing the marginalised posterior distributions and joint credible regions for the spin-axis orientation (α, δ) and rotational period P. Figure 5.13 displays the resulting body-fixed brightness map reconstructed from the maximum-likelihood parameter set, with the interpolated detections overlaid on the reference mirror geometry.

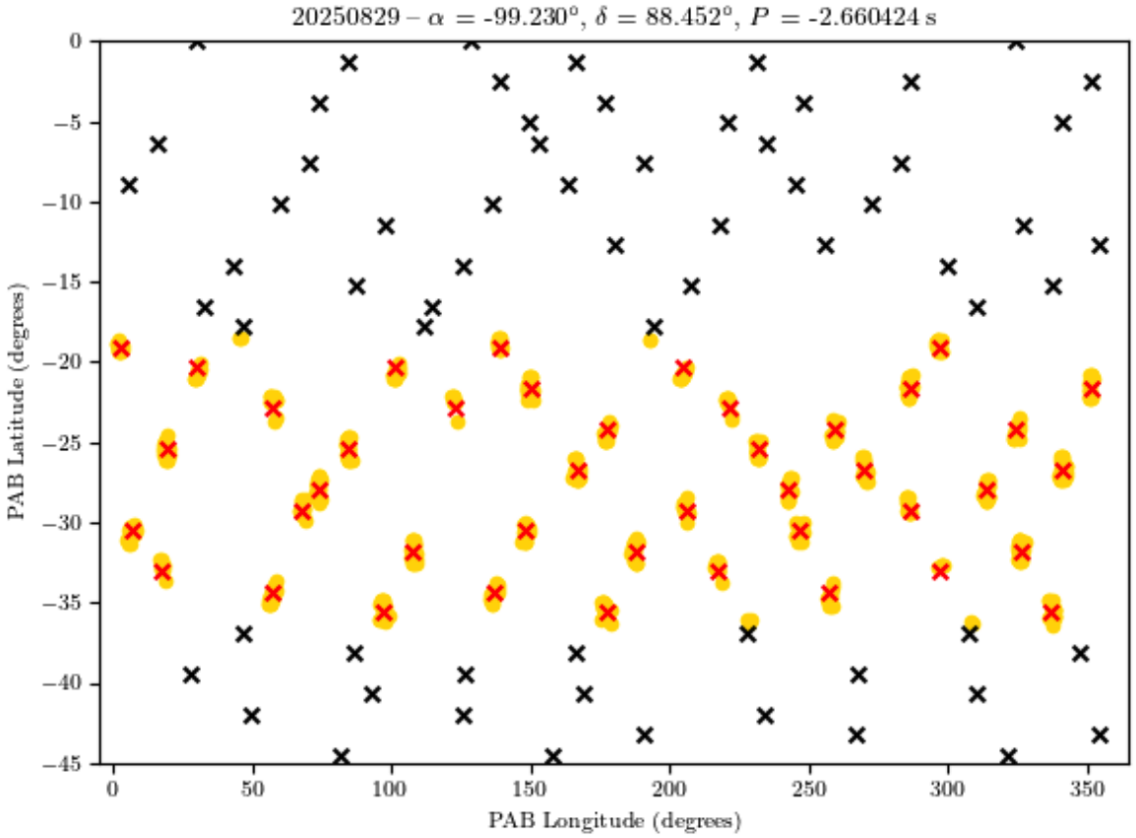

Figure 5.13: Brightness map generated using the MCMC-derived parameter sets (shown in the subplot title), with the maximum log-probability parameter set given above the brightness map. Red crosses denote the reference mirror points, overlapping with coloured points that indicate interpolated body-fixed lat-lon coordinates and associated measured detections. Mirror position data are taken from Koshkin et al. (2017), used with permission.

The recovered declination of the spin-pole, $\delta = 88.26^{\circ}{}^{+0.15}_{-0.32}$, is in close agreement with the empirical value and stated uncertainty of $88.01 \pm 0.03^{\circ}$ from the predicted empirical models in Shakun et al. (2026), with the two estimates consistent at approximately the 1σ level. The rotational period, $|\mathrm{P}| = 2.6604^{+1.2\times10^{-5}}_{-1.0\times10^{-5}}$ s, however, differs from the empirical value of $2.6629 \pm 4.12 \times 10^{-4}$ s by approximately 2.5×10^{-3} s, which is larger than the stated empirical uncertainty. Thus, while the MCMC provides a precise constraint on the dominant periodicity present in the detections, the recovered period is not consistent with the empirical estimate within the quoted uncertainties. This discrepancy may result from the limited temporal baseline and sampling of the single-night observation, as well as from systematic effects in the timing or identification of the detected glints, which may bias the recovered periodicity despite the small formal MCMC uncertainty. Further observations of the Ajisai satellite with STING and careful refinement of the glint detection

algorithm would be needed to determine this.

It is also useful to compare the formal precision of the STING-derived solution with that obtained from the original SuperWASP observations in Table 5.3. The STING fit gives substantially smaller statistical uncertainties than the four SuperWASP fits, particularly for the rotational period. For the SuperWASP data, the recovered period uncertainties range from approximately $\sim 7 \times 10^{-5}$ to $\sim 10^{-3}$ s, whereas the STING observation gives a formal uncertainty of order $\sim 10^{-5}$ s. The spin-pole uncertainties are also reduced: the SuperWASP right-ascension uncertainties are typically several to tens of degrees, while the STING value is constrained at the level of ~ 1–2°; similarly, the STING declination uncertainty is smaller than most of the SuperWASP estimates.

This apparent improvement in precision should not, however, be attributed solely to the smaller pixel scale of STING. A finer pixel scale may improve the reliability of the detection extraction by reducing blending with background sources, improving centroiding along the satellite trail (in this case, more or less point sources), and allowing detections to be associated more accurately with the predicted TLE track. Nevertheless, the precision of the recovered rotational period is more fundamentally controlled by the temporal baseline of the observation, the number of sampled rotations, the cadence or timing precision of the detections, and the signal-to-noise ratio of the recovered glints. This is supported by the fact that additional STING observations on subsequent nights were too short to yield meaningful constraints. For the spin-pole parameters, the limiting factors are instead dominated by the observing geometry, the number and distribution of mirror detections, and degeneracies in the body-frame mapping. Therefore, although the STING observation provides improved formal precision relative to the SuperWASP measurements, the offset in α indicates that the true accuracy remains limited by systematic effects such as single-pass coverage, illumination geometry, and model degeneracy.

The right ascension of the spin-pole, $\alpha = -102.80^{\circ}{}^{+1.15}_{-2.20}$, however, differs from the empirical estimate of $-90.97 \pm 0.91^\circ$ by approximately 12°, well outside the formal MCMC uncertainties. Several factors may contribute to this offset. First, at declinations approaching $\delta = 90^\circ$, the spin-axis orientation becomes weakly constrained in α: small angular displacements of the pole correspond to large excursions in right ascension, so the posterior in α may be broader than the formal credible interval suggests. This is consistent with the relative scale of the α and δ uncertainties in Table 5.4. Second, the empirical pole solution in Shakun et al. (2026) (and similarly the empirical period) is derived from a long-baseline dataset that averages over many years of observation, whereas the present estimate is based on a single night and may therefore reflect short-timescale effects, such as seasonal illumination geometry or limited longitudinal coverage of the satellite's surface during the observation window, that bias the recovered pole and period estimates.

The corner plot in Figure 5.12 supports this interpretation: the empirical α marker does not coincide with a minimum of the marginalised posterior, suggesting that the likelihood surface for this single-night observation genuinely favours a different orientation rather than simply failing to converge near the empirical value. Further observations with STING, combined with observations from the CLASP instrument and its optical and SWIR capabilities over a significant temporal baseline, could therefore be useful for constraining the seasonal illumination geometries that may contribute to discrepancies in the recovered parameters, as well as for improving the empirical models.

Parameter	Value	Uncertainty
Night Observed	2025-08-29	–
α [°]	−102.80	$^{+1.15}_{-2.20}$
δ [°]	88.26	$^{+0.15}_{-0.32}$
|P| [s]	2.6604	$^{+1.2\times10^{-5}}_{-1.0\times10^{-5}}$
α_{emp} [°]	−90.97	0.91
δ_{emp} [°]	88.01	0.03
P_{emp} [s]	2.6629	4.12×10^{-4}

Table 5.4: Best estimates of parameters, their uncertainties, and comparison with empirical values with their uncertainties stated in Shakun et al. (2026).

5.6 Summary

This chapter has demonstrated the feasibility of determining the spin-state of the Ajisai satellite through high-cadence photometric observations obtained using SuperWASP (2019) and STING (2025). By transforming TLE-derived vectors from the ICRF into the satellite's body-fixed frame and generating latitude–longitude brightness maps, a novel method for mapping surface reflectivity is showcased.

The methodology combines precision light curve extraction from streaked images with an MCMC-driven fitting process that aligns brightness maps in the satellite's body-fixed frame with modelled mirror positions. This indicates that constrained spin-state estimations can be derived from high-cadence photometry without relying on complex BRDF modelling or detailed prior knowledge of the satellite's shape, consistent with earlier photometric studies of Ajisai. The spin-state of a satellite significantly influences the morphology of light curves, governing the periodicity and modulation patterns that can be further exploited using machine learning models such as boosted decision trees (Shrive et al., 2024).

The best-fit spin-state solutions are consistent with empirical model predictions and observed variations in spin-period, affirming the adequacy of the mirror models in capturing the satellite's reflective behaviour. Notably, these results were achieved using relatively modest ground-based hardware adapted for high temporal resolution.

The generalisability of the approach warrants comment, since the method as applied here uses Ajisai's known mirror geometry as its reference: the scoring metric rewards spin-state solutions that project glints onto the modelled mirror positions. The procedure can therefore be inverted. Rather than using a known geometry to predict the locations of glints, the observed brightness map can be treated as a measurement of the locations of the reflecting surfaces, with the spin-state sought by maximising the sharpness and consistency of the resulting glint spikes. Ajisai then serves as a validation case for which the underlying geometry is known but it is not representative for targets that would be of interest for ADR. For tumbling targets with diffuse or extended surfaces, however, the assumption that brightness maxima correspond to a small number of localised specular facets breaks down. The prerequisite is therefore not prior knowledge of the target's structure, but the presence of discrete specular features in its light curve, since the geometry required for a specular reflection

is well understood. Simultaneous multi-colour photometry of the kind demonstrated in Chapter 4 could further strengthen this inverted approach: glints arising from different materials, such as solar arrays, multi-layer insulation, and antenna surfaces, may exhibit distinct colour signatures. In principle, a brightness map constructed across multiple bands, or directly in colour space, could therefore associate glints with similar spectral signatures and constrain them to common body-frame locations across an observation, providing sufficient structure to recover the spin-state.

These results affirm the value of high-cadence optical photometry for SDA applications, particularly for characterising the attitude of tumbling or rapidly spinning LEO objects, a critical capability in an increasingly congested orbital environment.

6 Three Techniques for Untangling Gorizont-6's Spin-state

> "In America, there's plenty of lite beer and you can always find a party! …In Russia, the party always finds you!"
>
> Yakov Smirnoff, Miller Lite commercial, 1986

6.1 Motivations

This study was motivated to combine the multi-colour observations of STING with linear polarimetry from the Liverpool Telescope as part of successful proposals to investigate tumbling satellites (PI: Klaas Wiersema). The polarimetric reduction and analysis within this chapter was undertaken by Klaas Wiersema. In addition to this, an observation of Gorizont-6 was acquired from the NGTS instrument in Chile on the night of 12th December 2023, which motivated the use of this multi-site data to investigate the spin-axis of the target. The following sections of this chapter have been published in Airey et al. (2026). The motivation and explanation of the Yarkovsky-O'Keefe-Radzievskii-Paddack (YORP) effect within this section have been expanded for the thesis.

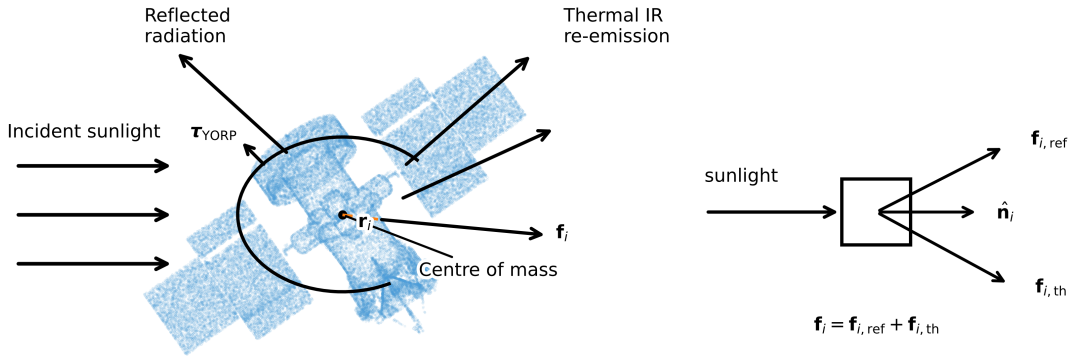

Figure 6.1: Schematic illustration of the YORP effect for the satellite **Gorizont-6**. **Left:** Projected shape model of Gorizont-6 with incident sunlight, reflected radiation, and thermal infrared re-emission shown schematically. The force on an individual surface element, f_i, acts at position r_i, relative to the centre of mass, producing a net YORP torque, M_{YORP}. **Right:** Decomposition of the recoil force on a facet into reflected and thermal components, $f_{i,ref}$ and $f_{i,th}$, together defined with respect to the surface normal $\hat{n}_i$.

The population of defunct, non-cooperative satellites in the near-Earth environment continues to increase, raising concerns for the long-term sustainability of Earth orbit and for the safety of operational spacecraft (see e.g., Blake, 2022; Buitrago-Leiva et al., 2024). Once attitude control is lost, a spacecraft will evolve under the action of external torques, and understanding those torques is essential for interpreting its observed rotational state. For highly asymmetric resident space objects (RSOs), one important long-term torque source is the YORP effect.

Physically, YORP arises because incident sunlight does not interact with a spacecraft uniformly. On each illuminated surface element, part of the incoming flux is reflected immediately, with a mixture of specular and diffuse reflection, while another part is absorbed and later re-emitted thermally in the infrared. If the body were perfectly symmetric, these small forces would largely cancel. For an irregular satellite with large appendages, antennas, solar panels, and an offset centre of mass, however, the forces act at different lever arms and do not cancel exactly. The resulting net moment can be written schematically as

$$\vec{M}_{\mathrm{YORP}} = \sum_{i=1}^{N} \vec{r}_i \times \vec{f}_i,$$

where $\mathbf{f}_i$ is the radiation force acting on facet i and $\mathbf{r}_i$ is the vector from the centre of mass to that facet. The magnitude and direction of this torque depend on the spacecraft geometry, optical properties, thermal properties, and Sun-object geometry, so the torque generally varies over both the spin cycle and the yearly solar cycle. Figure 6.1 illustrates the physical origin of the radiation-recoil torque acting on Gorizont-6. The total radiation-driven torque is obtained by summing the contributions from all surface facets over the spacecraft geometry. Owing to the highly non-spherical and asymmetric shape of Gorizont-6, these facet-level forces do not cancel exactly, leading to a non-zero net torque that can modify the rotational state over time.

This facet-based physical picture has been used to model defunct GEO satellites by discretising the spacecraft into surface facets and assigning each facet an area, normal vector, reflectivity, specular fraction, and emissivity. The resulting torques can then be propagated through Euler's rigid-body equations to follow the full attitude dynamics, or averaged over the spin/tumbling motion to obtain computationally efficient long-term evolution models. Work described in Albuja (2015) and Albuja et al. (2018) demonstrates that YORP can drive secular changes in spin rate and obliquity, and in sufficiently asymmetric cases can even move an object between principal-axis rotation and tumbling states. In GEO, where atmospheric drag is negligible and other environmental torques can be weak over long timescales, YORP may therefore play a major role in the rotational evolution of defunct spacecraft in graveyard orbits.

This mechanism is especially relevant for large, asymmetric communication satellites such as *Gorizont*. Its bus, solar arrays, and antenna structures provide a geometry for which asymmetric reflection and delayed thermal re-emission are expected to generate a persistent sunlight-driven torque. Over long timescales, even a very small net torque can accumulate, modifying the spin period, changing the spin-axis orientation, and potentially contributing to transitions between regular and non-principal-axis rotation. Establishing whether such behaviour is present is important not only for dynamical interpretation, but also for assessing future fragmentation risk and for informing planning for space situational awareness, on-orbit servicing, and active debris removal.

Utilising photometric measurements to estimate the spin-state of RSOs has been an active area of research (see e.g., Earl et al., 2018; Chote et al., 2019; Blake et al., 2021; Koshkin et al., 2024; Meyer et al., 2024; Airey et al., 2025b; Trelia et al., 2026). The rotational motion of these objects produces periodic signatures in optical light curves, from which candidate rotation periods can be derived. Multiple observations can also reveal changes in the rotational state through variations in period or light curve morphology (Papushev et al., 2009). A persistent difficulty, however, is identifying the true fundamental period, since harmonics such as $P/2$ or $2P$ can produce similarly convincing phase-folded light curves. In such cases, physically informed modelling of the spacecraft geometry and its illumination response can help distinguish the true rotation period from its aliases (Binz et al., 2014).

Additionally, long-baseline two-site or multi-site observations can help constrain possible spin-axis solutions for slowly rotating objects (see e.g., Koshkin et al., 2020; Reichegger et al., 2022).

In this study, the combined impact of simultaneous multi-colour and polarimetric measurements was investigated with particular emphasis on breaking the degeneracy between the fundamental rotation period and its harmonics, while long-baseline simultaneous two-site observations are used to probe possible spin-axis solutions for *Gorizont*. Interpreting these observations in the context of sunlight-driven torques provides a physically motivated framework for understanding the long-term rotational evolution of this defunct satellite.

6.2 GORIZONT-6 BACKGROUND

Gorizont-6 (NORAD: 13624, COSPAR: 1982-103A, also known as Gorizont No. 16L) was a Soviet geostationary communications satellite launched in 1982 and raised into graveyard orbit after its operations ended in 1989[1].

The family of Gorizont satellites all have a common design built upon the KAUR-3 bus, shown in Figure 6.2. This design has a cylindrical body with radius of 2 m and length of 5 m and two solar panels that are approximately 3.73 m × 5.45 m each (Albuja, 2015). All Gorizont satellites are now retired and they either move along graveyard orbits or drift through the GEO-belt.

Gorizont-6 was chosen for this study due to its favourable position in the sky from both Chile and La Palma on the night of 12th December 2023, and because previous observations (Chote et al., 2019) indicated that Gorizont light curves generally feature sharp specular glints, which could potentially be used to measure parallax-induced timing offsets. Subsequently, the object was incorporated into a broader polarimetric survey of GEO graveyard objects (Wiersema et al., in prep.).

An example of Gorizont satellite light curves are shown in Figure 6.3. Given these satellites are built upon the same KAUR-3 bus, it is expected that there will be a general similarity in light curve morphologies. This is seen in the form of very sharp glints and the appearance of doublet features in the Gorizont-30 and Gorizont-33 light curves in particular, which were present in some of the Gorizont-6 light curves.

[1] http://www.astronautix.com/g/gorizont.html

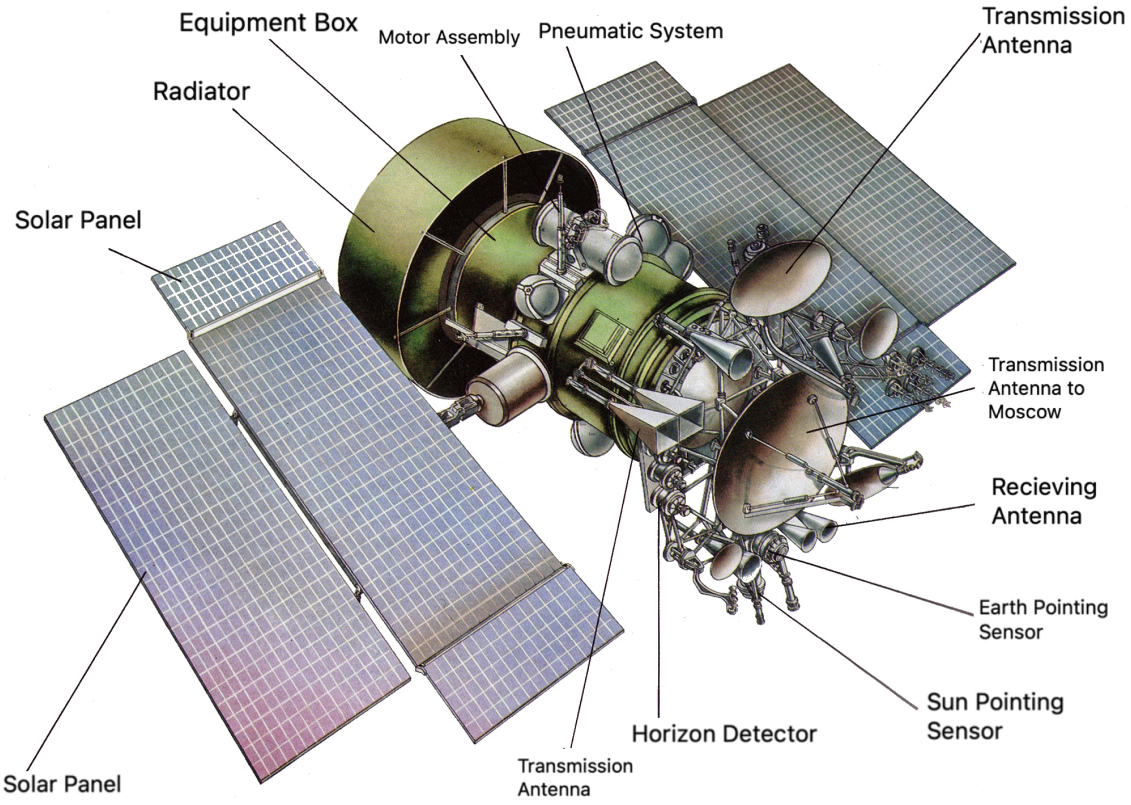

Figure 6.2: An example illustration of one of the satellites belonging to the Gorizont family. The key features of the KAUR-3 bus design are annotated. This figure is reproduced and translated from Pillet (2017).

6.3 Polarimetry

For satellites and space debris, optical photometry and spectroscopy form the conventional basis for characterisation, constraining their attitude, surface properties and behaviour (see e.g., Hall and Kervin, 2013; Jolley et al., 2015; Schmitt, 2020; Meredith et al., 2023; Chote et al., 2025).

In addition to multi-colour light curves and spectra, a powerful observational diagnostic to study objects in space is the optical linear polarisation of the received light. Reflection of sunlight off a solid body induces linear polarisation because the Fresnel reflection coefficients differ for the field components parallel and perpendicular to the plane of incidence ($R_p \neq R_s$). The main parameters that determine the amount of observed linear polarisation for specular reflection off a body in space are the refractive indices of the reflecting material(s), the reflection geometry (i.e., the angle(s) of the reflecting component(s) with respect to the observer and the Sun) and the shape (at both small and large size scales) of the reflecting surface. Additionally, the polarisation state of the light prior to reflection has an important influence; note that sunlight is generally unpolarised ($< 10^{-7}$) (Kemp et al., 1987). This leads to the possibility of using optical linear polarimetry (as a function of solar phase angle) to study both the orientation of the reflector and its (surface) material properties.

In stabilised, active satellites, the polarisation properties have received significant observational attention (see e.g., Speicher, 2015; Cegarra Polo et al., 2016; Wiersema et al., 2022), but the true potential for debris objects is still mostly unexplored.

As an object rotates, the angle(s) of the reflecting part(s) with respect to the Sun and observer

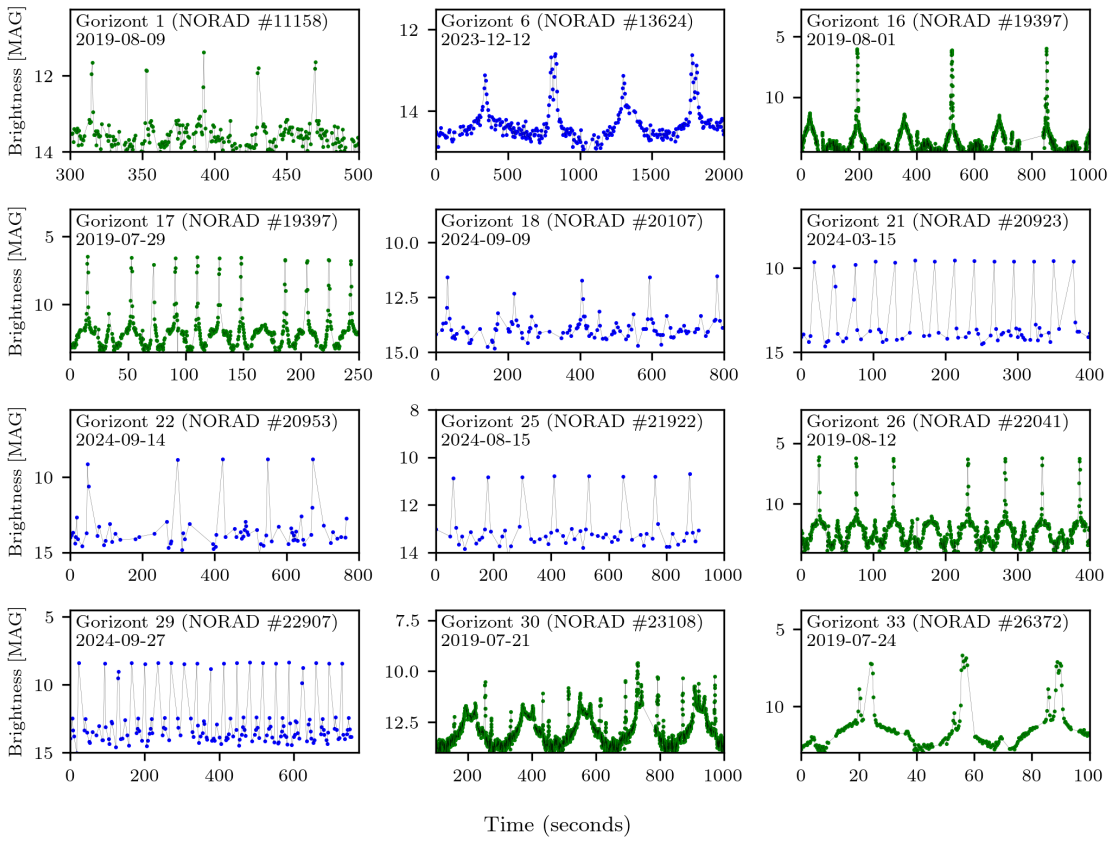

Figure 6.3: A collection of light curves of different Gorizont satellites observed with a RASA (Rowe Ackermann Schmidt Astrograph) in 2019 as part of a prototype survey of GEO (green scatter) and STING during the PE2 follow-up survey (blue scatter). The bandpass filter for the 2019 observations was $\mathtt{Gaia_G}$ for the RASA and for STING, the B_{RGB} band light curves are displayed with the exception of Gorizont-25, which only had a G_{RGB} band light curve due to visibility constraints of the enclosure. All light curves are shown within some arbitrary set of x-axis limits to best show their features.

rapidly change. Where light curves and colours can show degeneracy between different orientations of the object (particularly when the shape of the object is symmetric along one or more planes), the polarisation fraction and polarisation angle are not degenerate in the same way. Polarimetry may therefore provide a means to break spin-state degeneracies present in light curves. If more information is known about the satellite itself and/or its spin-state, polarisation may be able to uniquely identify the material and orientation of the reflecting components that contribute to specific light curve features.

6.4 Observational Strategy

The observational campaign was designed around a baseline of simultaneous four-colour photometry, enhanced by two complementary techniques to strengthen rotation-state identification: multi-site high-cadence photometry and broadband linear polarimetry. On the 12th December 2023, this baseline was augmented by a simultaneous observation with NGTS (Wheatley et al., 2017), which provided long-baseline, second-site photometry and by the Liverpool Telescope (LT; (Steele et al., 2004)) polarimetry, with the full combination of instruments used over several epochs shown in

Table 6.1.

The Next Generation Transit Survey is a ground-based observatory located at ESO's Paranal Observatory in Chile. The observatory consists of twelve 20 cm wide-field Newtonian telescopes, focused on the discovery and characterisation of extra-solar planets. In December 2023, one of the telescopes was reconfigured with an Andor Marana CMOS camera as part of a study into the use of scientific CMOS devices for high-precision photometry (Apergis et al., 2025). The telescope setup hosting the Marana camera operated independently of the main NGTS system, controlled by the same `rockit` software that operates STING (and other Warwick facilities on the ORM, (see e.g., Cooke et al., 2023; Mitchell et al., 2025)). A modified version of the STING pipeline was used to reduce the images and extract single-colour photometry in the standard NGTS filter ($520-890$ nm).

The observation of Gorizont-6 presented here was acquired during the on-site commissioning of this test instrument. The target was observed for a duration of $\sim$ 20 minutes with the *high dynamic range* mode[2] and 5 second exposures on the night of 12th December 2023. This exposure time was set to provide a high enough cadence to sample short-timescale variations and glints.

Polarimetric observations were obtained using the 2.0 m Liverpool Telescope under proposal numbers PL23B03 and PL24A03 (PI Wiersema). The dual-beam imaging polarimeter MOPTOP (Multi-colour OPTimised Optical Polarimeter) (Jermak et al., 2016, 2018; Shrestha et al., 2020) was used, following the observational strategy and data reduction described in Wiersema et al. (2022). Observations were acquired in `SLOW` mode, giving a polarimetric cadence of 20 s. This choice was driven primarily by operational constraints on the LT: `FAST`-mode sequences are limited to a maximum duration of ~800 s (file storage/I/O), whereas `SLOW` mode allows uninterrupted time series of essentially arbitrary length. Since the object's rotational modulation period was not known *a priori*, longer continuous sequences reduce the risk of sampling an incomplete cycle. In addition, the exposure times per angle in `SLOW` mode are $\sim 10\times$ longer than in `FAST`, improving the S/N for the target outside of glints; `FAST` is therefore only advantageous for very bright sources and/or signals on timescales much shorter than 800 s.

The `MOP-R` filter was chosen due to the high quality of the existing calibration values; the availability of extensive comparison data for active and inactive geostationary satellites (Wiersema et al. in prep); its lower sensitivity to moonlight compared to 'bluer' filters; and the instrument's good sensitivity at these wavelengths.

The tabulated values on the MOPTOP website[3] for the *cam3* and *cam4* instrument configuration were used. Pseudo-light curves were constructed from the MOPTOP data by computing the average flux $f'_i = (f_{cam3,i} + f_{cam4,i} + f_{cam3,i+1} + f_{cam4,i+1})/4$ (in contrast to the flux *ratios* used to construct the polarimetric measurements; see Wiersema et al. (2022) for further details), where two sets of waveplate angles for each average flux point (i.e., waveplate angles $0°, 22.5°$ make one average flux point, $45°, 67.5°$ make the next, etc). This results in a light curve with twice the cadence of the polarimetric data, since beam-switching required for accurate polarimetric calibration is not needed when measuring Stokes I (total intensity). Uncalibrated instrumental magnitudes are then computed as $m = -2.5\log(f')$ and an arbitrary zero-point offset is used to align it to the STING

[2] For further information regarding this mode see Apergis et al. (2025)
[3] https://telescope.livjm.ac.uk/TelInst/Inst/MOPTOP/

photometry.[4]

STING and the Liverpool Telescope are located at the same observatory on La Palma, separated by less than 250 m. This ensures a virtually identical viewing angle geometry for satellites in or near the geostationary belt, allowing a direct comparison of observed features. Simultaneous observations from La Palma and Paranal (where the NGTS telescope is situated) provide a large baseline (~8170 km). This offers a different viewing geometry such that reflections from the same components on the spacecraft will be received with a noticeable phase difference ($\gg$ timing uncertainties). This can then be leveraged in order to make constraints on the spin-state of the satellite. The locations of the sites are shown in Figure 6.4.

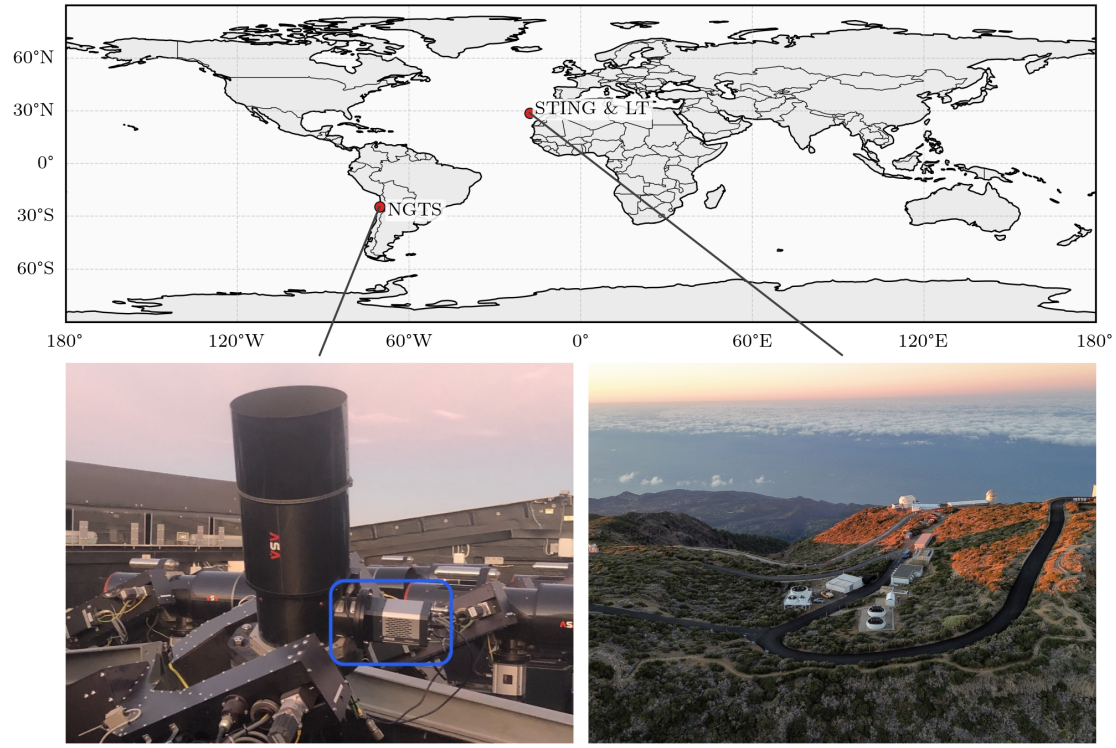

Figure 6.4: Top: Cartographic map illustrating the site locations for the NGTS (Paranal), STING (La Palma) and LT (La Palma) instruments. **Bottom Left:** A photograph of the Marana sCMOS camera, highlighted within the blue rectangle, mounted on one of the NGTS telescopes. This photograph is reproduced from Apergis et al. (2026) with permission. **Bottom Right:** A photograph of the ORM on La Palma, where STING and the LT are located.

6.5 Results

6.5.1 Period Analysis

The rotational period of Gorizont-6 was investigated using a Phase Dispersion Minimisation (PDM) (Stellingwerf, 1978) analysis. PDM is well suited to detecting periodic signals with non-sinusoidal light curve morphology (Fernández et al., 2005) and therefore provides a robust alternative to standard Lomb-Scargle techniques.

[4]The polarimetric reduction and analysis within this chapter was undertaken by Klaas Wiersema.

Night Starting	NGTS [UTC]	STING [UTC]	LT [UTC]
2023-12-12	00:34 – 00:54	23:44 – 02:29	00:12 – 00:33
2023-12-29	–	22:58 – 01:59	22:25 – 22:43
2024-01-21	–	03:02 – 04:27	04:03 – 04:23
2024-05-13	–	20:54 – 23:59	22:50 – 23:10
2024-07-19	–	–	21:34 – 21:54
2024-09-02	–	20:49 – 22:21	–

Table 6.1: Overview of the Gorizont-6 observations. The UTC start and end times for each observation are provided for the La Palma night (i.e., times after midnight were observed on the following day).

For this analysis, a custom PDM algorithm was implemented by myself based upon `PHMIN` [5]. An initial grid of candidate periods was defined based on visual inspection of each light curve, spanning 100–2400 s and sampled linearly with 30,000 trial periods. The best-fit period was identified as the period that minimised the PDM statistic:

1. For a given trial period P, phase-fold the light curve and fit a linear baseline whose subtraction minimises the PDM statistic $\Theta(P)$.

2. Assign the de-trended, phase-folded residuals to overlapping phase bins with approximately equal population.

3. Compute the intra-bin variances and form the PDM statistic $\Theta(P)$ from their ratio to the total variance.

4. Repeat for all trial periods and select $P_{\rm best} = \arg\min \Theta(P)$.

Conditional de-trending at each trial period was utilised because the baseline trend is small relative to the periodic signal and therefore cannot be robustly fit independently. While this pragmatic approach is common, it is not fully statistically rigorous: separating de-trending from period inference can mildly underestimate parameter uncertainties because baseline-model uncertainty is not propagated. A fully self-consistent treatment would fit a joint Bayesian model and marginalise over baseline parameters; this is not pursued in this study.

Uncertainties on the derived period were estimated using bootstrap resampling. For each bootstrap realisation, the light curve was resampled with replacement and the full PDM period search repeated. The resulting distribution of best-fit periods was approximately Gaussian; the mean period and its standard deviation as the 1σ uncertainty are reported.

Further details of the PDM methodology and statistical foundations can be found in Schwarzenberg-Czerny (1997).

The periods recovered from the individual photometric bands and from the colour light curves are summarised in Table 6.2. In several observing epochs, particularly where the single-band light curves exhibit strong harmonic structure, the periods inferred from individual bands differ by factors of approximately two or four, indicating ambiguity between the fundamental rotation period and its harmonics.

[5] `https://github.com/erikstacey/phmin`

Figure 6.5 demonstrates this ambiguity between harmonics. The trial period $P' \equiv P'_G$ is defined as the PDM-derived period from the G_{RGB} band (here $P'_G = 481.0$ s); as shown by the right-hand panels, folding on P' yields a poor fold with clear bifurcations, indicating that P' is not the true rotational period but a harmonic of it.

The left column shows stacked, normalised PDM Θ curves centred on the first four harmonics (P', $2P'$, $3P'$, $4P'$) for each band. In the single-band filters (B_{RGB}, G_{RGB}, R_{RGB}, i) there is comparatively little discrimination between these candidate periods (the Θ minima are broadly similar across the four harmonics), whereas the colour curve $B_{RGB} - i$ shows a much clearer preference for the longer-period solution. The $B_{RGB} - i$ colour index is adopted (as opposed to other colour indices) as it spans the largest wavelength separation available in the STING data, providing the strongest colour contrast between spectrally distinct reflecting components (maximising the visibility of 'bluer'/'redder' excursions associated with different facets).

The right panels show the corresponding folded, de-trended light curves for each trial period hP'. While the single-band folds remain double-valued or bifurcated at P' (and are not fully resolved at $2P'$ and $3P'$), folding on $4P'$ produces a coherent modulation across all bands and a consistent colour behaviour in $B_{RGB} - i$. Therefore, $P = 4P' \simeq 1926$ s is adopted as the true rotational period for this epoch.

The estimated period for the night of 21st January 2024 in the $B_{RGB} - i$ band is clearly unreliable due to the length of the observation arc. For this night, Gorizont-6 was observable by STING for ~ 5000 s, the shortest of the epochs. The limited length of this observation arc is comparable to twice the estimated period; in this case, twice the i-band bootstrapped estimate for that night is selected.

Overall, these results demonstrate that incorporating colour information into the period analysis can provide a more robust means of discriminating between the fundamental rotation period and its harmonics, particularly in cases where the single-band light curves alone yield ambiguous or multi-modal period solutions.

Figure 6.6 shows the evolution of the PDM-derived rotation period of Gorizont-6 as a function of time (days since the night of 12th December 2023), with each point representing the best-fit period from an observing epoch and vertical bars indicating the 1σ uncertainties. Over the observing baseline, the measured period increases asymptotically, suggesting a gradual approach toward a longer, quasi-stable spin period. One epoch (2024-07-19; orange scatter point) consists only of LT MOPTOP data; the period on this night via bootstrap resampling is estimated, restricting the period search bounds to lie between the lower (upper) 1σ limits of the preceding (following) observing epochs. In addition to the data, Figure 6.6 overlays several candidate trend models evaluated from weighted least-squares fits: an exponential approach-to-plateau model (solid blue curve), together with a quadratic fit (green), a sinusoidal fit (orange dotted curve), and a linear fit (grey dash-dotted curve) for comparison.

The asymptotic increase motivates an exponential approach model, which naturally captures a rapid early change followed by flattening towards a limiting value. This choice is supported by the goodness of fit: the exponential approach provides the smallest residual scatter among the functional forms tested, with $\chi^2 = 14.0$ for $\mathrm{DOF} = 6 - 3 = 3$ (i.e., $\chi^2/\mathrm{DOF} = 4.65$). By comparison, the alternative models in Figure 6.6 are substantially less consistent with the data given the quoted

Night	Band	P_{best} [s]	$P_{mean} \pm \sigma_P$ [s]
2023-12-12	B_{RGB}	481.0	481.1 ± 0.43
	G_{RGB}	481.0	481.1 ± 0.41
	R_{RGB}	481.1	481.1 ± 0.47
	i	1922.4	1924.5 ± 2.43
	$B_{RGB} - i$	1926.0	1925.2 ± 3.12
2023-12-29	B_{RGB}	501.6	501.3 ± 0.43
	G_{RGB}	501.4	501.2 ± 0.45
	R_{RGB}	501.7	501.2 ± 0.39
	i	501.5	501.4 ± 0.48
	$B_{RGB} - i$	2003.2	2001.2 ± 7.83
2024-01-21	B_{RGB}	1022.9	1022.4 ± 3.12
	G_{RGB}	1022.5	1023.5 ± 2.86
	R_{RGB}	1025.2	1023.8 ± 2.52
	i	1025.2	1023.1 ± 2.15
	$B_{RGB} - i$	2400.0	2234.3 ± 162.27
2024-05-13	B_{RGB}	1109.9	1109.2 ± 0.94
	G_{RGB}	1109.5	1109.3 ± 1.20
	R_{RGB}	1109.9	1111.1 ± 0.87
	i	1109.9	1110.1 ± 0.76
	$B_{RGB} - i$	2222.8	2223.4 ± 14.36
2024-09-02	B_{RGB}	1127.9	1127.1 ± 2.73
	G_{RGB}	1126.3	1127.8 ± 1.84
	R_{RGB}	1125.9	1126.3 ± 2.17
	i	1127.9	1126.1 ± 2.56
	$B_{RGB} - i$	2254.1	2258.9 ± 9.73

Table 6.2: PDM-derived periods and bootstrap mean periods with 1σ uncertainties for all STING observation epochs and bandpasses (and $B_{RGB} - i$ colour). The PDM arrived at the wrong harmonic solution (value highlighted in blue) for the night of the 21st of January 2024 in the $B_{RGB} - i$ colour index. Therefore, this can be corrected for by manually selecting twice the i band bootstrapped estimate.

uncertainties (linear: $\chi^2/\mathrm{DOF} \approx 131$; quadratic: $\chi^2/\mathrm{DOF} \approx 15$; sinusoid: $\chi^2/\mathrm{DOF} \approx 23$). The quadratic model is equivalent to the conventional spin-evolution parametrisation $P(x) = P_0 + \dot{P}x + \frac{1}{2}\ddot{P}x^2$, and is therefore directly comparable as a three-parameter fit, but it is disfavoured on physical as well as statistical grounds: as Figure 6.6 shows, it turns over within the observed baseline and predicts a decreasing period thereafter, for which there is no motivation under a monotonic de-spinning torque. While $\chi^2/\mathrm{DOF} > 1$ indicates that the scatter about even the best-fitting curve exceeds that expected from the formal 1σ errors, suggesting additional epoch-to-

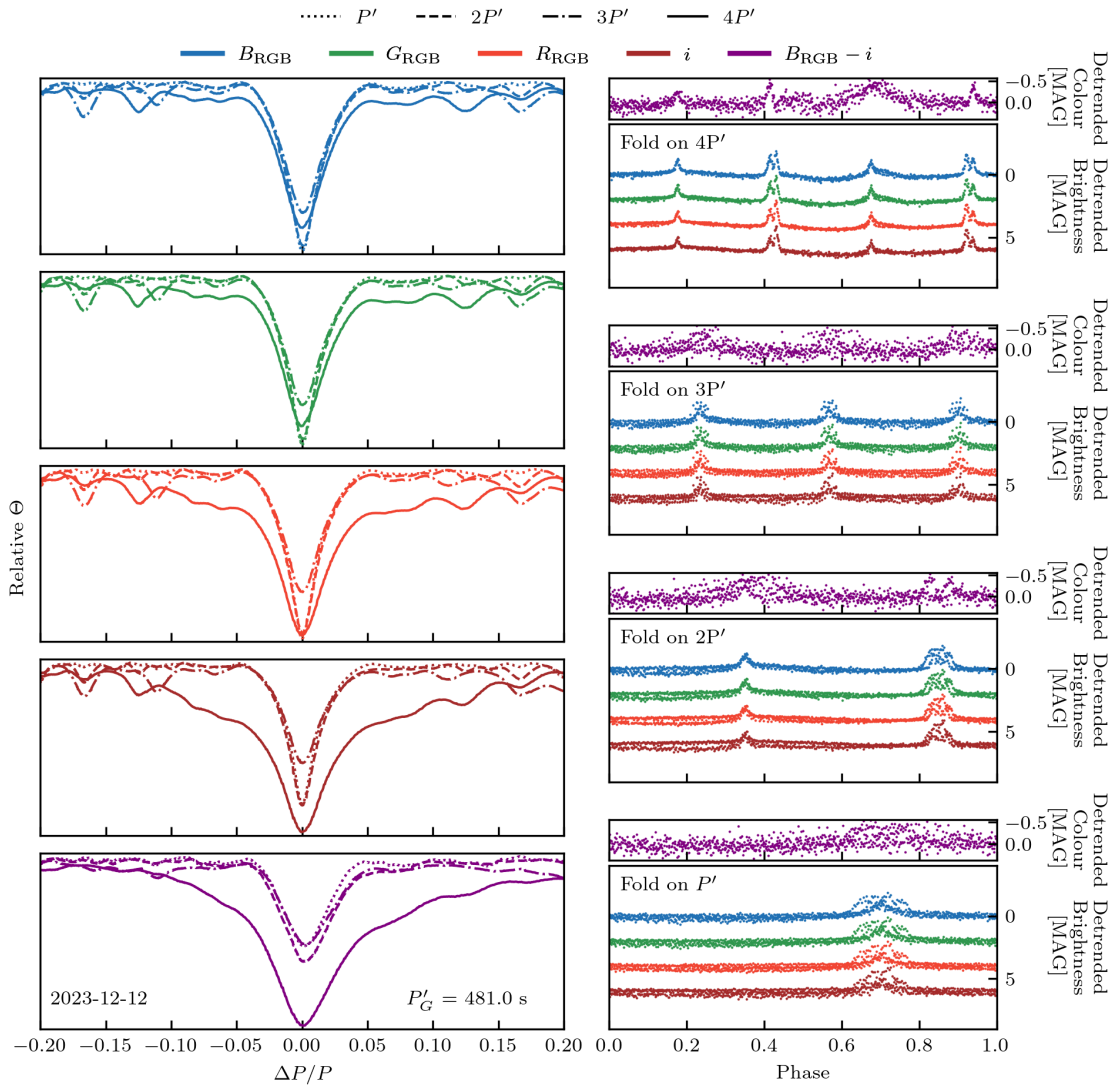

Figure 6.5: Left: Period-search statistic $\Theta(P)$ for the night 2023–12–12, shown as five panels (one per band: B_{RGB}, G_{RGB}, R_{RGB}, i, and $B_{RGB} - i$). In each panel slices of $\Theta(P)$ centred on the *trial* PDM period from the G_{RGB} band, $P'_G = 481.0$ s, and its first three harmonics ($2P'_G$, $3P'_G$, $4P'_G$) are plotted to demonstrate how well the PDM is able to distinguish between the different fold periods. **Right:** Phase-folded de-trended light curves in the B, G, R, and i bands (with $B_{RGB} - i$ shown above each panel), folded on the same set of trial periods hP'_G. Light curves are median-centred and vertically offset between bands for clarity; no additional amplitude scaling is applied.

epoch variance and/or unmodelled systematics, the exponential form nonetheless provides the most adequate empirical representation of the behaviour in this dataset.

Motivated by this, the long-term trend is quantified by fitting the data with an exponential approach model of the form

$$f(x) = a - be^{-kx}, \tag{6.1}$$

where x is time in days since 12th December 2023 and $f(x)$ is the predicted spin period. The best-fitting exponential curve is shown as the solid line in Figure 6.6. The fitted parameters are

$$a = 2267.7 \pm 7.3\,\mathrm{s}, \quad b = 341.1 \pm 7.2\,\mathrm{s}, \quad k = 0.0110 \pm 0.0005\,\mathrm{day}^{-1},$$

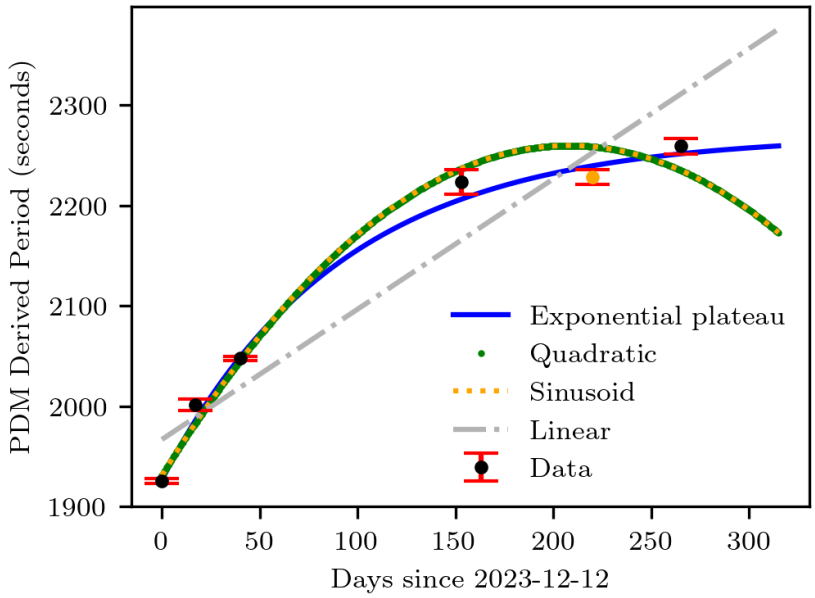

Figure 6.6: Evolution of the PDM-derived rotation period of Gorizont-6 as a function of time (days since the night of the 12th December 2023). Points show the best-fit periods from each observing epoch with 1σ uncertainties; the solid curve shows the best-fitting exponential approach model (Equation 6.1). The MOPTOP observation is shown by the orange scatter point.

with uncertainties corresponding to 1σ errors from the square roots of the diagonal elements of the fit covariance matrix. In this parametrisation, a is the asymptotic spin period approached within the current spin-cycle, b sets the initial offset from this plateau, and k controls the rate of approach. The associated e-folding timescale is $\tau = 1/k = 90.9 \pm 4.1\,\mathrm{days}$, meaning the difference $a - f(x)$ decreases by a factor of e approximately every ~ 91 days. With the best-fit parameters, the implied initial period at the start of the baseline is $P(t{=}0) = a - b \approx 1927\,\mathrm{s}$, increasing toward an asymptotic value of $\approx 2268\,\mathrm{s}$. The model also retains the physical content of the conventional parametrisation, since it yields a period derivative $\dot{f}(x) = bke^{-kx}$ that decays with time, with an initial value of $\dot{f}(0) = bk \approx 3.8\,\mathrm{s\,day^{-1}}$ at the start of the baseline. The exponential form is empirical rather than derived from a torque model, but an approach to a limiting period is what would be expected if the de-spinning torque diminishes as the spin-state evolves, since the YORP torque depends on the orientation of the spin axis relative to the Sun and is known to drive bodies toward asymptotic states in which the net torque vanishes. It should be emphasised that with six epochs spanning 310 days the data constrain only a decelerating increase in period, and cannot distinguish a genuine asymptote from a continued but slower increase; the fitted value of a should therefore be treated as an extrapolation beyond the observed baseline.

6.5.2 Phase Folded Light Curves

Figure 6.7 displays the phase-folded photometric and polarimetric measurements from STING and LT for the six observation epochs.

Considering only the measurements on the night of 12th December 2023 for the moment, the colour index ($B_{RGB} - i$) helps us distinguish between reflections originating from different facets on

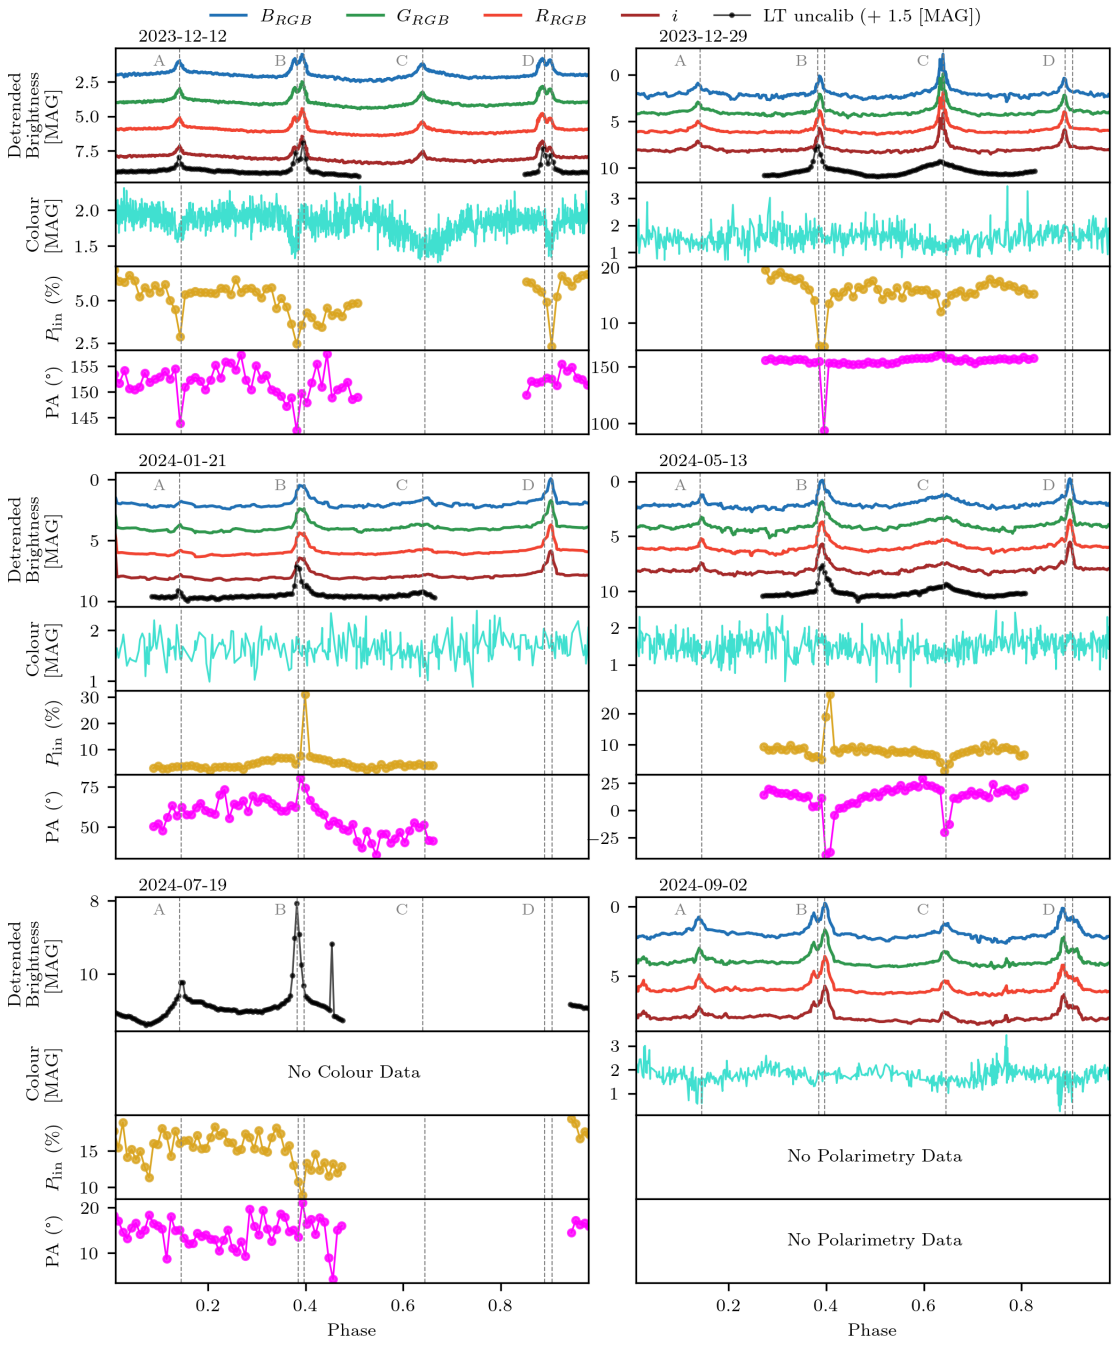

FIGURE 6.7: The six observations of Gorizont-6 with the STING telescope and polarimetry measurements (linear polarisation and polarisation angle) from the LT. The photometry (in all four STING bands) and polarimetry measurements are displayed as a function of rotational phase following a phase fold over the estimated bootstrapped rotational periods from the $B_{RGB} - i$ colour. In the case of the unreliable measurement on the night of the 21st January 2024, twice the bootstrapped estimate recovered from the i band data are taken. The individual STING band brightnesses were de-trended using the best fit linear baseline coefficients from the analysis prior, re-centred and offset for visual clarity. Like features across the photometry and polarimetry are annotated by the grey dashed vertical lines and are assigned A-D.

the spacecraft. Figure 6.7, peak A shows a narrow blueward colour shift compared to peak C, which is broader and overall 'bluer'. It can be speculated that peaks A and C arise from reflections off the back and front of the solar panels, respectively. The remaining doublet signatures (features B and D) are consistent with reflections from two surfaces on the spacecraft body with slightly different surface properties: in each doublet, the fainter of the two peaks is strongly blue, and the ordering of the bright and faint components flips between the two doublets, consistent with viewing the same feature from opposite sides as the object rotates.

This is also seen in the following observed epochs, where the reflection from the front of the solar panels eventually disappears due to the geometry not being favourable for specular reflections. This behaviour may suggest that the spin-axis has rotated away and is now approaching the polarisation angle for the material on the body given the increase in linear polarisation. Furthermore, the specular features begin to return in the latter two epochs.

6.5.2.1 Polarisation behaviour

The panels of Figure 6.7 show a rich polarimetric phenomenology, in line with expectations of a rotating object. The epochs with better quality data show a clear polarimetric signature associated with the glints, particularly obvious in the two epochs of December 2023 and May 2024. These generally appear as very short-lived dips in linear polarisation and a short-lived change in polarisation angle. These changes within glints are very short-lived, often limited to just one or two datapoints. This indicates that the obtained cadence of MOPTOP in SLOW mode (20 seconds as explained above), is generally insufficient to fully resolve the glint in polarisation space, particularly in polarisation angle. Observations in FAST mode (giving a 2 sec cadence) would have likely resolved this better, but at the cost of a loss of signal-to-noise.

Glints from a reflecting element on a satellite occur when the satellite's flat, reflective surface is oriented such that the Sun's reflection (light from the Sun) is directed straight towards an observer. In a single specular reflection of a simple material of a satellite it may be expected that the maximum linear polarisation will occur close to the Brewster angle (e.g., Stryjewski et al. (2010)), and the resulting light curve and polarisation curve of an active attitude-controlled satellite may look fairly smooth and have a simple shape with a low polarisation at small phase angles and high polarisation at high phase angles. This has indeed been observed for several satellites (e.g., Kosaka et al. (2020)); more detailed, high-cadence observations show deviations from this general behaviour, which often can be ascribed to the contribution of multiple reflecting elements to the received polarisation (Wiersema et al. in prep.).

For rotating satellites one can use the assumption of single reflection to generate wavelength-dependent linear polarisation expectations for a variety of shapes and refractive indices, see for example the work of Stryjewski et al. (2010) who provide examples of a rotating rocket body and the Hubble Space Telescope. As can be seen in the simulations by Stryjewski et al. (2010), the observed polarisation behaviour is strongly wavelength dependent, and dependent on the rotation axis and viewing geometry. While the different reflecting elements each have their own observed polarisation behaviour (as they each have their own orientation and material), they are tied together in that the configuration of the satellite is fixed. The received polarisation of a rotating object

also depends somewhat on the solar phase angle (Stryjewski et al., 2010). For relatively slowly rotating objects like Gorizont-6, the change in solar phase angle over the observed rotation period (or observation period) can be significant enough to play a role.

The polarimetry dataset covers a period of over 6 months, and a wide solar phase angle range. The six observations broadly appear to show a behaviour where the polarisation decreases with decreasing absolute phase angle. This is commonly observed in attitude controlled active satellites (see e.g., Kosaka et al., 2020; Papadogiannakis et al., 2022); the polarisation curve can be asymmetric with respect to the sign of the phase angle (Papadogiannakis et al., 2022). This behaviour has not been observed in a tumbling satellite before, and may imply a specific component of rotation.

Each polarimetric observation in Figure 6.7 shows strong, distinct features occurring on timescales of minutes. In most cases, changes in polarisation are accompanied by changes in polarisation angle, as expected from variability from reflecting components rotating in and out of view (i.e., they are seen under rapidly changing phase angles).

The two observations of December 2023 and the observation in May 2024 show clear sharp decreases of linear polarisation, occurring on short timescales, consisting of a shallow broad component and a sharp short duration spike, the latter seen in only one or two datapoints. These sharp depolarisation spikes seem to occur close in time to light curve flares (see Figure 6.7), both the broad light curve spike (see the polarisation data of 13th May 2024) as well as the narrow one (see the polarisation data of 29th December 2023). This is consistent with the expectation of a light curve glint at an angle that is far from the Brewster angle, contributing mainly unpolarised (or minimally polarised) light, depolarising the received signal. The polarisation signal appears to broadly mimic the general light curve shape of the flare in some cases (e.g., 29th December 2023), though this appears not always to be the case. The two December 2023 observations also show clear polarisation angle changes associated with the sharp polarisation dips, although the polarisation degree remains sufficiently high that the polarisation angle is still well-defined, consistent with this scenario. It is also clear from the data that the cadence of the polarimetry is too low to resolve these glints in polarisation angle space, and some of these spikes may be missed because of this relatively coarse (20 sec) sampling (19th July 2024). The high "baseline" polarisation at large phase angles improves the constraint on the polarisation position angle: in the high-S/N regime, $\sigma_\theta \approx \frac{1}{2}\,\sigma_P/P_{lin}$ (radians), whereas for P_{lin} comparable to the q, u noise level the position angle becomes effectively undefined (polarimetric sensitivity limit).

6.5.3 SPIN AXIS CONSTRAINT

Figure 6.8 displays the light curves from the joint NGTS and STING observation of Gorizont-6 on the night of the 12th December 2023. The glints are visibly offset between the two sites.

The PAB (see §2.1.1.3) is a useful quantity in determining whether a reflection has occurred off a surface as the specular condition holds true that a reflection will occur when the PAB vector and surface normal, $\vec{n}$ to a flat facet nearly coincide (Hall and Kervin, 2013).

The geometric construction used to obtain the candidate spin-axis directions $\pm S$ from two matched glint pairs is illustrated in Figure 6.9, and follows the methodology and geometrical logic outlined in Koshkin et al. (2024).

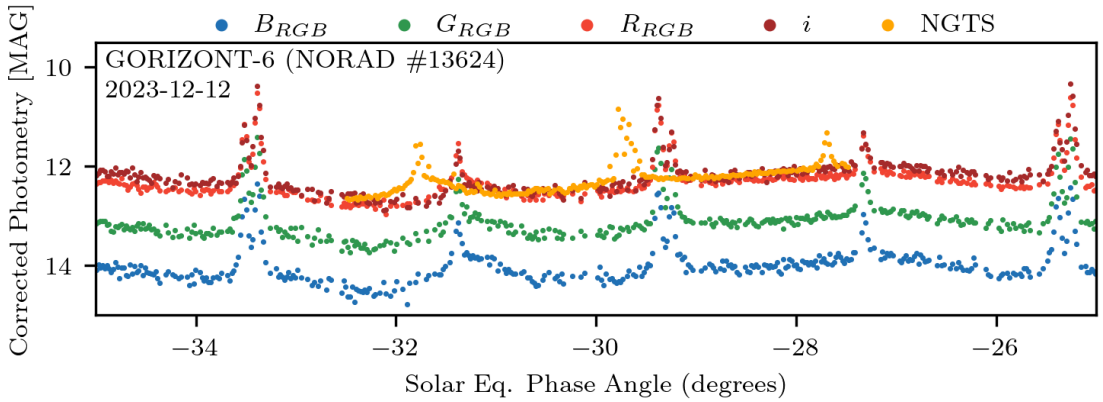

Figure 6.8: Multi-colour light curves as a function of UTC time. The simultaneous multi-instrument observations of Gorizont-6 were taken on the 12th of December 2023 with STING and NGTS.

The approach is to find the spin-pole as the intersection of two great circles on a unit sphere. Starting with at least two recognised glint events from the two light curves that can be matched, PAB vectors from both sites that correspond to that glint are acquired. These can be called $\vec{\nu}_1$ and $\vec{\nu}_2$, the normal to this glint-plane can be calculated, which is $\vec{\nu}_{12} = \vec{\nu}_1 \times \vec{\nu}_2$ and the bisector vector to the arc that passes through them, $\vec{b}_{12} = \vec{\nu}_1 + \vec{\nu}_2$. The normal vector to these two planes can be found by $\vec{n}_{12} = \vec{b}_{12} \times \vec{\nu}_{12}$ and this is repeated for another set of glint events occurring in both light curves ending up with a normal vector $\vec{n}_{34}$. The spin-axis will lie on the line of intersection of these planes, $\vec{S} = \vec{n}_{12} \times \vec{n}_{34}$.

A cost-function can then be defined to find a spin-axis that lies at the same pairwise angular distance from the PAB vectors at the instants of the glint events. This cost-function is shown in equation 6.2.

$$F = \sum (\beta_{i,j}(\alpha, \delta) - \beta_i(\alpha, \delta))^2 \tag{6.2}$$

Here $\beta_{i,j}$ refers to the angular separation between a trial spin-axis in the ICRS (International Celestial Reference System) and the calculated PAB vector at time of pattern glint. β_i refers to the mean PAB latitude of the glint pattern. Given the pattern consists of a glint seen at NGTS and then repeated later by STING, the mean estimate from these two observations can only be taken. The resultant cost surface from this least-squares fit is displayed in Figure 6.10.

The best fit spin-axis (minimisation of the cost function) occurs at $\alpha = 185.03°$, $\delta = -45.5°$. The antipode of this spin-axis is found at $\alpha = 5.03°$, $\delta = 45.5°$. No formal uncertainty is quoted: as Figure 6.10 shows, the cost surface has a broad, degenerate valley rather than a sharply defined minimum, reflecting the short observation arcs and the small number of glint pairs matched between the two sites. The result should therefore be read as a plausible orientation consistent with the paired glint events rather than a precise determination; tightening it would require longer simultaneous baselines, more matched glint pairs, or a third observing site.

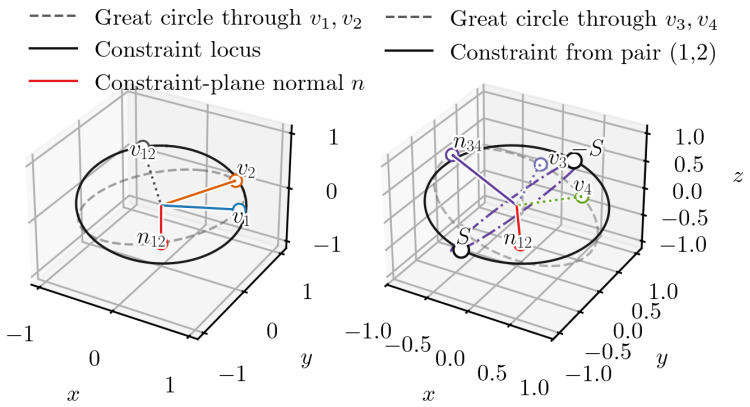

Figure 6.9: A diagram explaining the spin-pole estimation from matched glint pairs on the unit sphere. Unit PAB directions from two observing sites corresponding to the same glint are plotted as vectors $\vec{v}_i$. **Left:** A matched pair $(\vec{v}_1, \vec{v}_2)$ defines the plane spanned by the two directions (dashed great circle) and yields a great-circle locus of admissible spin poles given by the constraint plane with normal $\vec{n}_{12} = \vec{b}_{12} \times (\vec{v}_1 \times \vec{v}_2)$, where $b_{12} = \vec{v}_1 + \vec{v}_2$ is the bisector direction. **Right:** A second matched pair $(\vec{v}_3, \vec{v}_4)$ produces a second constraint great circle (purple); the candidate spin axes are the antipodal intersection points $\pm\vec{S}$, with $\vec{S} = \vec{n}_{12} \times \vec{n}_{34}$. A refined estimate is obtained by minimising the cost function in Eq. 6.2 over trial spin-axis directions.

6.6 Discussion

The observed spin-down period evolution of Gorizont-6 is characteristic of defunct satellites under the influence of the YORP effect. Such objects are expected to undergo cycles of uniform rotation and tumbling, repeatedly transitioning between the two modes (Benson and Scheeres, 2019). Over the course of the observations, Gorizont-6 is most likely in a uniform rotation cycle, consistent with the exponential dampening of its spin period and the morphology of the phase-folded light curves. Previous work (see Albuja et al., 2018; Benson, 2021) has shown how satellites can re-enter tumbling states from such conditions, and the findings here suggest that Gorizont-6 may similarly undergo transitions between stable spin, tumbling, and re-stabilisation phases. This behaviour parallels the YORP-driven cycles proposed for GOES-8, in which spin-up, energy dissipation, and realignment about different principal axes create a repeating dynamical evolution. Thus, the case of Gorizont-6 provides further evidence that the YORP effect plays a central role in shaping the long-term rotational dynamics of defunct satellites.

The attempt to constrain the spin-pole of Gorizont-6 yielded a range of possible solutions. The minimisation of the least-squares function identifies a plausible spin-pole orientation, but the constraint remains limited by the relatively short observation arcs and the small number of clear glint pairings between the two sites. These limitations would be alleviated by longer observational baselines or by including measurements from a third location.

While wide-field observations from instruments such as STING and NGTS enable the simultaneous monitoring of many objects and provide high-precision light curves, they remain fundamentally limited by the inherent degeneracies in photometric inversion. Symmetric spacecraft geometries, harmonic ambiguities in period determination, and the difficulty of associating specific light curve features with distinct reflecting surfaces all mean that even high-cadence multi-colour photometry cannot always yield a unique rotational solution. In contrast, single-target polarimetry offers a

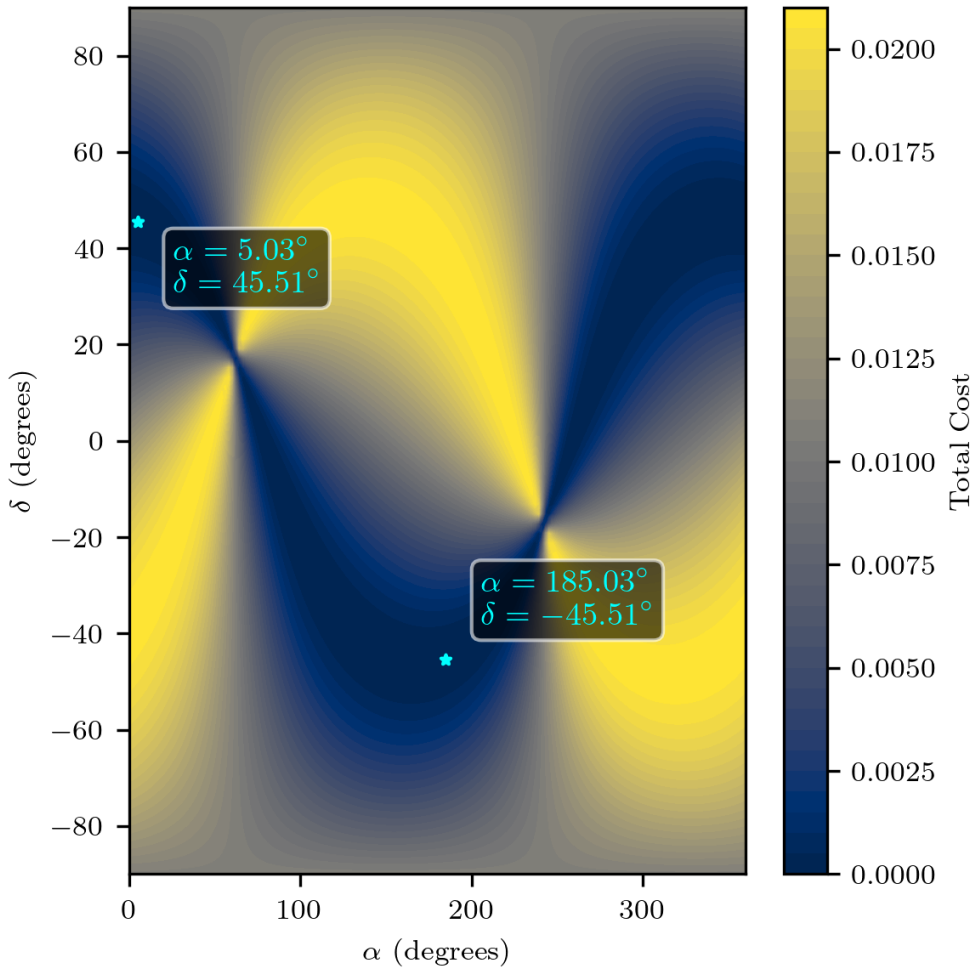

Figure 6.10: Total cost surface as a function of trial spin-axes constructed from α,δ of the spin-pole in the ICRS. Cyan stars mark the point on the grid where the cost function is minimised and the inversion of this optimised spin-axis. Darker regions are representative of trial spin-axes that are more likely to have caused the observed paired glint events from both sites.

substantially more diagnostic observable: the linear polarisation degree and angle are not subject to the same degeneracies as brightness, and they respond sensitively to changes in viewing geometry, surface specularity, and material refractive indices. These results for Gorizont-6 demonstrate this clearly—distinct dips in polarisation and rapid swings in polarisation angle pinpoint specific surface interactions even where photometry alone might admit multiple interpretations. Thus polarimetric observations provide the decisive constraints needed to break rotation-state degeneracies, validate the fundamental period, and link light curve features to physical spacecraft components. The optimal strategy for rotational characterisation therefore lies in a hybrid approach: wide-field multi-colour photometry for coverage and cadence, complemented by targeted polarimetry for unambiguous physical interpretation. In this case, it is possible to distinguish between components on the main bus and solar panels.

Taken together, these results reinforce the view that the YORP effect does play a role in the long-term rotational dynamics of defunct satellites in graveyard orbits, producing behaviours that parallel those seen in small-body populations. While Gorizont-6 provides a detailed case study, the broader

Gorizont series offers an opportunity to test whether satellites sharing a common bus architecture evolve along similar dynamical pathways. In total, 33 Gorizont satellites were launched[6], all built on the KAUR-3 platform; if their large-scale geometries and reflective surfaces are broadly comparable, one would expect recurring light curve morphologies and preferred spin-state outcomes across the population. To fully evaluate this hypothesis, longer observational baselines and systematic monitoring of the full series are essential. Looking forward, because spin-state strongly shapes light curve morphology—dictating both periodicity and modulation patterns—multi-colour time-series data from coordinated campaigns could also be leveraged by machine-learning approaches such as boosted decision trees (Shrive et al., 2024).

6.7 SUMMARY

A coordinated multi-instrument observation campaign was conducted which combined wide-field photometry with simultaneous high-cadence linear polarimetry to characterise the rotational evolution of the inactive geostationary satellite Gorizont-6. Multi-colour photometry enabled robust period estimation across six epochs, revealing a trend which was best fitted by a damped exponential curve, showing spin-down over nine months, consistent with YORP-driven rotational evolution. The simultaneous polarimetric data provided additional, independent constraints: rapid polarisation dips coincident with photometric glints traced reflections from specific spacecraft surfaces, breaking degeneracies between possible harmonics of the period and allowing confident identification of the true rotational period. Combined STING–NGTS observations further enabled a geometric inversion of paired glint events, yielding a plausible range of spin-axis orientations. Future work should involve combining these measurements with BRDF-derived light curve models to further constrain plausible spin-states, alongside dedicated long-term targeted monitoring of Gorizont/KAUR-3-class satellites to probe their dynamical evolution and quantify the impact of YORP on the Graveyard environment.

Taken together, these results demonstrate that simultaneous multi-colour photometry offers a cost-effective and efficient means of obtaining robust period measurements, especially across large satellite samples. However, achieving unambiguous confidence in the true rotation state — and in linking light curve features to physical components on the spacecraft — ultimately requires polarimetry, the sensitivity of which to surface geometry and material properties provides a unique diagnostic power. Future large-scale surveys of defunct spacecraft will therefore benefit most from a strategy in which wide-field photometry is combined with targeted polarimetry to provide accurate attitude characterisation.

[6] https://space.skyrocket.de/doc_sdat/gorizont.htm

7 Summary

> For wisdom is more profitable than silver. And her wages are worth more than gold.
>
> Proverbs 3:14

This thesis has investigated how high-precision photometry can be used to characterise resident space objects across both the geostationary and low-Earth-orbit regimes. In doing so it has addressed questions of instrument design, survey strategy, spin-state estimation and the complementary role of polarimetry, contributing new data, methods and physical insight to the field of space domain awareness.

Chapter 3 presented the characterisation of the STING instrument, the four-colour wide-field system that replaced SuperWASP-N at the Roque de los Muchachos Observatory. The photometric noise performance of each camera was quantified and shown to approach the theoretical limit set by photon statistics and detector read noise. The astrometric accuracy of the instrument was established by comparing measured streak positions of GPS satellites against their precise SP3 ephemerides, which provide millimetre-level positional precision and therefore serve as an effectively exact reference. GPS satellites were used because they are the only class of resident space object with an independent precise ephemeris of this quality; the calibration, background modelling, astrometric solution and timing chain being tested are common to all targets, and only the final centroiding step differs, with the trailed GPS streaks representing a harder case than the point-like appearance of a tracked GEO satellite. These measurements showed that STING achieves a sub-pixel precision when centroiding. Empirical horizon-limit surveys established the azimuth-dependent visibility envelope of the instrument, providing a data-driven basis for automated scheduling and ensuring that future observations are confined to regions of the sky where calibration quality is demonstrably reliable.

Chapter 4 reported the first large systematic colour survey of the GEO belt, producing full-night simultaneous four-colour light curves for over 100 active satellites observed across 17 fields over five weeks. The survey demonstrated that simultaneous colour measurements substantially reduce the errors that arise from sequential filter cycling during rapid brightness changes. Colour light curve maps revealed distinctive signatures associated with specific bus families: STAR-2 satellites were found to be collectively ‘bluer’ than any other class, BSS-702 satellites were the reddest, and there was tentative evidence that older Eurostar-3000 satellites are redder than their successors, possibly reflecting manufacturing changes in the solar-panel production process. Short-timescale glinting features were found to be unique in profile to each satellite and showed a characteristic

blue colour excess relative to the surrounding light curve, suggesting reflections from inclined flat surfaces such as solar panels. If these features prove to be seasonal, they offer a powerful route to individual object identification. Additionally, manoeuvres were detected and fitted in the light curve of the LDPE-3A satellite using a Cook–Torrance BRDF model, and a population of 54 solar-panel offset glints was identified and their colour responses catalogued.

Chapter 5 demonstrated that the spin-state of the geodetic satellite Ajisai can be recovered from high-cadence optical photometry obtained with modest ground-based hardware. Four SuperWASP observations from August 2019, and an additional observation with the STING instrument were analysed using TLE-derived geometry to transform sub-second specular glints into latitude–longitude brightness maps in the satellite's body-fixed frame. An MCMC-driven fitting process aligned these maps with modelled mirror positions, yielding spin-state solutions consistent with empirical predictions and independently measured period evolution — all without recourse to a detailed BRDF model or prior shape knowledge. The methodology establishes that constrained attitude estimation from photometry alone is feasible for well-characterised targets.

Chapter 6 extended the photometric toolkit to include both colour and polarimetry through a coordinated multi-instrument campaign targeting the defunct geostationary satellite Gorizont-6. Multi-colour photometry across six epochs revealed a monotonic spin-down best described by a damped exponential, consistent with YORP-driven rotational evolution. Simultaneous linear polarimetry from the Liverpool Telescope provided independent constraints that were critical for breaking the harmonic degeneracies inherent in photometric period estimation: rapid dips in linear polarisation coincident with photometric glints traced reflections from specific spacecraft surfaces, allowing confident identification of the true rotational period. Paired glint events observed simultaneously from STING and NGTS enabled a geometric inversion yielding candidate spin-axis orientations. The case of Gorizont-6 provides further evidence that the YORP effect plays a central role in shaping the long-term rotational dynamics of defunct satellites in graveyard orbits.

Taken together, these results establish that high-precision multi-colour photometry offers a cost-effective, scalable and physically informative route to the characterisation of non-cooperative space objects, and that its diagnostic power is greatly enhanced when combined with targeted polarimetry.

7.1 MEETING THE REQUIRED CAPABILITIES

§2.2.1 derived four observational requirements for satellite characterisation from the science itself, in the absence of formal performance thresholds attached to the published SDA capability requirements. Each has been addressed by the work presented here.

The *photometric* requirement, percent-level per-exposure precision with simultaneously recorded colour bands, was quantified in Chapter 3 by comparing a theoretical noise model against the measured scatter of satellite light curves under both bright and dark sky conditions. The resulting precision of approximately one per cent for targets brighter than 9th magnitude, degrading to ten per cent at 14th magnitude, proved sufficient to resolve the colour differences between bus families reported in Chapter 4 and the glint colour signatures analysed there. The expectation that the coarse plate scale would penalise only the sky-limited faint end was borne out.

The *temporal* requirement operated on two scales. For GEO targets, the ten-second survey cadence resolved the short timescale glints in Chapter 4, and full-night coverage across six epochs spanning ~ 9 months allowed the spin-down of Gorizont-6 to be traced in Chapter 6. For Ajisai, whose glints last of order ten milliseconds and are therefore beyond any practical frame cadence, streak photometry was shown in Chapter 5 to deliver effective time resolution of order one hundred milliseconds, sufficient to resolve the spacing between successive glints even though the individual glint profiles are unresolved.

The *astrometric* requirement was verified directly against an independent reference. Measured positions of GPS satellites agreed with their precise SP3 ephemerides to sub-arcsecond level, with one-to-two-arcsecond scatter, an order of magnitude inside the along-track errors of propagated TLEs at GEO, confirming that catalogue association is secure and that centroiding recovers positions to a small fraction of a pixel despite the coarse sampling.

The *spatial* requirement was met by design and quantified empirically. The wide-field allowed seventeen tiles to cover the visible GEO belt in a five-week campaign yielding light curves for 305 unique objects, in contrast to the handfuls of targets characterised by narrow-field programmes, while the horizon-limit mapping established what fraction of the belt is practically accessible and encoded those limits as scheduling constraints.

Against the capability requirements of the Ministry of Defence and UK Space Agency (Ministry of Defence and UK Space Agency, 2026), the thesis contributes to each of the characterisation objectives identified in §2.2.1: effective rotation rate and rotation axis through the spin-state reconstructions of Chapter 5 and the period and spin-axis analyses of Chapter 6; bus type and class through the colour survey of Chapter 4; manoeuvre history through the detection and fitting of the LDPE-3A manoeuvres; the solar panel offsets measured in Chapter 4 are also suited to repeated monitoring: tracking an individual satellite's solar panel offset across successive observing seasons would establish whether it is stable, and any change would be a non-cooperative indication that something about the spacecraft's power generation has altered.

7.2 Future Outlook

The work presented in this thesis opens several avenues for further research, both in extending the observational programmes described here and in developing the analysis techniques needed to exploit the growing volume of photometric data on the resident space object population.

Seasonal and long-baseline GEO surveys. The five-week STING survey reported in Chapter 4 represents a single-epoch snapshot of the GEO belt. Repeating the survey at intervals of several months would allow the seasonal dependence of colour signatures and glint profiles to be quantified. If the short-timescale glinting features identified in this work prove to recur with predictable phase-angle dependence, they could serve as robust fingerprints for individual satellites — a capability that would be of considerable value for identifying uncooperative or mis-catalogued objects. An extended temporal baseline would also clarify whether the colour evolution tentatively observed within the Eurostar-3000 class is a genuine manufacturing trend or a coincidence of the limited sample.

Extending to GEO debris and inclined objects. The survey presented here focused primarily on active, low-inclination satellites because the vertical field was cropped to 1° to reduce data volume. Widening the vertical field or adopting a dedicated tiling strategy for inclined objects would bring defunct satellites, rocket bodies and fragmentation debris into the sample. These populations are of particular interest for space sustainability assessments, and multi-colour photometry could provide constraints on surface material composition and degradation state that are otherwise difficult to obtain at GEO distances. Targets lacking reliable orbital elements, which includes much of the fragmentation debris population, would require a blind streak-detection mode instead. Fainter targets would need deeper stacks and correspondingly lower detection thresholds, and objects on inclined orbits, which traverse the field rather than remaining near a fixed position, would require a tiling and scheduling strategy designed around their apparent motion.

BRDF-forward modelling and light curve inversion. Chapter 4 demonstrated the use of a Cook–Torrance BRDF model to reconstruct manoeuvres in the LDPE-3A light curve, and Chapter 6 identified the need to combine photometric measurements with BRDF-derived light curve models to further constrain the spin-state of Gorizont-6. A natural next step is to develop a systematic framework in which parametric BRDF models — incorporating wavelength-dependent reflectance — are fitted jointly to multi-colour light curves. Such an approach could simultaneously recover attitude, shape and surface-property information, moving beyond the current practice of treating colour as a post-hoc diagnostic. The multi-colour capability of STING is well suited to this, as it provides instantaneous spectral leverage that single-band instruments lack.

Machine-learning classification from multi-colour light curves. The colour differences between bus families identified in Chapter 4, together with the unique glint profiles of individual satellites, suggest that supervised classification algorithms could be trained on the STING dataset to assign unknown objects to bus classes or even to identify specific spacecraft. Recent work has shown that boosted decision trees can classify LEO satellite platforms from single-band light curves (Shrive et al., 2024); the addition of simultaneous colour information is expected to increase discriminating power substantially. As the STING survey database grows with repeated observations, it will provide an increasingly rich training set for such approaches. The same framework could be extended to anomaly detection — flagging light curve behaviour that deviates from the learned norm for a given satellite, which could indicate attitude changes, component deployment or degradation.

Systematic monitoring of the Gorizont/KAUR-3 population. The YORP-driven spin-down observed for Gorizont-6 in Chapter 6 raises the question of whether similar dynamical evolution is occurring across the broader Gorizont series. All 33 Gorizont satellites were built on the KAUR-3 platform; if their large-scale geometries and reflective surfaces are broadly comparable, one would expect recurring light curve morphologies and preferred spin-state outcomes across the population. A dedicated monitoring programme targeting the accessible members of this series, using the same multi-colour photometric and polarimetric strategy developed here, would test whether YORP-driven cycles of spin-up, dissipation and tumbling are a generic feature of this satellite class and, by extension, of other families of defunct GEO spacecraft.

Polarimetry as a standard complement to photometry. The results of Chapter 6 demonstrate unambiguously that polarimetry breaks degeneracies that photometry alone cannot resolve, providing the decisive constraints needed to validate the fundamental period, link light curve features

to physical spacecraft components and distinguish between surface materials. Moving forward, the optimal observing strategy for rotational characterisation lies in a hybrid approach: wide-field multi-colour photometry for coverage and cadence, complemented by targeted polarimetry for unambiguous physical interpretation. Embedding polarimetric follow-up as a routine element of STING-led campaigns — for example by scheduling simultaneous Liverpool Telescope or MOPTOP observations for high-priority targets identified photometrically — would allow the most informative constraints to be obtained efficiently. As robotic polarimeters become more widely available, the cost and scheduling overhead of this complementary mode will continue to decrease.

Application to active debris removal and on-orbit servicing. The ability to determine the attitude state of a non-cooperative object is a prerequisite for most proposed active debris removal and on-orbit servicing techniques, which typically require a service vehicle to match the target's tumble and attach to it. The photometric and polarimetric methods developed in this thesis — particularly the MCMC-based spin-state estimation demonstrated for Ajisai (Chapter 5) and the multi-epoch period tracking applied to Gorizont-6 (Chapter 6) — are directly applicable to pre-mission characterisation of ADR targets. As the regulatory and commercial landscape for ADR matures, demand for ground-based attitude estimates of candidate targets is expected to grow, and the scalability of wide-field instruments such as STING positions them well to support this need.

Extension of the Ajisai methodology to other targets. The body-fixed brightness mapping technique developed in Chapter 5 exploited Ajisai's well-known geometry and mirror layout to validate the spin-state recovery. The same framework could be applied to other objects whose geometries are at least partially known — for instance, defunct LEO satellites with available technical documentation or CAD-derived shape models. For targets with unknown geometry, the brightness maps themselves encode shape information that could, in principle, be inverted to recover approximate three-dimensional form. Combining this with the multi-colour capability of STING would add a spectral dimension to the mapping, offering constraints on material composition alongside geometry. In particular, there could be an extension of this methodology to the Haiyang series of satellites that have a periodically scanning instrument on-board.

Coordination with international networks. The geometric spin-axis inversion demonstrated for Gorizont-6 relied on simultaneous observations from two sites (La Palma and Paranal). Extending this to a network of three or more geographically separated stations would tighten the constraints on spin-pole orientation and reduce the ambiguity inherent in two-site solutions. Several groups worldwide are developing or operating wide-field instruments for space domain awareness, and the standardisation of data formats and reduction pipelines would facilitate the construction of a coordinated observing network. Such a network would also increase the total available observing time and provide redundancy against weather losses, enabling the kind of sustained, multi-epoch monitoring that the Gorizont programme requires.

The mega-constellation era. The rapid deployment of large satellite constellations in LEO is fundamentally altering the orbital environment and, with it, the demands placed on space surveillance systems. The sheer number of objects — combined with frequent orbital manoeuvres, batch deployments and de-orbit operations — creates a need for automated, high-throughput characterisation pipelines capable of processing large volumes of photometric data with minimal human intervention. The wide-field, multi-colour, high-cadence architecture of STING is inherently suited

to this task, and the colour-based classification techniques explored in this thesis provide a foundation on which such pipelines could be built. At the same time, the growing population of defunct mega-constellation satellites in graveyard or disposal orbits will itself become a target for the kind of long-term rotational monitoring described here.

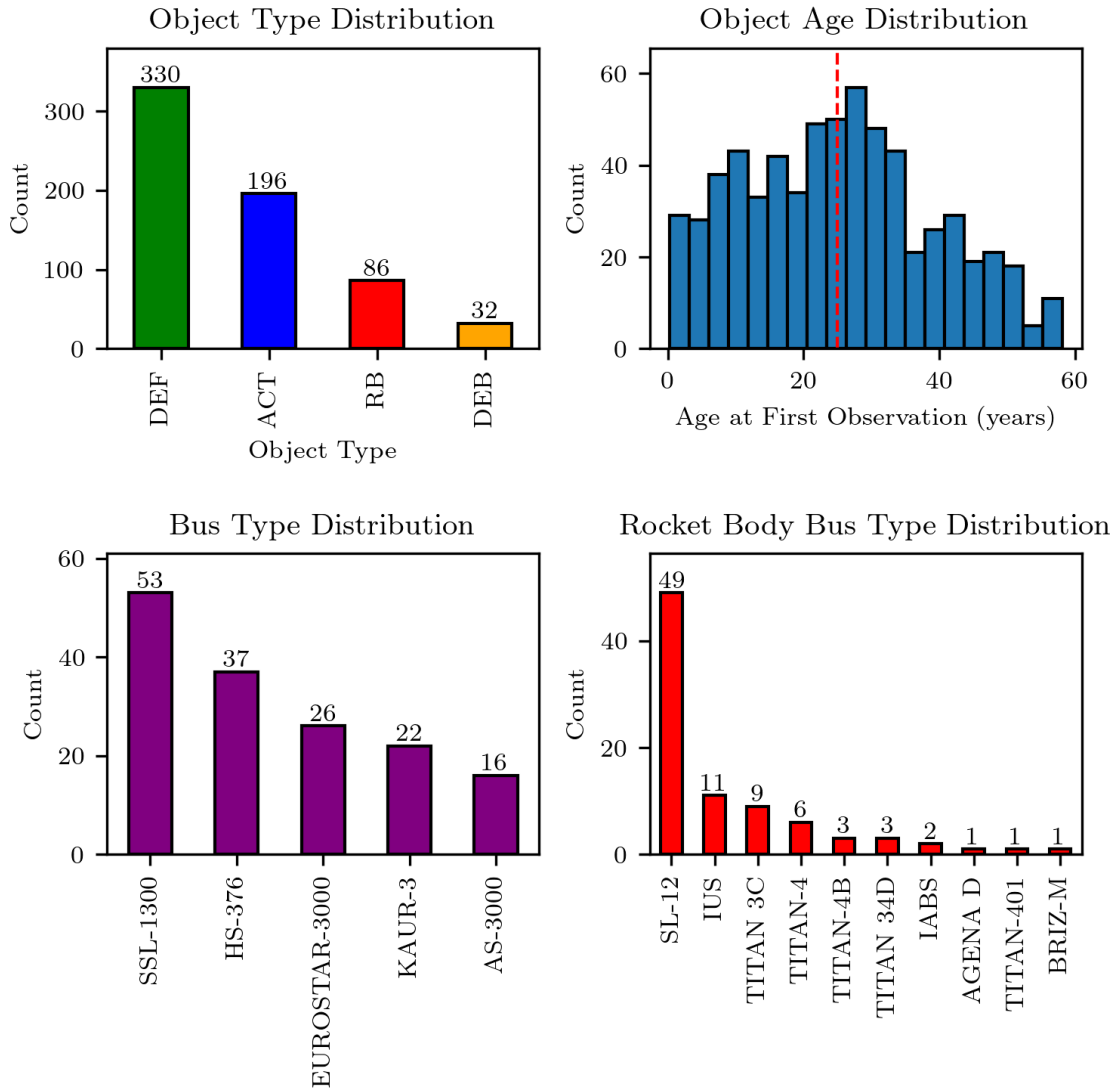

Figure 7.1: Demographics of unique targets observed with STING. **Upper Left**: Count of object type between defunct (DEF), active (ACT), Rocket Bodies (RB) and Debris (DEB). **Upper Right**: Distribution of the age of observed targets when they were first observed with STING. These are calculated using the launch dates provided within the master satellite list from McDowell (2024). The red dashed line refers to the median age of this dataset which sits at ~ 25 yr. **Lower Left**: Count of the 5 most populous bus types between active and defunct object types. **Lower Right**: Breakdown of different types of Rocket Bodies observed with STING.

Database of STING RSO observations.

Since SuperWASP-N was upgraded to STING, the amount of data collected on individual targets as part of larger surveys or follow up observations has grown. It is therefore vital to have an easily accessible look-up table for object type/count and light curve viewing.

A database was created with DASH (Parmer et al., 2025) to achieve the above. At time of writing the total number of unique light curves acquired with STING stands at **7629** (644 unique targets). Figure 7.1 displays some of the demographics of the unique objects observed.

In summary, the instruments, datasets and methods developed in this thesis provide a platform from which several of these directions can be pursued. The combination of wide-field multi-colour photometry with targeted polarimetry, supported by BRDF-forward modelling and machine-learning classification, offers a physically grounded and operationally scalable route to characterising the increasingly crowded near-Earth environment, and does so within a performance envelope that has been measured rather than assumed.

Bibliography

Herman Ader, G.J. Mellenbergh, and David Hand. Advising on research methods: A consultant companion. *Johannes Van Kessel Publ*, 01 2008.

Robert J.S. Airey, Paul Chote, James A. Blake, Benjamin F. Cooke, James McCormac, Phineas Allen, Alex MacManus, Don Pollacco, Billy Shrive, and Richard West. A comprehensive survey of the geo-belt using simultaneous four-colour observations with sting. *Advances in Space Research*, 75(7):5757–5780, April 2025a. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2025.02.027. URL https://www.sciencedirect.com/science/article/pii/S0273117725001474.

Robert J.S. Airey, Paul Chote, James A. Blake, James McCormac, Billy Shrive, Don Pollacco, and Benjamin F. Cooke. Lord of the (sub-)Rings: mapping the surface reflectance and spin-axis of Ajisai. *RAS Techniques and Instruments*, 4:rzaf058, November 2025b. doi: 10.1093/rasti/rzaf058.

Robert J.S. Airey, Paul Chote, Klaas Wiersema, Ioannis Apergis, James McCormac, James A. Blake, Benjamin F. Cooke, Isobel S. Lockley, Peter J. Wheatley, Daniel Bayliss, Samuel Gill, Christopher A. Watson, Stefano Covino, Frans Snik, Jon Marchant, Justyn Maund, Brooke Simmons, and Iain A. Steele. Brightness, colour, polarisation: a multi-instrument observation campaign of the gorizont-6 satellite. *Advances in Space Research*, 2026. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2026.03.106. URL https://www.sciencedirect.com/science/article/pii/S0273117726004424.

Miguel R. Alarcon, Javier Licandro, Miquel Serra-Ricart, Enrique Joven, Vicens Gaitan, and Rebeca de Sousa. Scientific CMOS Sensors in Astronomy: IMX455 and IMX411. *Publications of the Astronomical Society of the Pacific*, 135(1047):055001, May 2023. doi: 10.1088/1538-3873/acd04a.

Antonella A. Albuja. *Rotational Dynamics of Inactive Satellites as a Result of the YORP Effect*. PhD thesis, University of Colorado, Boulder, January 2015.

Antonella A. Albuja, Daniel J. Scheeres, Rita L. Cognion, William Ryan, and Eileen V. Ryan. The YORP effect on the GOES 8 and GOES 10 satellites: A case study. *Advances in Space Research*, 61(1):122–144, January 2018. doi: 10.1016/j.asr.2017.10.002.

American Association of Variable Star Observers. *AAVSO Guide to CCD/CMOS Photometry with Monochrome Cameras*, 2022. URL https://www.aavso.org/sites/default/files/publications_files/ccd_photometry_guide/CCDPhotometryGuide.pdf.

Paul V. Anderson, Darren S. McKnight, Frank Di Pentino, and Hanspeter Schaub. Operational considerations of geo debris synchronization dynamics. In *Proceedings of the 66th International Astronautical Congress*, 2015. URL https://api.semanticscholar.org/CorpusID:89613560.

Andor Technology / Oxford Instruments. Sensitivity and noise of ccd, emccd and scmos sensors, 2023. URL https://andor.oxinst.com/learning/view/article/sensitivity-and-noise-of-ccd-emccd-and-scmos-sensors. Accessed 12 Dec 2025.

Andor Technology / Oxford Instruments. scmos vs cmos cameras: What is the difference? https://andor.oxinst.com/learning/view/article/scmos-vs-cmos-cameras-what-is-the-difference, 2025. Accessed 12 Dec 2025.

Phillip D. Anz-Meador, John N. Opiela, Debra Shoots, and J.-C. Liou. History of On-Orbit Satellite Fragmentations, 15th Edition. NASA Technical Memorandum, NASA/STI Accession number: 20180008451, July 2018.

Ioannis Apergis, Daniel Bayliss, Leonidas Asimakoulas, Paul Chote, James McCormac, Morgan A Mitchell, Sam Gill, Philip G Steen, and Peter J Wheatley. High-precision photometry with a scientific cmos camera: I lab testing of the marana camera. *RAS Techniques and Instruments*, 4: rzaf049, 10 2025. ISSN 2752-8200. doi: 10.1093/rasti/rzaf049. URL https://doi.org/10.1093/rasti/rzaf049.

Ioannis Apergis, Daniel Bayliss, Paul Chote, James McCormac, Peter J Wheatley, Morgan A Mitchell, Jorge Fernández Fernández, Sam Gill, Edward M Bryant, Toby Rodel, Leonidas Asimakoulas, David R Anderson, James A Blake, Sarah L Casewell, Fintan Eeles-Nolle, Faith Hawthorn, James S Jenkins, Monika Lendl, Isobel S Lockley, Maximiliano Moyano, Sean M O'Brien, Suman Saha, Alexis M S Smith, Philip G Steen, Jose I Vines, Richard G West, and Tafadzwa Zivave. High-precision photometry with a scientific cmos camera: Ii on-sky testing of the marana camera at the ngts facility. *RAS Techniques and Instruments*, 5:rzag022, 01 2026. ISSN 2752-8200. doi: 10.1093/rasti/rzag022. URL https://doi.org/10.1093/rasti/rzag022.

George Arfken and Albert A. Mullin. Mathematical Methods for Physicists. *American Journal of Physics*, 35(11):1097–1098, November 1967. doi: 10.1119/1.1973757.

Astropy Collaboration. AltAz (astropy.coordinates) — api reference. Astropy Documentation (online), a. URL https://docs.astropy.org/en/stable/api/astropy.coordinates.AltAz.html.

Astropy Collaboration. TETE (astropy.coordinates) — api reference. Astropy Documentation (online), b. URL https://docs.astropy.org/en/stable/api/astropy.coordinates.TETE.html.

Astropy Collaboration. Time (astropy.time) — api reference. Astropy Documentation (online), c. URL https://docs.astropy.org/en/stable/api/astropy.time.Time.html.

Astroscale Ltd. Elsa-d: End-of-life services by astroscale demonstration. Technical report, Astroscale, 2022. URL https://astroscale.com/elsa-d/. Accessed: 2026-01-19.

Eijiro Atarashi, Hiroki Inoue, and Genki Fujii. Ultra-close RPO on-orbit demonstration of the ADRAS-J program. In *75th International Astronautical Congress (IAC)*, number IAC-24,A6,6,9,x87881, 2024.

Kyle Barbary. SEP: Source Extractor as a library. *The Journal of Open Source Software*, 1(6):58, October 2016. doi: 10.21105/joss.00058.

Mr. Bayes and Mr. Price. An Essay towards Solving a Problem in the Doctrine of Chances. By the Late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a Letter to John Canton, A. M. F. R. S. *Philosophical Transactions of the Royal Society of London Series I*, 53:370–418, January 1763.

Petr Beckmann and Andre Spizzichino. *The scattering of electromagnetic waves from rough surfaces*. Artech House, 1987.

D. Bedard and G. Wade. Interpretation of Spectrometric Measurements of Active Geostationary Satellites. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E31, September 2014.

C. J Benson. *Solar Torque and Dissipation Dynamics for Tumbling Bodies: Theory and Observations*. PhD thesis, University of Colorado, 2021.

C. J. Benson and D. J. Scheeres. The yorp effect for tumbling defunct geo satellites. In *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Portland, ME, 2019. AAS 19-858.

E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *Astronomy & Astrophysics Supplement Series*, 117:393–404, June 1996. doi: 10.1051/aas:1996164.

C. Binz, M. Davis, B. Kelm, and C. Moore. Optical Survey of the Tumble Rates of Retired GEO Satellites. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E61, September 2014.

J. A. Blake, D. Pollacco, R. Airey, P. Chote, B. F. Cooke, J. McCormac, B. Shrive, R. G. West, J. Davis, W. Feline, et al. Exploring SDA sensor architectures for the surveillance of geosynchronous spacecraft. In *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2023.

James A Blake. Looking out for a sustainable space. *Astronomy & Geophysics*, 63(2):2.14–2.20, 4 2022. doi: 10.1093/astrogeo/atac022. URL https://doi.org/10.1093%2Fastrogeo%2Fatac022.

James A. Blake, Paul Chote, Don Pollacco, William Feline, Grant Privett, Andrew Ash, Stuart Eves, Arthur Greenwood, Nick Harwood, Thomas R. Marsh, Dimitri Veras, and Christopher Watson. DebrisWatch I: A survey of faint geosynchronous debris. *Advances in Space Research*, 67(1):360–370, January 2021. doi: 10.1016/j.asr.2020.08.008.

Alexander Bolonkin. New methods of removing space debris, March 2014. URL https://www.gsjournal.net/Science-Journals/Research%20Papers-Engineering%20(Applied)/Download/5382. General Science Journal; document dated 11 March 2014.

Christophe Bonnal. Design and operational practices for the passivation of spacecraft and launchers at the end of life. *Proceedings of The Institution of Mechanical Engineers Part G-journal of Aerospace Engineering - PROC INST MECH ENG G-J A E*, 221:925–931, 12 2007. doi: 10.1243/09544100JAERO231.

Achira Boonrath, Tarunraj Singh, and Eleonora M. Botta. Identification of parameters for tethered satellite system to emulate net-captured debris towing. *Acta Astronautica*, 225:676–688, December 2024. doi: 10.1016/j.actaastro.2024.09.022.

Ben K. Bradley and Penina Axelrad. Lightcurve inversion for shape estimation of geo objects from space-based sensors. In *Proceedings of the 24th International Symposium on Space Flight Dynamics*, pages 1–20, Laurel, MD, USA, May 2014. URL https://issfd.org/ISSFD_2014/ISSFD24_Paper_S10-3_Bradley.pdf.

Jeimmy Nataly Buitrago-Leiva, Mohamed El Khayati Ramouz, Adriano Camps, and Joan A. Ruiz-de-Azua. Towards a second life for Zombie Satellites: Anomaly occurrence and potential recycling assessment. *Acta Astronautica*, 217:238–245, April 2024. doi: 10.1016/j.actaastro.2024.01.048.

Bureau International des Poids et Mesures. International atomic time, 2011. URL https://web.archive.org/web/20110110113226/http://www.bipm.org/en/scientific/tai/tai.html.

Grzegorz Bury, Krzysztof Sośnica, Radosław Zajdel, and Dariusz Strugarek. Toward the 1-cm Galileo orbits: challenges in modeling of perturbing forces. *Journal of Geodesy*, 94(2):16, February 2020. doi: 10.1007/s00190-020-01342-2.

Carlo Calatroni, Gilles Métris, Clément Courde, Duy-Hà Phung, Julien Chabé, Mourad Aimar, Nicolas Maurice, and Hervé Mariey. A new model to predict ajisai satellite reflected sunlight flashes and application to the determination of its rotation parameters. *Advances in Space Research*, 2025. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2025.02.054. URL https://www.sciencedirect.com/science/article/pii/S0273117725001899.

Michel Capderou. *Satellites: Orbits and missions*. Springer Science & Business Media, 2005.

P. J. Castro, T.E. Payne, L.W Weisenbach, T.J Godar, V. L. Wiley, J. Frith, and S.P. Milster. Analysis of Induced Color Index Error Due to Sequential Filter Photometry. In *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2022.

M. Cegarra Polo, A. Alenin, I. Vaughn, and A. Lambert. GEO Satellite Characterization through Polarimetry using Simultaneous Observations from nearby Optical Sensors. In Sandy Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page 23, September 2016.

Delphine Cerutti-Maori, Claudio Carloni, Vassilis Karamanavis, Felix Rosebrock, Frank Schlichthaber, and Jeremias Schroer. Observation of the fragmentation cloud of the intelsat 33e breakup with the space observation radar tira. In *Proceedings of the 9th European Conference on Space Debris (SDC9)*. ESA Space Debris Office, 2025. URL https:

//conference.sdo.esoc.esa.int/proceedings/sdc9/paper/288/SDC9-paper288.pdf. Paper 288, ESA SDO Conference Proceedings.

P. Charlot, C. S. Jacobs, D. Gordon, S. Lambert, A. de Witt, J. Böhm, A. L. Fey, R. Heinkelmann, E. Skurikhina, O. Titov, E. F. Arias, S. Bolotin, G. Bourda, C. Ma, Z. Malkin, A. Nothnagel, D. Mayer, D. S. MacMillan, T. Nilsson, and R. Gaume. The third realization of the International Celestial Reference Frame by very long baseline interferometry. *Astronomy & Astrophysics*, 644:A159, December 2020. doi: 10.1051/0004-6361/202038368.

Paul Chote, James A Blake, and Don Pollacco. Precision Optical Light Curves of LEO and GEO Objects. In S. Ryan, editor, *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 52, September 2019.

Paul Chote, Robert Airey, James McCormac, Don Pollacco, Richard West, Krzysztof Ulaczyk, Martin J. Dyer, Alexander Agathanggelou, William Feline, Simon George, Calum Meredith, and Grant Privett. High-precision light curves of geostationary objects: The PHANTOM ECHOES 2 RPO campaign. *Advances in Space Research*, 76(2):764–783, July 2025. doi: 10.1016/j.asr.2025.05.026.

Arthur C. Clark. Extra-terrestrial relays: Can rocket stations give world-wide radio coverage? In Richard B. Marsten, editor, *Communication Satellite Systems Technology*, volume 19 of *Progress in Astronautics and Rocketry*, pages 3–6. Elsevier, 1966. doi: https://doi.org/10.1016/B978-1-4832-2716-0.50006-2. URL https://www.sciencedirect.com/science/article/pii/B9781483227160500062.

Rita L. Cognion. Large phase angle observations of GEO satellites. In Khanh D. Pham, Joseph L. Cox, Richard T. Howard, and Genshe Chen, editors, *Sensors and Systems for Space Applications VI*, volume 8739, page 87390K. International Society for Optics and Photonics, SPIE, 2013. doi: 10.1117/12.2014623. URL https://doi.org/10.1117/12.2014623.

M. Compean and T. Small. Modified Solar Cell BRDF. In S. Ryan, editor, *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 42, September 2023.

R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1):7–24, January 1982. ISSN 0730-0301. doi: 10.1145/357290.357293. URL https://doi.org/10.1145/357290.357293.

Benjamin F. Cooke, Paul Chote, Don Pollacco, Richard West, James A. Blake, James McCormac, Robert Airey, and Billy Shrive. Simulated recovery of LEO objects using sCMOS blind stacking. *Advances in Space Research*, 72(4):907–921, August 2023. doi: 10.1016/j.asr.2023.05.003.

James Cooley and John Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.

H. Cowardin, P. Anz-Meador, J. Melo, M. Castaneda, C. Cruz, J. Opiela, C. Cline, and E. Christiansen. Analysis of Darkened Fragments Resulting from Laboratory Hypervelocity Experiments. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 1, September 2024a.

Heather Cowardin, Eric Christiansen, Mark Matney, Joshua Miller, Bruce (Alan) Davis, Corbin Cruz, John Seago, Austen King, John Opiela, and Alyssa Manis. Orbital Debris Shape Effect Investigations for Mitigating Risk. *Journal of the Astronautical Sciences*, 71(6):49, November 2024b. doi: 10.1007/s40295-024-00471-6.

Alessandra Di Cecco, Alberto Buzzoni, Maddalena Mochi, Giuseppe Altavilla, Sara Brambilla, Marco Castronuovo, Fiore De Luise, Silvia Galleti, José Guichard, Cosimo Marzo, and Gaetano Valentini. Geo Satellites' Manufacturing Markers from Photometric Colors. In *2nd NEO and Debris Detection Conference*, page 64, January 2023.

Peter Duffett-Smith and Jonathan Zwart. Time scales: Ut1, utc, tai, tt, gps time, 2011. URL `https://stjarnhimlen.se/comp/time.html`. Accessed: 2026-01-19.

Anita N. Dunsmore, Joshua A. Key, Ryan M. Tucker, Evan M. Weld, Francis K. Chun, and Roger D. Tippets. Spectral measurements of geosynchronous satellites during glint season. *Journal of Spacecraft and Rockets*, 54(2):349–355, 2017. doi: 10.2514/1.A33583. URL `https://doi.org/10.2514/1.A33583`.

Michael A. Earl, Philip W. Somers, Konstantin Kabin, Donald Bédard, and Gregg A. Wade. Estimating the spin axis orientation of the Echostar-2 box-wing geosynchronous satellite. *Advances in Space Research*, 61(8):2135–2146, April 2018. doi: 10.1016/j.asr.2018.01.032.

A. Einstein. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 322(6):132–148, January 1905. doi: 10.1002/andp.19053220607.

William Emery and Adriano Camps. Chapter 7 - orbital mechanics, image navigation, and cartographic projections. In William Emery and Adriano Camps, editors, *Introduction to Satellite Remote Sensing*, pages 565–596. Elsevier, 2017. ISBN 978-0-12-809254-5. doi: https://doi.org/10.1016/B978-0-12-809254-5.00007-5. URL `https://www.sciencedirect.com/science/article/pii/B9780128092545000075`.

T. Endo, T. Tsuchikawa, T. Anada, H. Ono, and H. Tsuji. Simulating the Photometric Light Curve of Artificial Satellites in GEO used with a Ray-Tracing. In S. Ryan, editor, *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 94, September 2023.

Yuta Endo, Hirohisa Kojima, and Pavel M. Trivailo. New formulation for evaluating status of space debris capture using tether-net. *Advances in Space Research*, 70(10):2976–3002, November 2022. doi: 10.1016/j.asr.2022.09.024.

eoPortal – Earth Observation Portal. Egs (experimental geodetic satellite) / ajisai. https://www.eoportal.org/satellite-missions/egs, 2012. Accessed: 2025-07-05.

ESA Space Debris Office. ESA's annual space environment report. Technical Report Issue 10, European Space Agency, European Space Operations Centre, Darmstadt, Germany, 5 2026. URL https://www.sdo.esoc.esa.int/publications/Space_Environment_Report_I10R0_20260501.pdf.

EU SST Partnership. EU space surveillance and tracking. https://www.eusst.eu/. Accessed: 2026-08-15.

European Space Agency. Mechanical Platform Improvement Plan. https://connectivity.esa.int/projects/mechanical-platform-improvement-plan, 2014. Retrieved 2024-05-13.

European Space Agency. Clearspace-1: First active debris removal mission. Technical report, European Space Agency, 2021. URL https://www.esa.int/Space_Safety/Clean_Space/ClearSpace-1. Accessed: 2026-01-19.

European Space Agency. Esa space environment report. Technical report, European Space Agency, 2023a. URL https://www.esa.int/Space_Safety/Space_Debris. Accessed: 2026-01-19.

European Space Agency. Zero debris charter. Technical report, European Space Agency, 2023b. URL https://www.esa.int/Space_Safety/Zero_Debris. Accessed: 2026-01-19.

European Space Agency. Esa's annual space environment report. Technical report, ESA ESOC, 2025a.

European Space Agency. Space environment statistics. https://sdup.esoc.esa.int/discosweb/statistics/, 2025b. Last accessed: 2025-12-04.

European Space Agency (ESA), ESOC. Esa's fragmentation database: Modelling. https://fragmentation.esoc.esa.int/home/modelling. Accessed: 2025-10-28.

Forrest Fankhauser, J. Anthony Tyson, and Jacob Askari. Satellite Optical Brightness. *Astronomical Journal*, 166(2):59, August 2023. doi: 10.3847/1538-3881/ace047.

Federal Communications Commission. Mitigation of orbital debris in the new space age. Technical report, Federal Communications Commission, 2022. URL https://www.fcc.gov/document/fcc-adopts-5-year-orbital-debris-rule. FCC 22-74; Accessed: 2026-01-19.

Y. R. Fernández, S. C. Lowry, P. R. Weissman, B. E. A. Mueller, N. H. Samarasinha, M. J. S. Belton, and K. J. Meech. New near-aphelion light curves of Comet 2P/Encke. *Icarus*, 175(1): 194–214, May 2005. doi: 10.1016/j.icarus.2004.10.019.

United States Space Force, 2022. URL https://digitalcommons.usu.edu/smallsat/2022/all2022/124/. Last accessed 14 April 2025.

Daniel Foreman-Mackey. corner.py: Scatterplot matrices in python. *The Journal of Open Source Software*, 1(2):24, 6 2016. doi: 10.21105/joss.00024. URL https://doi.org/10.21105/joss.00024.

Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306, March 2013. doi: 10.1086/670067.

Carolin Früh and Thomas Schildknecht. Accuracy of two-line-element data for geostationary and high-eccentricity orbits. *Journal of Guidance, Control, and Dynamics*, 35(5):1483–1491, 2012. doi: 10.2514/1.55843. URL https://doi.org/10.2514/1.55843.

Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Akti Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC Texts in Statistical Science Series. CRC, Boca Raton, Florida, third edition, 2013. URL https://stat.columbia.edu/~gelman/book/.

Doug George. Scientific imaging camera performance. https://diffractionlimited.com/scientific-imaging-camera-performance/, June 2024. Blog post. Published 25 June 2024, last updated 12 August 2025. Accessed 27 August 2026.

Herbert Goldstein. *Classical mechanics*. Addison-Wesley, second edition, 1980.

Mauricio M. Guelman. Geostationary satellites autonomous closed loop station keeping. *Acta Astronautica*, 97:9–15, 2014. ISSN 0094-5765. doi: https://doi.org/10.1016/j.actaastro.2013.12.009. URL https://www.sciencedirect.com/science/article/pii/S0094576513004621.

D. Hall and P. Kervin. Analysis of Faint Glints from Stabilized GEO Satellites. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E34, September 2013.

D. Hall, J. Africano, D. Archambeault, B. Birge, D. Witte, and P. Kervin. AMOS Observations of NASA's IMAGE Satellite. In S. Ryan, editor, *The Advanced Maui Optical and Space Surveillance Technologies Conference*, page E76, January 2006.

Bruce Hapke. *Theory of Reflectance and Emittance Spectroscopy*. Cambridge University Press, 2012. doi: 10.1017/CBO9781139025683.

Akihisa Hattori and Toshimichi Otsubo. Time-varying solar radiation pressure on Ajisai in comparison with LAGEOS satellites. *Advances in Space Research*, 63(1):63–72, January 2019. doi: 10.1016/j.asr.2018.08.010.

Arne A. Henden, Stephen Levine, Dirk Terrell, and Douglas L. Welch. APASS - The Latest Data Release. In *American Astronomical Society Meeting Abstracts #225*, volume 225 of *American Astronomical Society Meeting Abstracts*, page 336.16, January 2015.

Caleb Henry. Intelsat-29e satellite suffers fuel leak, spotted drifting along GEO arc. *SpaceNews*, April 2019. URL https://spacenews.com/intelsat-29e-satellite-suffers-fuel-leak-spotted-drifting-along-geo-arc/. Accessed: 2026-04-24.

Steve Hilla. The extended standard product 3 orbit format (sp3-d). Technical report, National Geodetic Survey, National Ocean Service, NOAA, Silver Spring, MD, USA, February 2016. URL https://gssc.esa.int/wp-content/uploads/2018/07/sp3d.pdf.

Steve B. Howell. Two-Dimensional Aperture Photometry: Signal-to-Noise Ratio of Point-Source Observations and Optimal Data-Extraction Techniques. *Publications of the Astronomical Society of the Pacific*, 101:616, June 1989. doi: 10.1086/132477.

Steve B. Howell. *Handbook of CCD Astronomy*. Cambridge Observing Handbooks for Research Astronomers. Cambridge University Press, 2 edition, 2006.

F. C. Hoyt. Stationary earth orbits. *AIAA Journal*, 2(8):1470–1471, August 1964. doi: 10.2514/3.2609.

IAU SOFA Board. *IAU SOFA Software Collection*. International Astronomical Union. https://www.iausofa.org.

IERS Conventions. Iers conventions (2010). Technical report, International Earth Rotation and Reference Systems Service, 2010.

Inter-Agency Space Debris Coordination Committee. Iadc space debris mitigation guidelines. Technical report, IADC, 2007. URL https://www.iadc-home.org/documents/. Revision 1; Accessed: 2026-01-19.

International Astronomical Union. Resolution b2 of the XXIIIrd general assembly: On the International Celestial Reference System (ICRS). Kyoto, Japan, 1997. URL https://www.iers.org/SharedDocs/Publikationen/EN/IERS/resolutions/IAU/IAU1997_Resolutions.pdf?__blob=publicationFile&v=3.

International Astronomical Union. Resolution B1.3 of the XXIVth general assembly: Definition of Barycentric Celestial Reference System and Geocentric Celestial Reference System. Manchester, UK, 2000. URL https://www.iers.org/SharedDocs/Publikationen/EN/IERS/resolutions/IAU/IAU2000_Resolutions_B1-B3.pdf?__blob=publicationFile&v=3.

International Astronomical Union. Impact of satellite constellations on astronomy. Technical report, IAU, 2022. URL https://www.iau.org/science/scientific_bodies/working_groups/344/. Accessed: 2026-01-19.

International Earth Rotation and Reference Systems Service (IERS). Iers technical note 36: Chapter 5 — transformation between the ITRS and the GCRS. Chapter PDF (online), 2010. URL https://www.iers.org/SharedDocs/Publikationen/EN/IERS/Publications/tn/TechnNote36/tn36_043.

International Telecommunication Union. Radio regulations and the geostationary orbit. Technical report, ITU, 2020. URL https://www.itu.int/en/ITU-R/space. Accessed: 2026-01-19.

James R. Janesick. *Scientific charge-coupled devices*. SPIE Press, 2001.

James R. Janesick. *Photon Transfer DN → λ*. SPIE Press, 2007.

Helen Jermak, Iain A. Steele, and Robert J. Smith. MOPTOP: a multi-colour optimised optical polarimeter. In Christopher J. Evans, Luc Simard, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy VI*, volume 9908 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99084I, August 2016. doi: 10.1117/12.2232324.

Helen Jermak, Iain A. Steele, and Robert J. Smith. Optical design of the Liverpool Telescope Multicolour OPTimised Optical Polarimeter (MOPTOP). In Christopher J. Evans, Luc Simard, and Hideki Takami, editors, *Ground-based and Airborne Instrumentation for Astronomy VII*, volume 10702 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107024Q, July 2018. doi: 10.1117/12.2312132.

A. Jolley, G. Wade, and D. Bedard. Multicolour Optical Photometry of Active Geostationary Satellites. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, page 74, January 2015.

Harry W. Jones. The recent large reduction in space launch cost. In *48th International Conference on Environmental Systems (ICES)*, Albuquerque, New Mexico, 2018. ICES-2018-81.

Sergey Karpov, Armelle Bajat, Asen Christov, Michael Prouza, and Grigory Beskin. Evaluation of scientific CMOS sensors for sky survey applications. In Andrew D. Holland and James Beletic, editors, *X-Ray, Optical, and Infrared Detectors for Astronomy IX*, volume 11454, page 114540G. International Society for Optics and Photonics, SPIE, 2020. doi: 10.1117/12.2561834. URL https://doi.org/10.1117/12.2561834.

T. S. Kelso. Celestrak: Current gp element sets. https://celestrak.org/NORAD/elements/, 2026. Accessed: 15 March 2026.

J. C. Kemp, G. D. Henson, C. T. Steiner, and E. R. Powell. The optical polarization of the Sun measured at a sensitivity of parts in ten million. *Nature*, 326(6110):270–273, March 1987. doi: 10.1038/326270a0.

Emma Kerr, Gabriele Falco, Nina Maric, David Petit, Patrick Talon, Elisabeth Geistere Petersen, Chris Dorn, Stuart Eves, Noelia Sánchez-Ortiz, Raul Dominguez Gonzalez, and Jaime Nomen-Torres. Light Curves for GEO object characterisation. In *8th European Conference on Space Debris*, page 49, April 2021.

Donald J. Kessler and Burton G. Cour-Palais. Collision frequency of artificial satellites: The creation of a debris belt. *Journal of Geophysical Research: Space Physics*, 83(A6):2637–2646, 1978. doi: https://doi.org/10.1029/JA083iA06p02637. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JA083iA06p02637.

Georg Kirchner, Walter Hausleitner, and Elena Cristea. Ajisai spin parameter determination using graz kilohertz satellite laser ranging data. *Geoscience and Remote Sensing, IEEE Transactions on*, 45:201 – 205, 02 2007. doi: 10.1109/TGRS.2006.882254.

N. Kosaka, Y. Itoh, T. Saito, M. Tozuka, T. Endo, and T. Ando. Imaging Polarimetry of Geostationary Satellite Express-AM5. *Odessa Astronomical Publications*, 33:115, January 2020. doi: 10.18524/1810-4215.2020.33.207036.

N. Koshkin, L. Shakun, N. Burlak, E. Korobeynikova, S. Strakhova, S. Melikyants, S. Terpan, and A. Ryabov. Ajisai spin-axis precession and rotation-period variations from photometric observations. *Advances in Space Research*, 60(7):1389–1399, October 2017. doi: 10.1016/j.asr.2017.06.045.

N. Koshkin, L. Shakun, S. Melikyants, E. Korobeynikova, S. Strakhova, A. Bilinsky, Ye. Vovchyk, V. Kudak, I. Motrunich, I. Neubauer, O. Kozhukhov, Ya. Romanyuk, A. Ryabov, S. Terpan, V. Dragomiretsky, and T. Golubovskaya. Simulatneous Multi-Site Photometry of Leo Satellites to Characterise Their Rotation States. *Odessa Astronomical Publications*, 33:119, January 2020. doi: 10.18524/1810-4215.2020.33.216464.

Nikolay Koshkin, Leonid Shakun, Elena Korobeynikova, Seda Melikyants, Svetlana Strakhova, Vladislav Dragomiretsky, Andrey Ryabov, Sergey Terpan, and Tatiana Golubovskaya. Determination of the spacecraft's spin axis orientation. Photometric patterns method. *Advances in Space Research*, 74(11):5725–5744, December 2024. doi: 10.1016/j.asr.2024.08.038.

Vladimir Kouprianov. Distinguishing features of CCD astrometry of faint GEO objects. *Advances in Space Research*, 41(7):1029–1038, January 2008. doi: 10.1016/j.asr.2007.04.033.

E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley, 2006. ISBN 9780471728979. URL `https://books.google.co.uk/books?id=3t7DQgAACAAJ`.

Kevin Krisciunas and Bradley E. Schaefer. A Model of the Brightness of Moonlight. *Publications of the Astronomical Society of the Pacific*, 103:1033, September 1991. doi: 10.1086/132921.

D. Kucharski, G. Kirchner, S. K. Flegel, M. K. Jah, and J. C. Bennett. High-Definition Photometry — An Optical Method for Remote Detection of MMOD Impacts on Satellites. In *First International Orbital Debris Conference*, volume 2109 of *LPI Contributions*, page 6075, December 2019a.

D. Kucharski, Georg Kirchner, Toshimichi Otsubo, Hiroo Kunimori, Moriba K. Jah, Franz Koidl, James C. Bennett, Hyung-Chul Lim, Peiyuan Wang, Michael Steindorfer, and Krzysztof Sośnica. Hypertemporal photometric measurement of spaceborne mirrors specular reflectivity for Laser Time Transfer link model. *Advances in Space Research*, 64(4):957–963, aug 2019b. doi: 10.1016/j.asr.2019.05.030.

Daniel Kucharski, Georg Kirchner, Toshimichi Otsubo, Hyung-Chul Lim, James Bennett, Franz Koidl, Young-Rok Kim, and Joo-Yeon Hwang. Confirmation of gravitationally induced attitude drift of spinning satellite Ajisai with Graz high repetition rate SLR data. *Advances in Space Research*, 57(4):983–990, February 2016. doi: 10.1016/j.asr.2015.12.010.

Daniel Kucharski, Georg Kirchner, Toshimichi Otsubo, Sven K. Flegel, Hiroo Kunimori, Moriba K. Jah, Franz Koidl, James C. Bennett, Michael Steindorfer, and Peiyuan Wang. Quanta Photogrammetry of Experimental Geodetic Satellite for remote detection of micrometeoroid and orbital

debris impacts. *Acta Astronautica*, 174:24–31, September 2020. doi: 10.1016/j.actaastro.2020.04.042.

J.B. Kuipers. *Quaternions and Rotation Sequences: A Primer with Applications to Orbits, Aerospace, and Virtual Reality*. Princeton paperbacks. Princeton University Press, 1999. ISBN 9780691102986. URL https://books.google.co.uk/books?id=_2sS4mC0p-EC.

Johann Heinrich Lambert. *Photometria, sive de mensura et gradibus luminis, colorum et umbrae*. Eberhard Klett, Augsburg, 1760.

M. H. Lane. The development of an artificial satellite theory using power-law atmospheric density representation. In *2nd Aerospace Sciences Meeting*, 1965. doi: 10.2514/6.1965-35. URL https://arc.aiaa.org/doi/abs/10.2514/6.1965-35.

Max H. Lane and K. H. Cranford. An improved analytical drag theory for the artificial satellite problem. In *Astrodynamics Conference*, number AIAA 1969-925, Princeton, NJ, 1969. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.1969-925. URL https://arc.aiaa.org/doi/abs/10.2514/6.1969-925.

Dustin Lang. *Astrometry.net: Automatic recognition and calibration of astronomical images*. PhD thesis, University of Toronto, Canada, January 2009.

Peter M. Lee. *Bayesian Statistics: An Introduction*. Wiley Publishing, 4th edition, 2012. ISBN 1118332571.

Creon Levit and William Marshall. Improved orbit predictions using two-line elements. *Advances in Space Research*, 47(7):1107–1115, 2011. ISSN 0273-1177. doi: 10.1016/j.asr.2010.10.017.

Hugh G. Lewis, Arrun Saunders, Graham Swinerd, and Rebecca J. Newland. Effect of thermospheric contraction on remediation of the near-Earth space debris environment. *Journal of Geophysical Research (Space Physics)*, 116:A00H08, August 2011. doi: 10.1029/2011JA016482.

Richard Linares, Moriba K. Jah, John L. Crassidis, and Christopher K. Nebelecky. Space Object Shape Characterization and Tracking Using Light Curve and Angles Data. *Journal of Guidance Control Dynamics*, 37(1):13–25, January 2014. doi: 10.2514/1.62986.

N. R. Lomb. Least-Squares Frequency Analysis of Unequally Spaced Data. *Astrophysics and Space Science*, 39(2):447–462, February 1976. doi: 10.1007/BF00648343.

P. Longa-Peña, J. Tregloan-Reed, C. Adam, J. Chamoun-Contreras, E. Unda-Sanzana, C. Ávalos-Vega, F. Amadio, M. Andersen, M. Bonavita, V. Bozza, B. Campos Estrada, M. Dominik, A. Donaldson, R. Figuera Jaimes, J. Fynbo, T. C. Hinse, M. Hundertmark, U. G. Jørgensen, E. Khalouei, M. Kretlow, V. Molina, N. Peixinho, M. Rabus, S. Rahvar, A. R. Rajkumar, M. I. Romero-Colmenares, P. Rota, S. Sajadian, J. Skottfelt, C. Snodgrass, J. Southworth, H. Alarcon, G. Blanc, P. Bocaz, J. P. Colque, J. Cortes, C. Flores Quintana, P. Garcia, R. González, S. Kim, S. Martín, T. Nakos, E. Ortiz, A. Otarola, J. Plaza Hernández, P. Sanhueza, G. Siringo, and M. Soto. Brightness evolution of LEO Starlink mega-constellation satellites from 2021 to 2023:

a multiyear ground-based photometric study. *Monthly Notices of the Royal Astronomical Society*, 548(2):stag552, May 2026. doi: 10.1093/mnras/stag552.

Yao Lu. BRDF-Based Photometric Modeling of LEO Constellation Satellite from Massive Observations. *Universe*, 10(5):215, May 2024. doi: 10.3390/universe10050215.

Delphine Ly, Romain Lucken, and Damien Giolito. Correcting tles at epoch: Application to the gps constellation. *Journal of Space Safety Engineering*, 7(3):302–306, 2020. ISSN 2468-8967. doi: https://doi.org/10.1016/j.jsse.2020.07.032. URL https://www.sciencedirect.com/science/article/pii/S2468896720300963. Space Debris: The State of Art.

C. Ma, E. F. Arias, T. M. Eubanks, A. L. Fey, A.-M. Gontier, C. S. Jacobs, O. J. Sovers, B. A. Archinal, and P. Charlot. The International Celestial Reference Frame as Realized by Very Long Baseline Interferometry. *Astronomical Journal*, 116(1):516–546, July 1998. doi: 10.1086/300408.

Alexander Macmanus, Don Pollacco, Paul Chote, and Calum Meredith. Dual-band Characterisation of Low-Earth Orbit Satellites Using Machine Learning Techniques. In S. Ryan, editor, *The Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 23, September 2025.

Stephen Malys, John H. Seago, Nikolaos K. Pavlis, P. Kenneth Seidelmann, and George H. Kaplan. Why the Greenwich meridian moved. *Journal of Geodesy*, 89(12):1263–1272, December 2015. doi: 10.1007/s00190-015-0844-y.

A. Manis, C. Vavrin, et al. An overview of ground-based radar and optical measurements utilized by the nasa orbital debris program office. In *Proceedings of the Second International Orbital Debris Conference (IOC II)*, Sugar Land, Texas, USA, 2023. URL https://www.hou.usra.edu/meetings/orbitaldebris2023/pdf/6045.pdf. Full paper, paper number 6045. Available at https://www.hou.usra.edu/meetings/orbitaldebris2023/pdf/6045.pdf.

Mark Matney. Measuring small debris – what you can't see can hurt you. Technical Report 20160011226, NASA Johnson Space Center, Orbital Debris Program Office, 2016. URL https://ntrs.nasa.gov/api/citations/20160011226/downloads/20160011226.pdf. NASA Technical Reports Server (NTRS).

Dennis D. McCarthy and P. Kenneth Seidelmann. Time: From earth rotation to atomic physics. *Cambridge University Press*, 2008.

Dennis D. McCarthy, Christine Hackman, and Robert A. Nelson. The Physical Basis of the Leap Second. *Astronomical Journal*, 136(5):1906–1908, November 2008. doi: 10.1088/0004-6256/136/5/1906.

Gary A. McCue, James G. Williams, and Joan M. Morford. Optical characteristics of artificial satellites. *Planetary and Space Science*, 19(8):851–868, August 1971. doi: 10.1016/0032-0633(71)90137-1.

Jonathan McDowell. The united nations registry of space objects - notes. http://planet4589.org/space/un/un_notes.html, 2002. Accessed: 2025-06-29.

Jonathan C. McDowell. "Master Satellite List". Jonathan's Space Report., 2024. Retrieved 2024-05-13, from https://www.planet4589.org/space/gcat/data/cat/satcat.html.

Ian S. McLean. *Electronic Imaging in Astronomy: Detectors and Instrumentation (Second Edition)*. Springer, 2008. doi: 10.1007/978-3-540-76583-7.

C. Meredith, P. Chote, and R. Airey. Tracking Merged Objects within Non-Resolved Imagery. In S. Ryan, editor, *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 44, September 2023.

T. Meyer, D. Keil, D. Traub, S. Scharring, W. Riede, T. Dekorsy, M. Nussbaum, M. Lengowski, R. Schweigert, and S. Klinkner. High-Fidelity Light Curve Simulation and Validation Using Empirical Data. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 115, September 2024.

Ministry of Defence and UK Space Agency. Cross-government space domain awareness (SDA) requirements. UK Government policy paper, 2026. URL https://www.gov.uk/government/publications/space-domain-awareness-requirements. Version 3.0, updated 16 June 2026. Accessed [date].

Moutaman Mirghani, Hoyam Abobaker, Eshraga Adel, and Asia Habila. Orbit design of cube satellite. In *10th Scientific Conference of the National Center for Research*, Khartoum, Sudan, December 2015.

Morgan A. Mitchell, Paul Chote, James McCormac, Don Pollacco, Ioannis Apergis, J. D. Lyman, Isobel S. Lockley, Samuel Gill, James A. Blake, Alastair B. Claringbold, D. T. H. Steeghs, and J. Casares. A search for photometric variability towards the globular cluster M3 with the TWenty Inch Survey Telescope. *Monthly Notices of the Royal Astronomical Society*, November 2025. doi: 10.1093/mnras/staf2100.

Oliver Montenbruck and Gill Eberhard, editors. *Satellite orbits : models, methods, and applications*, January 2000.

Benoit Moulin. Fitting ABg scattering coefficients from raw measured data. ANSYS Optics Knowledge Base, 2021. URL https://optics.ansys.com/hc/en-us/articles/4307113242139 5-Fitting-ABg-scattering-coefficients-from-raw-measured-data. Retrieved November 4, 2025.

Sangeetha Nandakumar, Siegfried Eggl, Jeremy Tregloan-Reed, Christian Adam, Jasmine Anderson-Baldwin, Michele T. Bannister, Adam Battle, Zouhair Benkhaldoun, Tanner Campbell, J. P. Colque, Guillermo Damke, Ilse Plauchu Frayn, Mourad Ghachoui, Pedro F. Guillen, Aziz Ettahar Kaeouach, Harrison R. Krantz, Marco Langbroek, Nicholas Rattenbury, Vishnu Reddy, Ryan Ridden-Harper, Brad Young, Eduardo Unda-Sanzana, Alan M. Watson, Constance E. Walker,

John C. Barentine, Piero Benvenuti, Federico Di Vruno, Mike W. Peel, Meredith L. Rawls, Cees Bassa, Catalina Flores-Quintana, Pablo García, Sam Kim, Penélope Longa-Peña, Edgar Ortiz, Ángel Otarola, María Romero-Colmenares, Pedro Sanhueza, Giorgio Siringo, and Mario Soto. The high optical brightness of the BlueWalker 3 satellite. *Nature*, 623(7989):938–941, November 2023. doi: 10.1038/s41586-023-06672-7.

NASA Orbital Debris Program Office. Frequently asked questions. https://orbitaldebris.jsc.nasa.gov/faq/. Accessed: 2026-08-15.

NASA Orbital Debris Program Office. Stability of the future leo environment. Technical report, NASA Johnson Space Center, 2018. URL https://orbitaldebris.jsc.nasa.gov/. Accessed: 2026-01-19.

NASA Orbital Debris Program Office. Optical measurements (modest, es-mcat, etc.). https://orbitaldebris.jsc.nasa.gov/measurements/optical.html, 2023. https://orbitaldebris.jsc.nasa.gov/measurements/optical.html.

National Geospatial Intelligence Agency. World geodetic system 1984 datasheet. https://www.unoosa.org/documents/pdf/icg/2012/template/WGS_84.pdf, 2021. United Nations Office for Outer Space Affairs. Retrieved November 6th, 2025.

National Geospatial-Intelligence Agency. Gnss and gps precise ephemeris data, 2025. URL https://earth-info.nga.mil/index.php?action=gnss&dir=gnss. Office of Geomatics, accessed March 2026.

National Research Council. *Limiting Future Collision Risk to Spacecraft: An Assessment of NASA's Meteoroid and Orbital Debris Programs*. The National Academies Press, Washington, DC, 2011. doi: 10.17226/13244.

Fred E. Nicodemus. Directional reflectance and emissivity of an opaque surface. *Applied Optics*, 4 (7):767, July 1965. doi: 10.1364/AO.4.000767.

NOAA Office of Space Commerce. Traffic coordination system for space (TraCSS). https://space.commerce.gov/traffic-coordination-system-for-space-tracss/. Accessed: 2026-08-15.

Jaime Nomen Torres, Noelia Sánchez-Ortiz, Raúl Domínguez-González, Nuria Guijarro López, and Pablo A. Quiles Ibernón. Accurate Optical Observations of Space Objects in GEO, and applicability to closer LEO regimes. In *7th European Conference on Space Debris*, page 173, April 2017.

Orbital ATK. *GEOStar-2 Bus : A Fully Redundant, Flight-Proven, Spacecraft Bus Designed for Geosynchronous Missions*, 2017.

Orbital Radar. What is Delta-V? The currency of Spaceflight, 2026. Retrieved 2026-08-14, from https://orbitalradar.com/glossary/delta-v.

J. Osborn, D. Föhring, V. S. Dhillon, and R. W. Wilson. Atmospheric scintillation in astronomical photometry. *Monthly Notices of the Royal Astronomical Society*, 452(2):1707–1716, July 2015. ISSN 1365-2966. doi: 10.1093/mnras/stv1400. URL http://dx.doi.org/10.1093/mnras/stv1400.

T. Otsubo, J. Amagai, H. Kunimori, and M. Elphick. Spin motion of the AJISAI satellite derived from spectral analysis of laser ranging data. *IEEE Transactions on Geoscience and Remote Sensing*, 38(3):1417–1424, May 2000. doi: 10.1109/36.843036.

Semeli Papadogiannakis, Torbjorn Sundberg, Per Hagg, and Hanna Sundberg. Polarimetry and spectroscopy on geostationary satellites with the nordic optical telescope. In Stephen Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2022. URL https://amostech.com/TechnicalPapers/2022/Poster/Papadogiannakis.pdf.

P. Papushev, Yu. Karavaev, and M. Mishina. Investigations of the evolution of optical characteristics and dynamics of proper rotation of uncontrolled geostationary artificial satellites. *Advances in Space Research*, 43(9):1416–1422, 2009. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2009.02.007. URL https://www.sciencedirect.com/science/article/pii/S0273117709001021.

William E. Parker and Richard Linares. Satellite drag analysis during the may 2024 gannon geomagnetic storm. *Journal of Spacecraft and Rockets*, 61(5):1412–1416, 2024. doi: 10.2514/1.A36164.

Chris Parmer, Philippe Duval, and Alex Johnson. A data and analytics web app framework for python, no javascript required., July 2025. URL https://github.com/plotly/dash. Version 3.2.0.

T. Payne, S. Gregory, J. Tombasco, K. Luu, and L. Durr. Satellite Monitoring, Change Detection, and Characterization Using Non-Resolved Electro-Optical Data from a Small Aperture Telescope. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E50, January 2007.

Tamara E. Payne, Stephen A. Gregory, and K. Kim Luu. Ssa analysis of geos photometric signature classifications and solar panel offsets. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2006. URL https://amostech.com/TechnicalPapers/2006/NROC/Payne.pdf.

Bui Tuong Phong. Illumination for computer generated pictures. *Commun. ACM*, 18(6):311–317, June 1975. ISSN 0001-0782. doi: 10.1145/360825.360839. URL https://doi.org/10.1145/360825.360839.

Nicolas Pillet. Les satellites gorizont, 5 2017. URL https://www.kosmonavtika.com/satellites/gorizont/gorizont.html. Accessed: 2025-06-30.

PixInsight Development Team. Ccd and cmos image calibration. https://pixinsight.com/doc/docs/Calibration/Calibration.html, 2022. CMOS calibration section.

D. L. Pollacco, I. Skillen, A. Collier Cameron, D. J. Christian, C. Hellier, J. Irwin, T. A. Lister, R. A. Street, R. G. West, D. R. Anderson, W. I. Clarkson, H. Deeg, B. Enoch, A. Evans, A. Fitzsimmons, C. A. Haswell, S. Hodgkin, K. Horne, S. R. Kane, F. P. Keenan, P. F. L. Maxted, A. J. Norton, J. Osborne, N. R. Parley, R. S. I. Ryans, B. Smalley, P. J. Wheatley, and D. M. Wilson. The WASP Project and the SuperWASP Cameras. *Publications of the Astronomical Society of the Pacific*, 118:1407–1418, October 2006. doi: 10.1086/508556.

William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical recipes in C. The art of scientific computing*. Cambridge University Press, 1992.

T. Pultarova. Every spacex starlink satellite has to dodge a collision almost weekly, and experts fear the worst, 2026. URL https://www.space.com/space-exploration/satellites/every-spacex-starlink-satellite-has-to-dodge-a-collision-almost-weekly-and-experts-fear-the-worst. Accessed: 2026-08-15.

M. Quazi, Fazal M. A., A. S. M. A Haseeb, Farazila Yusof, H.H. Masjuki, and Arslan Ahmed. Laser-based surface modifications of aluminum and its alloys. *Critical Reviews in Solid State and Materials Sciences*, 41:1–26, 10 2015. doi: 10.1080/10408436.2015.1076716.

G.D. Reeves, G.L. Delzanno, and Q.R Marksteiner. Earth's Radiation Belts: The Hazards to Satellites and What Can Be Done to Mitigate the Risks? In *Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2020.

Manik Reichegger, Anja Schlicht, and Urs Hugentobler. Space debris: Extraction of the rotational state from multistatic light curves. In *Proceedings of the 22nd International Workshop on Laser Ranging (IWLR)*, Guadalajara, Mexico, 2022. Session S07-05 (paper).

Reliasat. Satellite communications evolution – from geo to leo, 2024. URL https://reliasat.com/satellite-communications-evolution-from-geo-to-leo/.

Brandon Rhodes. Skyfield: High precision research-grade positions for planets and Earth satellites generator. Astrophysics Source Code Library, record ascl:1907.024, July 2019.

Holli. Riebeek. "Catalog of Earth Satellite Orbits". NASA Earth Observatory., 2009. Retrieved 2025-03-19, from https://earthobservatory.nasa.gov/features/OrbitsCatalog/page1.php.

Alejandro. Rivera and Alphonso. Stewart. Study of spacecraft deployables failures. Technical report, NASA, 2021.

M. Rosengren. ERS-1 - An earth observer that exactly follows its chosen path. *ESA Bulletin*, 72: 76–82, November 1992.

Alessandro Rossi, Noelia Sánchez-Ortiz, Emmanuelle David, Roberto Opromolla, and Dmitriy Grishko. Future activities in the near-earth space in the face of ever-increasing space traffic. *Acta Astronautica*, 225:891–897, December 2024. doi: 10.1016/j.actaastro.2024.09.063.

Shabadini Sampath and Jinglang Feng. Intelligent and robust control of space manipulator for sustainable removal of space debris. *Acta Astronautica*, 220:108–117, July 2024. doi: 10.1016/j.actaastro.2024.04.024.

Minoru Sasaki and Hidekazu Hashimoto. Launch and Observation Program of the Experimental Geodetic Satellite of Japan. *IEEE Transactions on Geoscience and Remote Sensing*, 25(5): 526–533, September 1987. doi: 10.1109/TGRS.1987.289830.

J. D. Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *Astrophysical Journal*, 263:835–853, December 1982. doi: 10.1086/160554.

Henrique R. Schmitt. Multi wavelength optical broad band photometric properties of a representative sample of geostationary satellites. *Advances in Space Research*, 65(1):326–336, 2020. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2019.09.036. URL https://www.sciencedirect.com/science/article/pii/S0273117719307148.

A. Schwarzenberg-Czerny. The Correct Probability Distribution for the Phase Dispersion Minimization Periodogram. *Astrophysical Journal*, 489(2):941–945, November 1997. doi: 10.1086/304832.

Günter Seeber. *Satellite Geodesy*. De Gruyter, Berlin, New York, 2003. ISBN 9783110200089. doi: doi:10.1515/9783110200089. URL https://doi.org/10.1515/9783110200089.

L. Sehnal. The perturbations of the orbit of the stationary satellite of the Earth. *Bulletin of the Astronomical Institutes of Czechoslovakia*, 11:132, January 1960.

P. Kenneth Seidelmann, editor. *Explanatory Supplement to the Astronomical Almanac*. University Science Books, 1992.

Patrick Seitzer, Ed Barker, Brent Buckalew, Andrew Burkhardt, Heather Cowardin, James Frith, Catherine Kaleida, Susan M. Lederer, and Chris H. Lee. The Small Size Debris Population at GEO from Optical Observations. In *7th European Conference on Space Debris*, page 3, April 2017.

J. Serrano, A. Gallego, Á. Martínez, and D. Moreno. Monitoring of Rendezvous & Proximity Operations With SST and SDA Techniques Combination. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 134, September 2024.

SES S.A. MEO: Medium Earth Orbit: Performance Without Compromise, 2026. URL https://www.ses.com/network-and-technology/meo. Accessed: 2026-04-24.

Jainisha Shah, Gregory Badura, Ryan Hoffmann, and Daniel Engelhart. Bidirectional Reflectance Distribution Function (BRDF) Measurement of Materials Aged Under Simulated Space Environment. *Journal of the Astronautical Sciences*, 71(1):8, February 2024. doi: 10.1007/s40295-023-00429-0.

L. Shakun, N. Koshkin, E. Korobeynikova, S. Melikyants, S. Strakhova, V. Dragomiretsky, V. Marsakova, A. Ryabov, and S. Terpan. A new model of the ajisai spin axis motion derived from the photometric patterns. *Advances in Space Research*, 2026. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2026.04.061. URL https://www.sciencedirect.com/science/article/pii/S0273117726005326.

Minghe Shan, Jian Guo, and Eberhard Gill. Review and comparison of active space debris capturing and removal methods. *Progress in Aerospace Sciences*, 80:18–32, 2016. ISSN 0376-0421. doi: https://doi.org/10.1016/j.paerosci.2015.11.001. URL https://www.sciencedirect.com/science/article/pii/S0376042115300221.

Sanjib Sharma. Markov Chain Monte Carlo Methods for Bayesian Data Analysis in Astronomy. *Annual Review of Astronomy and Astrophysics*, 55(1):213–259, August 2017. doi: 10.1146/annurev-astro-082214-122339.

Manisha Shrestha, Iain A. Steele, Andrzej S. Piascik, Helen Jermak, Robert J. Smith, and Chris M. Copperwheat. Characterization of a dual-beam, dual-camera optical imaging polarimeter. *MNRAS*, 494(4):4676–4686, April 2020. doi: 10.1093/mnras/staa1049.

Billy Shrive. Classifying low altitude satellites with boosted decision trees. Unpublished, August 2024. URL http://webcat.warwick.ac.uk/record=b4712198.

Billy Shrive, Don Pollacco, Paul Chote, James A. Blake, Benjamin F. Cooke, James McCormac, Richard West, Robert Airey, Alex MacManus, and Phineas Allen. Classifying LEO satellite platforms with boosted decision trees. *RAS Techniques and Instruments*, 3(1):247–256, January 2024. doi: 10.1093/rasti/rzae018.

Jiri Silha, Matej Zigo, Danica Žilková, and Tomas Hrobar. Space debris characterization through photometry and spectroscopy. In *44th COSPAR Scientific Assembly. Held 16-24 July*, volume 44, page 3160, July 2022.

Roger W. Sinnott. Virtues of the Haversine. *Sky Telesc.*, 68:159, 1984.

Jovan Skuljan. Photometric measurements of geostationary satellites over the western pacific region. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2018. URL https://amostech.com/TechnicalPapers/2018/Poster/Skuljan.pdf.

Richard A. Snay and Tomás Soler. Modern terrestrial reference systems (part 1). https://www.ngs.noaa.gov/CORS/Articles/Reference-Systems-Part-1.pdf, 1999. Professional Surveyor, December 1999. Retrieved [Access Date].

E.M. Soop. *Handbook of Geostationary Orbits*. Space Technology Library. Springer Netherlands, 1994. ISBN 9780792330547. URL https://books.google.co.uk/books?id=hqhZKjLaYZUC.

A. Speicher. Identification of geostationary satellites using polarization. Master's thesis, University of Denver, 2015. Visible-spectrum polarimetric data of geosynchronous satellites.

Square Kilometre Array Observatory. Satellite constellations and radio astronomy interference. Technical report, SKAO, 2021. URL https://www.skao.int/en. Accessed: 2026-01-19.

D. Steeghs, D. K. Galloway, K. Ackley, M. J. Dyer, J. Lyman, K. Ulaczyk, R. Cutter, Y. L. Mong, V. Dhillon, P. O'Brien, G. Ramsay, S. Poshyachinda, R. Kotak, L. K. Nuttall, E. Pallé, R. P. Breton, D. Pollacco, E. Thrane, S. Aukkaravittayapun, S. Awiphan, U. Burhanudin, P. Chote, A. Chrimes, E. Daw, C. Duffy, R. Eyles-Ferris, B. Gompertz, T. Heikkilä, P. Irawati, M. R. Kennedy, T. Killestein, H. Kuncarayakti, A. J. Levan, S. Littlefair, L. Makrygianni, T. Marsh, D. Mata-Sanchez, S. Mattila, J. Maund, J. McCormac, D. Mkrtichian, J. Mullaney, K. Noysena, M. Patel, E. Rol, U. Sawangwit, E. R. Stanway, R. Starling, P. Strøm, S. Tooke, R. West, D. J. White, and K. Wiersema. The Gravitational-wave Optical Transient Observer (GOTO): prototype performance and prospects for transient science. *Monthly Notices of the Royal Astronomical Society*, 511(2):2405–2422, April 2022. doi: 10.1093/mnras/stac013.

Iain A. Steele, Robert J. Smith, Paul C. Rees, Ian P. Baker, S. D. Bates, Michael F. Bode, Mark K. Bowman, Dave Carter, Jason Etherton, Martyn J. Ford, Stephen N. Fraser, A. Gomboc, Robert D. J. Lett, Anthony G. Mansfield, Jonathon M. Marchant, Gustavo A. Medrano-Cerda, Christopher J. Mottram, D. Raback, A. B. Scott, M. D. Tomlinson, and R. Zamanov. The Liverpool Telescope: performance and first results. In Jr. Oschmann, Jacobus M., editor, *Ground-based Telescopes*, volume 5489 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 679–692, October 2004. doi: 10.1117/12.551456.

R. F. Stellingwerf. Period determination using phase dispersion minimization. *Astrophysical Journal*, 224:953–960, September 1978. doi: 10.1086/156444.

J. Stryjewski, D. Hand, D. Tyler, S. Murali, M. Roggemann, and N. Peterson. Real Time Polarization Light Curves for Space Debris and Satellites. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E39, September 2010.

Mark A. Sturza and Gemma Saura Carretero. Mega-constellations – a holistic approach to debris aspects. In T. Flohrer, S. Lemmens, and F. Schmitz, editors, *Proceedings of the 8th European Conference on Space Debris*, volume 8. ESA Space Debris Office, 2021. URL https://conference.sdo.esoc.esa.int/proceedings/sdc8/paper/227.

Technische Universität Wien. Vms base time origin. https://web.archive.org/web/20070606064704/http://vms.tuwien.ac.at/info/humour/vms-base-time-origin.txt, 2007. Archived June 06, 2007. Accessed January 19, 2026.

K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces*. *J. Opt. Soc. Am.*, 57(9):1105–1114, 9 1967. doi: 10.1364/JOSA.57.001105. URL https://opg.optica.org/abstract.cfm?URI=josa-57-9-1105.

Mădălina Trelia, Dan Alin Nedelcu, and Mirel Birlan. APSIS: Automated Photometric Survey of Inactive Satellites for rotational dynamics and lightcurve characterization. *Acta Astronautica*, 242:85–106, May 2026. doi: 10.1016/j.actaastro.2026.01.031.

United Nations Office for Outer Space Affairs. Guidelines for the long-term sustainability of outer space activities. Technical report, United Nations, 2019. URL https://www.unoosa.org/oosa/en/ourwork/topics/space-debris.html. Accessed: 2026-01-19.

U.S. Naval Observatory. Computing altitude and azimuth from greenwich apparent sidereal time. Astronomical Applications Department (FAQ web page), a. URL https://aa.usno.navy.mil/faq/alt_az.

U.S. Naval Observatory. Sidereal time. https://aa.usno.navy.mil/faq/docs/GAST.php, b. Accessed 2026.

David A. Vallado. *Fundamentals of Astrodynamics and Applications*. Microcosm Press, 4th edition, 2013.

Melissa F. Vess and Scott R. Starin. A study on the effects of j2 perturbations on a drag-free control system for spacecraft in low earth orbit. Technical Report NASA NTRS 20040031507, NASA Goddard Space Flight Center, Greenbelt, MD, 2003. URL https://ntrs.nasa.gov/citations/20040031507.

F. Vrba, D. Hutter, P. Shankland, J. Armstrong, H. Schmitt, R. Hindsley, M. Divittorio, and J. Benson. A Survey of Geosynchronous Satellite Glints. In S. Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E28, January 2009.

Yang Wang, Xiaoping Du, Jiguang Zhao, Zhiyong Yin, and Yishuo Song. What causes the abrupt changes in the light curve of a geo satellite? *Acta Astronautica*, 153:130–137, 2018. ISSN 0094-5765. doi: https://doi.org/10.1016/j.actaastro.2018.10.035. URL https://www.sciencedirect.com/science/article/pii/S0094576518313857.

Youcun Wang, Min Li, Kecai Jiang, Wenwen Li, Qile Zhao, Rongxin Fang, Na Wei, and Renhai Mu. Improving precise orbit determination of leo satellites using enhanced solar radiation pressure modeling. *Space Weather*, 21(1):e2022SW003292, 2023. doi: https://doi.org/10.1029/2022SW003292. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2022SW003292. e2022SW003292 2022SW003292.

Wayne H. Warren, Jr. *Fifth Fundamental Catalogue (FK5). Part I. Basic Fundamental Stars (Fricke, Schwan, and Lederle 1988). Documentation for the machine-readable version.* 1990.

D. Weisz, A. Dunsmore, J. Key, R. Tucker, E. Weld, F. Chun, and R. Tippets. Comparison of Geosynchronous Satellites Spectral Signatures During Glint Season. In Sandy Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page 126, September 2016.

Peter J Wheatley, Richard G West, Michael R Goad, James S Jenkins, Don L Pollacco, Didier Queloz, Heike Rauer, Stéphane Udry, Christopher A Watson, Bruno Chazelas, Philipp Eigmüller, Gregory Lambert, Ludovic Genolet, James McCormac, Simon Walker, David J Armstrong, Daniel Bayliss, Joao Bento, François Bouchy, Matthew R Burleigh, Juan Cabrera, Sarah L

Casewell, Alexander Chaushev, Paul Chote, Szilárd Csizmadia, Anders Erikson, Francesca Faedi, Emma Foxell, Boris T Gänsicke, Edward Gillen, Andrew Grange, Maximilian N Günther, Simon T Hodgkin, James Jackman, Andrés Jordán, Tom Louden, Lionel Metrailler, Maximiliano Moyano, Louise D Nielsen, Hugh P Osborn, Katja Poppenhaeger, Roberto Raddi, Liam Raynard, Alexis M S Smith, Maritza Soto, and Ruth Titz-Weider. The next generation transit survey (ngts). *Monthly Notices of the Royal Astronomical Society*, 475(4):4476–4493, 11 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx2836. URL https://doi.org/10.1093/mnras/stx2836.

Klaas Wiersema, Paul Chote, Jonathan Marchant, Stefano Covino, Justyn R. Maund, Alexander Agathanggelou, William Feline, Simon George, Grant Privett, Brooke Simmons, and Iain A. Steele. Short timescale imaging polarimetry of geostationary satellite thor-6: The nature of micro-glints. *Advances in Space Research*, 70(10):3003–3015, 2022. ISSN 0273-1177. doi: https://doi.org/10.1016/j.asr.2022.07.034. URL https://www.sciencedirect.com/science/article/pii/S027311772200655X.

Zhao Xiao-Fen, Yu Yong, Mao Yin-Dun, and Tang Zheng-hong. Long-term photometric signature study of two geo satellites. *Advances in Space Research*, 67:2241–2251, 2021. URL https://api.semanticscholar.org/CorpusID:233908407.

Andrew T. Young. Air mass and refraction. *Appl. Opt.*, 33(6):1108–1110, 2 1994. doi: 10.1364/AO.33.001108. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-33-6-1108.

Yili Zhao, Huimin Li, Chuntao Chen, and Jianhua Zhu. Verification and recalibration of HY-2A microwave radiometer brightness temperature. *Journal of Oceanology and Limnology*, 37(3): 968–981, 2019. doi: 10.1007/s00343-019-8073-5. URL https://archimer.ifremer.fr/doc/00503/61427/65584.pdf.

Matej Zigo, Jiří Šilha, Tomáš Hrobár, and Peter Jevčák. Space debris surface characterization through BVR_cI_c photometry. *Advances in Space Research*, 72(9):3802–3817, November 2023. doi: 10.1016/j.asr.2023.07.029.

Michał Żołnowski, Krzysztof Kamiński, Tobias Delbruck, Alicja Gąsiorowska, Justyna Gołębiewska, Marcin Gędek, Monika Kamińska, Joachim Kruger, Mikołaj Krużyński, Mikołaj Pieniowski, Rafał Reszelewski, Krzysztof Romanowski, Edwin Wnuk, Wojciech Dimitrow, and Wojciech Mich. Astrometric measurements of satellites with Dynamic Vision Sensor. In *8th European Conference on Space Debris*, page 176, April 2021.

A BRDF Models

If you can keep your head when all about you
    Are losing theirs and blaming it on you,
If you can trust yourself when all men doubt you,
    But make allowance for their doubting too;
If you can wait and not be tired by waiting,
    Or being lied about, don't deal in lies,
Or being hated, don't give way to hating,
    And yet don't look too good, nor talk too wise:

If you can dream—and not make dreams your master;
    If you can think—and not make thoughts your aim;
If you can meet with Triumph and Disaster
    And treat those two impostors just the same;
If you can bear to hear the truth you've spoken
    Twisted by knaves to make a trap for fools,
Or watch the things you gave your life to, broken,
    And stoop and build 'em up with worn-out tools:

If you can make one heap of all your winnings
    And risk it on one turn of pitch-and-toss,
And lose, and start again at your beginnings
    And never breathe a word about your loss;
If you can force your heart and nerve and sinew
    To serve your turn long after they are gone,
And so hold on when there is nothing in you
    Except the Will which says to them: 'Hold on!'

If you can talk with crowds and keep your virtue,
    Or walk with Kings—nor lose the common touch,
If neither foes nor loving friends can hurt you,
    If all men count with you, but none too much;
If you can fill the unforgiving minute
    With sixty seconds' worth of distance run,
Yours is the Earth and everything that's in it,
    And—which is more—you'll be a Man, my son!

If—Rudyard Kipling

This appendix section describes how some common BRDF models can be implemented using a consistent vector notation. Incident and outgoing unit directions are denoted by $\hat{\omega}_i$ and $\hat{\omega}_o$ respectively, and the surface normal by $\hat{n}$. The perfect reflection direction is:

$$\hat{\omega}_r = 2(\hat{n} \cdot \hat{\omega}_i)\,\hat{n} - \hat{\omega}_i. \tag{A.1}$$

All BRDFs are evaluated only over the positive hemisphere:

$$(\hat{\omega}_i \cdot \hat{n}) > 0 \quad \text{and} \quad (\hat{\omega}_o \cdot \hat{n}) > 0. \tag{A.2}$$

Lambertian

The Lambertian model assumes perfectly diffuse reflection:

$$f_r(\hat{\omega}_i, \hat{\omega}_o) = \frac{\rho}{\pi}. \tag{A.3}$$

It satisfies reciprocity and energy conservation with $0 \leqslant \rho \leqslant 1$.

ABG Scattering

The ABG model is described in Moulin (2021) and defined as:

$$f_r = \frac{A}{B + x^g}, \tag{A.4}$$

where:

$$\vec{\beta} = \hat{\omega}_o - (\hat{n} \cdot \hat{\omega}_o)\,\hat{n}, \tag{A.5}$$

$$\vec{\beta}_0 = \hat{\omega}_r - (\hat{n} \cdot \hat{\omega}_r)\,\hat{n}, \tag{A.6}$$

$$x = \left\| \vec{\beta} - \vec{\beta}_0 \right\|. \tag{A.7}$$

Here $\vec{\beta}$ and $\vec{\beta}_0$ are the projections of the outgoing and specular directions onto the surface plane, so that x measures the angular deviation from the specular direction. Parameters: A peak strength, B knee location, g exponential decay slope.

Gaussian Microfacet

The Gaussian specular BRDF model follows standard microfacet-based rough surface scattering theory (Torrance and Sparrow, 1967; Beckmann and Spizzichino, 1987). A Gaussian lobe around $\hat{\omega}_r$:

$$f_r = A \exp\left(\frac{\hat{\omega}_r \cdot \hat{\omega}_o - 1}{\sigma} \right), \tag{A.8}$$

where A is the peak height and σ controls the width of the lobe. Often used for polished optics.

Phong Model

The Phong model combines Lambertian diffuse and glossy specular reflection (Phong, 1975):

$$f_r(\hat{\omega}_i, \hat{\omega}_o) = \frac{K_d}{\pi} + K_s \frac{\alpha + 2}{2\pi} \left(\hat{\omega}_r \cdot \hat{\omega}_o \right)^{\alpha}, \tag{A.9}$$

where α is the specular exponent and K_d, K_s are the diffuse and specular reflectances. Energy conservation requires:

$$K_d + K_s \leqslant 1. \tag{A.10}$$

Cook–Torrance Microfacet

The Cook-Torrance model was introduced in Cook and Torrance (1982). It is defined using the microfacet distribution D, Fresnel F and geometric shadowing-masking G:

$$f_r = \frac{D(\hat{h})\, F(\hat{\omega}_i)\, G(\hat{\omega}_i, \hat{\omega}_o)}{4(\hat{\omega}_i \cdot \hat{n})(\hat{\omega}_o \cdot \hat{n})}, \tag{A.11}$$

where the half-vector is:

$$\hat{h} = \frac{\hat{\omega}_i + \hat{\omega}_o}{\|\hat{\omega}_i + \hat{\omega}_o\|}. \tag{A.12}$$

A common choice for Fresnel:

$$F(\hat{\omega}_i) = F_0 + (1 - F_0)\left(1 - \hat{h} \cdot \hat{\omega}_i\right)^5. \tag{A.13}$$

Masking-shadowing (Smith):

$$G(\hat{\omega}_i, \hat{\omega}_o) = G_1(\hat{\omega}_i)\, G_1(\hat{\omega}_o), \tag{A.14}$$

with:

$$G_1(\hat{\omega}) = \frac{2(\hat{n} \cdot \hat{\omega})(\hat{n} \cdot \hat{h})}{(\hat{\omega} \cdot \hat{h})}. \tag{A.15}$$

This model is reciprocal and energy conserving.

B Typical Frame Transforms

Before presenting the transforms themselves, it is worth briefly summarising the coordinate systems involved, since each is suited to a different purpose and they are easily conflated.

Geodetic coordinates (ϕ, λ, h) are the familiar latitude, longitude, and height triple used for Earth-based locations. Latitude and longitude are defined with respect to a reference ellipsoid (typically WGS84) rather than a true sphere, and the height h is the *ellipsoidal* height — the perpendicular distance above the ellipsoid surface. This is not the same as height above mean sea level: mean sea level is approximated by the geoid, an equipotential surface of Earth's gravity field, and the difference between the geoid and the WGS84 ellipsoid (the *geoid undulation*) can exceed $\pm 100\,\mathrm{m}$ in places. GPS receivers natively report ellipsoidal height, and a geoid model must be applied to recover orthometric (sea-level) height.

Earth-Centred, Earth-Fixed (ECEF) coordinates (X, Y, Z), of which ITRF is a specific realisation, are Cartesian coordinates with origin at Earth's centre of mass. The Z-axis points along the mean rotation axis toward the conventional North Pole, the X-axis points toward the intersection of the equator and the IERS Reference Meridian (near Greenwich), and the Y-axis completes a right-handed system. ECEF rotates with the Earth, so a fixed point on the surface has constant ECEF coordinates (modulo small effects like polar motion and tectonic drift). ECEF is convenient for describing positions of ground stations, satellite tracks relative to the surface, and anything where "where on Earth" matters more than "where in space".

Earth-Centred Inertial (ECI) coordinates, with J2000 being the standard realisation, are also Cartesian and centred on Earth, but the axes are fixed relative to the distant stars rather than rotating with the Earth. The J2000 frame's axes are aligned with the mean equator and mean equinox at the J2000.0 epoch (12:00 TT on 1 January 2000). Because Newton's laws of motion take their simple form only in an inertial (non-rotating, non-accelerating) frame, ECI is the natural frame in which to integrate orbital dynamics. Note that strictly speaking ECI is only *quasi*-inertial because Earth itself accelerates around the Sun, but for near-Earth dynamics this is negligible.

The International Celestial Reference System (ICRS) is the modern realisation of an inertial frame for astronomical purposes. Its axes are defined kinematically by the positions of distant extra-galactic radio sources rather than by Earth's equator and equinox, making it independent of Earth's orientation. For most practical purposes ICRS and J2000 are aligned to within milliarcsecond-level offsets and the two are often used interchangeably, but the distinction matters for high-precision astrometry.

Topocentric horizontal (alt-az) coordinates describe a direction as seen by an observer on Earth's surface: altitude a is the angle above the local horizon and azimuth A is the bearing measured around the horizon (here, East of North). Unlike the frames above, alt-az depends on

both the observer's location and the observation time, since the sky rotates overhead. This is the natural frame for pointing a ground-based telescope or antenna.

Geodetic Coordinates to ECEF

Geodetic coordinates (ϕ, λ, h) can be converted into ECEF coordinates (X, Y, Z) using:

$$
\begin{aligned}
X &= \big(N(\phi) + h\big)\cos\phi\cos\lambda,\\
Y &= \big(N(\phi) + h\big)\cos\phi\sin\lambda,\\
Z &= \left(\frac{b^2}{a^2}N(\phi) + h\right)\sin\phi = \big((1-e^2)N(\phi) + h\big)\sin\phi = \big((1-f)^2N(\phi) + h\big)\sin\phi,
\end{aligned}
$$

where:

$$
N(\phi) = \frac{a^2}{\sqrt{a^2\cos^2\phi + b^2\sin^2\phi}} = \frac{a}{\sqrt{1 - e^2\sin^2\phi}} = \frac{a}{\sqrt{1 - \frac{e^2}{1+\cot^2\phi}}}.
$$

Here:

$$
\begin{aligned}
&a \text{ is the equatorial radius (semi-major axis),}\\
&b \text{ is the polar radius (semi-minor axis),}\\
e^2 &= 1 - \frac{b^2}{a^2} \quad \text{(first numerical eccentricity squared),}\\
f &= 1 - \frac{b}{a} \quad \text{(flattening of the ellipsoid).}
\end{aligned}
$$

The prime vertical radius of curvature $N(\phi)$ is the distance from the ellipsoid surface to the Z-axis along the normal.

ECEF (ITRF) to ECI (J2000)

ECEF position, velocity, and acceleration vectors $(\mathbf{r}_{ECEF}, \boldsymbol{v}_{ECEF}, \boldsymbol{a}_{ECEF})$ are transformed into J2000 inertial coordinates $(\mathbf{r}_{J2000}, \boldsymbol{v}_{J2000}, \boldsymbol{a}_{J2000})$ using:

$$
\begin{aligned}
\vec{\mathbf{r}}_{J2000} &= PNR\,W\vec{\mathbf{r}}_{ECEF},\\
\vec{\boldsymbol{v}}_{J2000} &= PNR\big(W\vec{\boldsymbol{v}}_{ECEF} + \omega_E \times (W\vec{\mathbf{r}}_{ECEF})\big),\\
\vec{\boldsymbol{a}}_{J2000} &= PNR\Big(W\,\vec{\boldsymbol{a}}_{ECEF} + \omega_E \times (\omega_E \times \vec{\mathbf{r}}) + 2\,\omega_E \times \vec{\boldsymbol{v}}_{ECEF}\Big).
\end{aligned}
$$

$$
R_1(x) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos x & \sin x \\ 0 & -\sin x & \cos x \end{bmatrix} \tag{B.1}
$$

$$
R_2(x) = \begin{bmatrix} \cos x & 0 & -\sin x \\ 0 & 1 & 0 \\ \sin x & 0 & \cos x \end{bmatrix} \tag{B.2}
$$

$$R_3(x) = \begin{bmatrix} \cos x & \sin x & 0 \\ -\sin x & \cos x & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{B.3}$$

The composite rotation matrix is

$$PNR = R_{mod\to eci}\, R_{tod\to mod}\, R_{pef\to tod},$$

with

$$\begin{aligned} R_{pef\to tod} &= R_3(-\mathrm{GAST}), \\ R_{tod\to mod} &= R_1(-\bar{\epsilon})\, R_3(\Delta\psi)\, R_1(\epsilon), \\ R_{mod\to eci} &= R_3(\zeta)\, R_2(-\theta)\, R_3(z). \end{aligned}$$

Polar motion is given by

$$W = \begin{bmatrix} c_x & 0 & -s_x \\ s_x s_y & c_y & c_x s_y \\ s_x c_y & -s_y & c_x c_y \end{bmatrix}, \qquad \begin{aligned} c_x &= \cos(x_p), \\ s_x &= \sin(x_p), \\ c_y &= \cos(y_p), \\ s_y &= \sin(y_p). \end{aligned}$$

The Earth rotation vector is

$$\vec{\omega}_E = \begin{bmatrix} 0 \\ 0 \\ \omega_E \end{bmatrix}, \qquad \omega_E = \omega_0 \left(1 - \frac{\mathrm{LOD}}{86400}\right), \quad \omega_0 = 7.2921151467 \times 10^{-5}\ \mathrm{rad/s}.$$

where the Earth Orientation Parameters (EOP) provide:

$$\begin{aligned} x_p, y_p &\text{ polar motion parameters (arcseconds)}, \\ \mathrm{LOD} &\text{ length of day correction (seconds)}, \\ \mathrm{DUT1} = \mathrm{UT1} - \mathrm{UTC} &\text{ (seconds)}, \\ \Delta\psi, \Delta\epsilon &\text{ nutation corrections (arcseconds)}. \end{aligned}$$

Additional parameters are defined as:

$$\begin{aligned} \mathrm{GAST} &\text{ Greenwich Apparent Sidereal Time (radians)}, \\ \epsilon, \bar{\epsilon} &\text{ true and mean obliquity of the ecliptic}, \\ \zeta, \theta, z &\text{ IAU-76 precession angles}, \\ W &\text{ polar motion matrix}, \\ \mathrm{PNR} &\text{ combined precession–nutation–rotation matrix}. \end{aligned}$$

ICRS to Alt-Az

An ICRS sky position with right ascension α and declination δ can be transformed to topocentric horizontal coordinates (azimuth A, altitude a) for an observer at geodetic location (ϕ, λ, h) and observation time t as follows. λ is taken positive east of Greenwich, and azimuth A is defined to be measured East of North (N= 0°, E= 90°), consistent with the `Astropy` 'Altaz' convention (Astropy Collaboration, a).

First form the ICRS unit direction vector

$$\hat{\mathbf{p}}_{\mathrm{ICRS}} = \begin{bmatrix} \cos\delta\cos\alpha \\ \cos\delta\sin\alpha \\ \sin\delta \end{bmatrix}.$$

The hour-angle formulation for horizontal coordinates uses *apparent* sidereal time and therefore requires *apparent-of-date* equatorial coordinates ("apparent place"). Let $\mathrm{BPN}(t)$ denote the bias–precession–nutation rotation taking the geocentric ICRS/GCRS direction to the true equator and true equinox of date (TETE; apparent place) (International Earth Rotation and Reference Systems Service (IERS), 2010). Then

$$\hat{\mathbf{p}}_{\mathrm{app}} = \mathrm{BPN}(t)\,\hat{\mathbf{p}}_{\mathrm{ICRS}} = \begin{bmatrix} x_{\mathrm{app}} \\ y_{\mathrm{app}} \\ z_{\mathrm{app}} \end{bmatrix},$$

and the corresponding apparent-of-date right ascension and declination are

$$\alpha_{\mathrm{app}} = \mathrm{atan2}\big(y_{\mathrm{app}}, x_{\mathrm{app}}\big), \qquad \delta_{\mathrm{app}} = \mathrm{atan2}\Big(z_{\mathrm{app}}, \sqrt{x_{\mathrm{app}}^2 + y_{\mathrm{app}}^2}\Big).$$

(Equivalently, in `Astropy` one may obtain $(\alpha_{\mathrm{app}}, \delta_{\mathrm{app}})$ via the frame transform `ICRS→TETE(obstime=t)` (Astropy Collaboration, b).)

Compute the local apparent sidereal time $\theta_{\mathrm{LAST}}(t)$ at longitude λ:

$$\theta_{\mathrm{LAST}}(t) = \theta_{\mathrm{GAST}}(t) + \lambda,$$

and form the hour angle (USNO "HA") (U.S. Naval Observatory, a)

$$H = \theta_{\mathrm{LAST}}(t) - \alpha_{\mathrm{app}}.$$

(`Astropy` provides θ_{LAST} via `Time.sidereal_time('apparent', longitude=λ)` (Astropy Collaboration, c).)

Using observer latitude ϕ, compute the altitude a from (U.S. Naval Observatory, a)

$$\sin a = \sin\phi\,\sin\delta_{\mathrm{app}} + \cos\phi\,\cos\delta_{\mathrm{app}}\,\cos H, \qquad a = \arcsin(\sin a),$$

and compute the azimuth using a two-argument arctangent (robust quadrant handling) (U.S. Naval Observatory, a)

$$A = \mathrm{atan2}\big(-\sin H,\ \tan\delta_{\mathrm{app}}\cos\phi - \sin\phi\,\cos H\big), \qquad A \leftarrow A \bmod 2\pi.$$

The above yields geometric (vacuum) topocentric (A, a). If atmospheric refraction is required, an apparent altitude may be formed as $a_{\mathrm{refr}} = a + R(a; P, T, \ldots)$ for a chosen refraction model. In Astropy, refraction is included within the `AltAz` frame when `pressure` is non-zero (Astropy Collaboration, a).

C Ajisai Empirical Model Calculations

The orientation of the Ajisai spin axis can be expressed by a set of time-dependent equations, following the empirical model developed by Kucharski et al. (2016). The coordinates of a vector oriented at the azimuthal angle ϕ about the concentric cone of radius R, and with its axis at right ascension α and declination δ, are defined by the transformation:

$$S = R_3(-\alpha)\, R_2(\delta - 90^\circ)\, R_3(-\phi)\, R_2(-R), \tag{C.1}$$

where R_3 and R_2 are the rotation matrices about the z- and y-axes, respectively, as defined in Equations B.3 and B.2.
The inertial orientation of the nutation cone axis $\vec{\nu}_N$ is calculated as:

$$\hat{\nu}_N = S \cdot \hat{z}, \quad \text{with } \hat{z} = [0, 0, 1] \tag{C.2}$$

The parameters for the matrix S describing the precession cone are:

1. Right ascension of the cone axis: $\alpha = 88.90^\circ$
2. Declination of the cone axis: $\delta = -88.85^\circ$
3. Radius of the precession cone: $R = 1.08^\circ$
4. Precession angle (azimuth): $\phi = -0.0277404\, D + 452.803^\circ$

where D is the number of days since launch (MJD 46654.86).
The spin axis vector $\hat{\nu}_S$ is calculated using a time-varying nutation cone defined about $\hat{\nu}_N$, with parameters:

1. Nutation cone radius:

$$R = 7.66693 \times 10^{-9}\, D^2 - 0.0000656721\, D + 1.42878^\circ,$$

2. Nutation angle (azimuth):

$$\phi = 8.08453575 \times 10^{-7}\, D^2 + 3.07506\, D - 19238.5^\circ$$

Using these parameters, the transformation matrix S is recomputed with the updated α and δ (from $\hat{\nu}_N$), and the spin axis vector is given by:

$$\hat{\nu}_S = S \cdot \hat{z} \tag{C.3}$$

This set of empirical models was used for the 2019 dataset. For the 2025 dataset, the empirical model was updated using the recently published model of Shakun et al. (2026), which is based on Ajisai spin-pole estimates from the 2009–2024 interval. The functional form of the model remains unchanged, i.e., the spin axis is still represented as a nutation cone whose axis follows a precession cone. However, the reference epoch and the fitted coefficients are updated.

The main differences between the Kucharski et al. (2016) parameterisation and the Shakun et al. (2026) model are:

- The time variable is redefined. In the Kucharski (2016) model, D is measured from launch epoch, MJD 46654.86. In the Shakun et al. (2026) model, D is measured from epoch 2016.0, corresponding to MJD 57388.0.
- The precession-cone parameters are updated to

$$\alpha_p = 96.94^\circ, \qquad \delta_p = -89.00^\circ, \qquad R_p = 1.24^\circ.$$

- The precession phase is now given by the linear expression

$$\phi_p = -0.02858\,D + 164.0^\circ.$$

- The nutation-cone radius is no longer represented by a quadratic function of time. Instead, it is described by the linear relation

$$R_n = 7.368 \times 10^{-5}\,D + 1.7975^\circ.$$

- The nutation phase is also represented by a linear function,

$$\phi_n = 3.09124\,D + 181.40^\circ.$$

- The quadratic terms present in the Kucharski (2016) nutation radius and phase expressions are omitted in the updated fit.

Thus, the nutation-cone axis is first computed as

$$\hat{\nu}_N = R_3(-\alpha_p)R_2(\delta_p - 90^\circ)R_3(-\phi_p)R_2(-R_p)\hat{z}, \tag{C.4}$$

where $\hat{z} = [0, 0, 1]$. The resulting vector $\hat{\nu}_N$ is then converted back to right ascension and declination, (α_N, δ_N), and the final spin-axis vector is obtained from

$$\hat{\nu}_S = R_3(-\alpha_N)R_2(\delta_N - 90^\circ)R_3(-\phi_n)R_2(-R_n)\hat{z}. \tag{C.5}$$

SPIN PERIOD EVOLUTION

The spin period P, of Ajisai is modelled as an exponential trend function:

$$P(D) = 1.4934 \cdot \exp(0.000040553 \cdot D)\,[s], \tag{C.6}$$

which describes the gradual loss of rotational energy due to weak magnetic torque effects. This empirical model was used for both the 2019 and 2025 datasets

D Haversine Based Separation Calculation

The calculation of the quantity $sep(x_i, y_j)$ in chapter 5 is calculated using the Haversine formula for small distances (Sinnott, 1984). These are described by Equations D.1 and D.2 respectively. The subscripts of λ and ϕ denote the longitude and latitude components of x_i and y_j respectively.

$$\begin{aligned} A(x_i, y_j) = \sin^2\left(\frac{(y_j)_\phi - (x_i)_\phi}{2}\right) \\ + \cos\big((x_i)_\phi\big)\cos\big((y_j)_\phi\big)\sin^2\left(\frac{(y_j)_\lambda - (x_i)_\lambda}{2}\right) \end{aligned} \tag{D.1}$$

$$sep(x_i, y_j) = 2\arctan 2\left(\sqrt{A(x_i, y_j)}, \sqrt{1 - A(x_i, y_j)}\right) \tag{D.2}$$

E Ajisai Individual Corner Plots

SuperWASP Observations

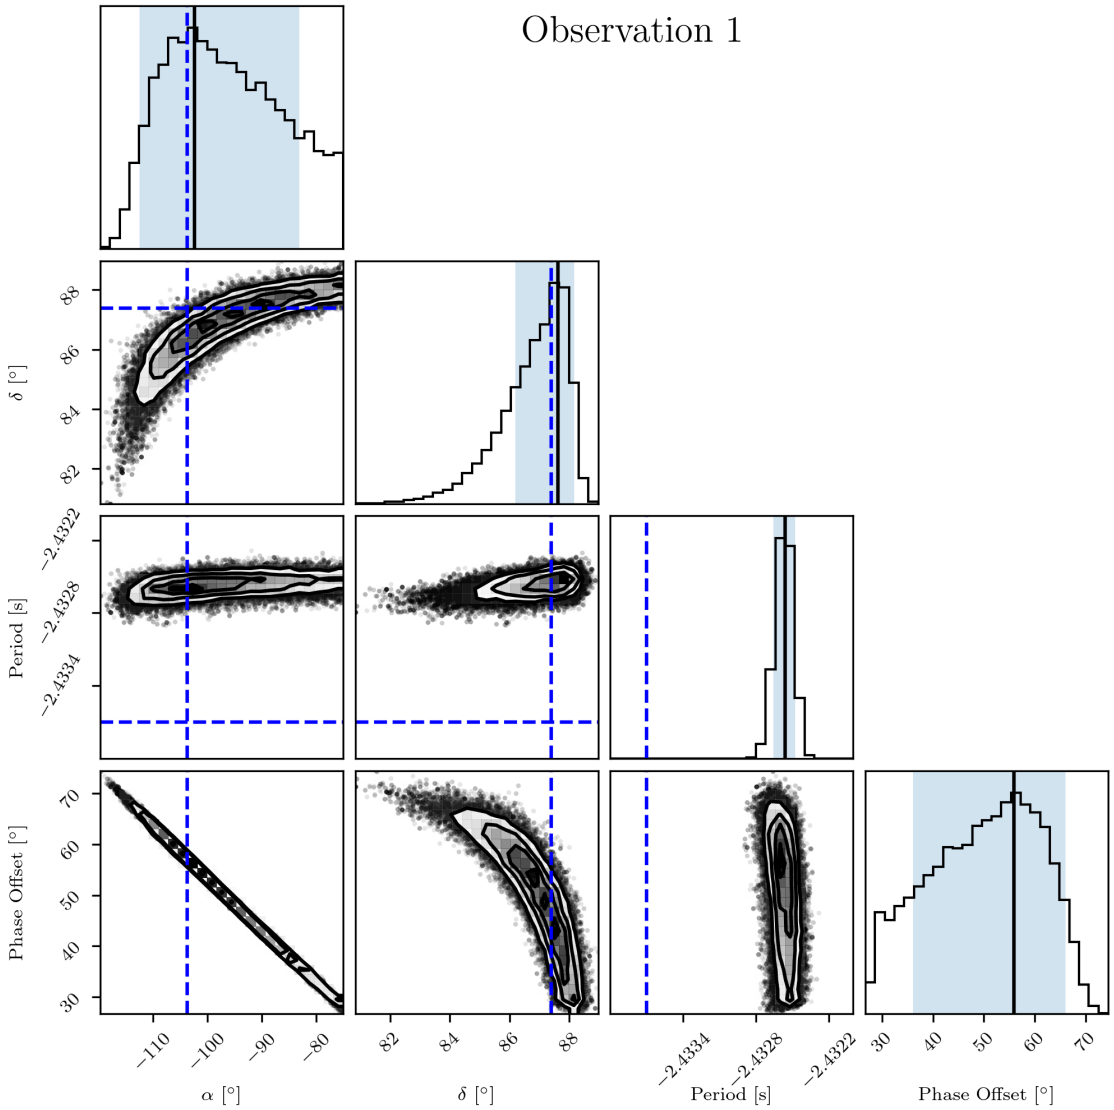

Figure E.1: Corner plot from the MCMC run on the data from the night of 2019-08-01. The blue dashed lines indicate the empirical value estimate from Kucharski et al. (2016), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimate (e.g., rotational period) do not pass through a minimum, this may be indicative of variations for example, seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density.

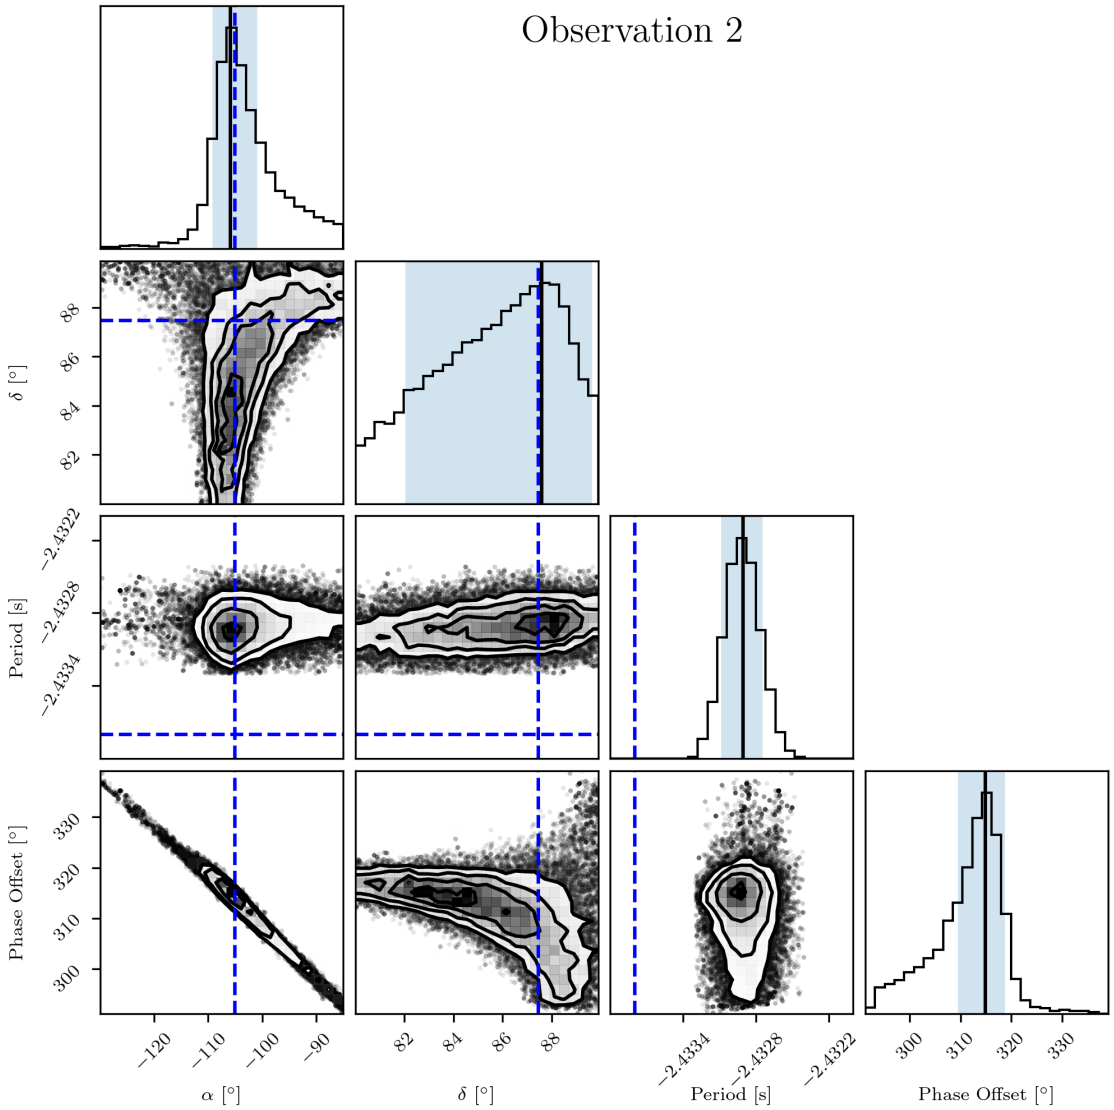

Figure E.2: Corner plot from the MCMC run on the data from the night of 2019-08-02. The blue dashed lines indicate the empirical value estimate from Kucharski et al. (2016), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimate (e.g., rotational period) do not pass through a minimum, this may be indicative of variations for example, seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density.

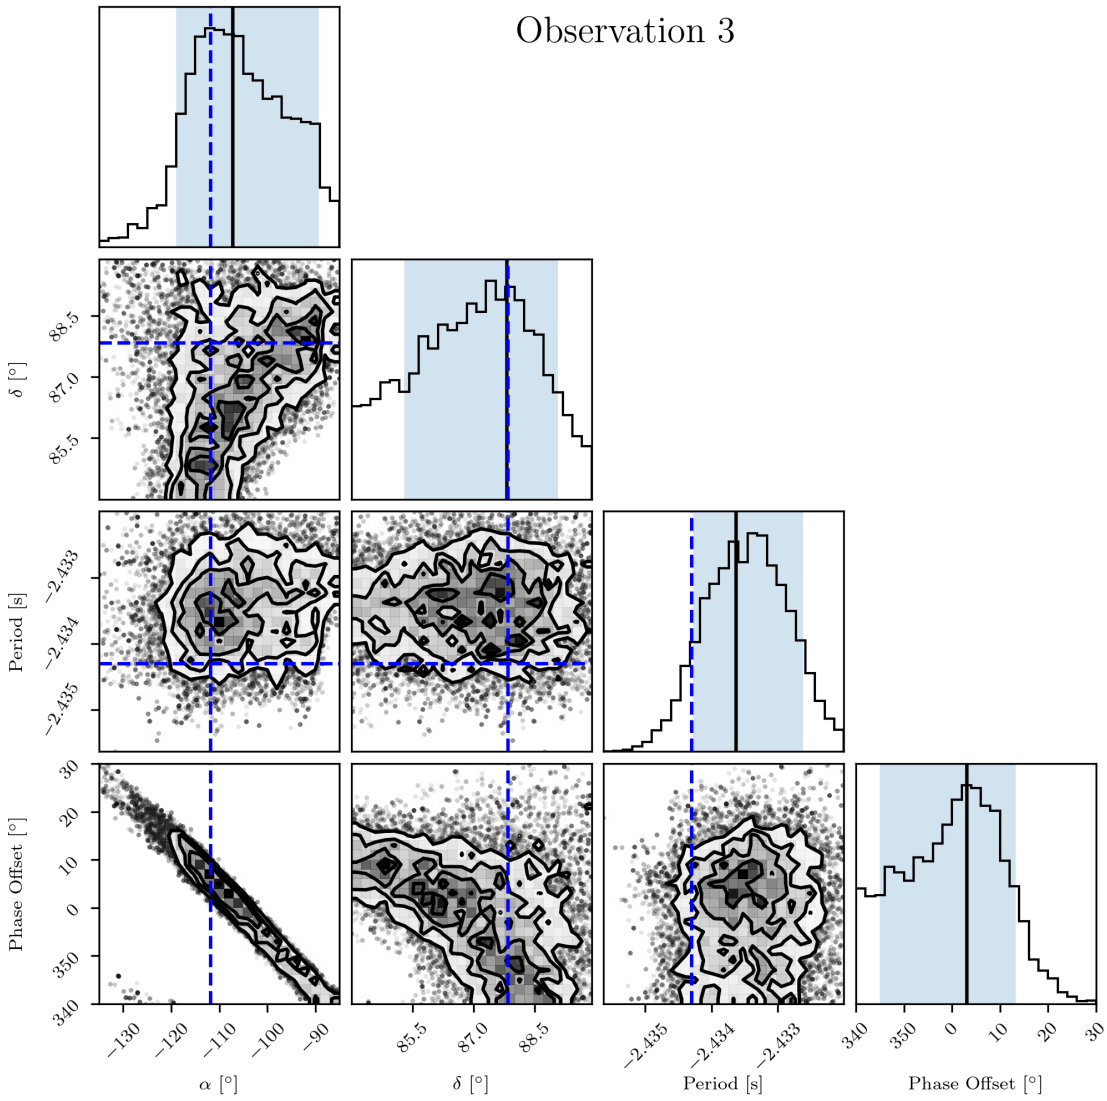

Figure E.3: Corner plot from the MCMC run on the data from the night of 2019-08-07. The blue dashed lines indicate the empirical value estimate from Kucharski et al. (2016), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimate (e.g., rotational period) do not pass through a minimum, this may be indicative of variations for example, seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density.

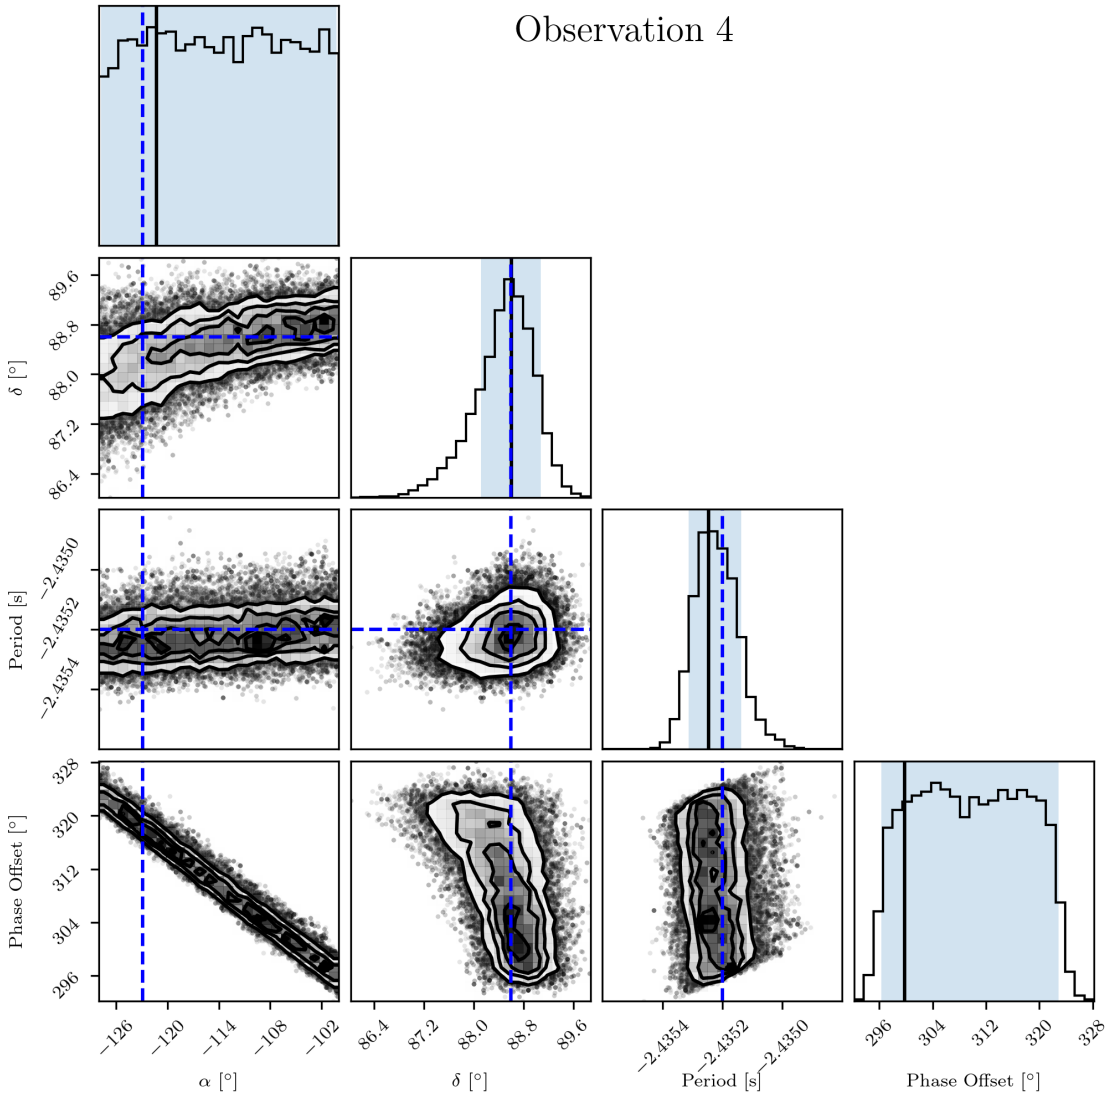

Figure E.4: Corner plot from the MCMC run on the data from the night of 2019-08-16. The blue dashed lines indicate the empirical value estimate from Kucharski et al. (2016), with the black lines and blue shaded regions representing the best estimate and error bounds from the mode and FWHM of the sample distributions. In some instances, the empirical estimate (e.g., rotational period) do not pass through a minimum, this may be indicative of variations for example, seasonal illumination changes. Contours show 2-D highest-posterior-density credible regions from the MCMC samples; darker shading indicates higher posterior density.