

Characterization of Resident Space Objects using Light Curves

M.Sc. David Vallverdú Cabrera

Vollständiger Abdruck der von der Fakultät für Luft- und Raumfahrttechnik der Universität der Bundeswehr München zur Erlangung des akademischen Grades eines

Doktor-Ingenieurs (Dr.-Ing.)

angenommenen Dissertation.

Gutachter/Gutachterin:

1. Prof. Dr.-Ing. Roger Förstner
2. Prof. Dr. phil. nat. Carolin Frueh

Die Dissertation wurde am 19. Juli 2023 bei der Universität der Bundeswehr München eingereicht und durch die Fakultät für Luft- und Raumfahrttechnik am 28. März 2024 angenommen. Die mündliche Prüfung fand am 06. Mai 2024 statt.

Abstract

The population of Earth-orbiting objects has grown steadily at an increasing pace. Unless this trend is reversed, exploiting the near-Earth space environment will eventually become unfeasible. Achieving a comprehensive portrait on how many objects there are, where, and what their characteristics are, is key to devise mitigation strategies for this problem. This thesis addresses the latter point: it explores which object characteristics can be estimated, and how, using time-series of brightness measurements (a.k.a. light curves). An initial literature review identified three main research areas in which this thesis could contribute scientifically.

First, a new rendering method used to simulate light curves of non-convex objects has been developed. Under the assumption that the object is a polyhedron with flat facets, it compares these facets with each other to determine which are observable and illuminated. When the object is composed by a few, big facets, it is shown that this new method outperforms the traditional ray-tracing approach, both in terms of accuracy and speed.

Next, an existing two-step process that estimates the shape of an object from its light curve and attitude ephemerides has been analysed. This method 1) extracts an Extended Gaussian Image (EGI) using Weighted Least Squares, from which it then 2) reconstructs a convex polyhedron using Minkowski minimization. This thesis' contribution shows that dense EGIs can be used even when the first step is not observable according to linear control theory; deduces an analytical expression of the Hessian of the volume constraint of the second step; and introduces the novel homotheticity, a criterion that enables assessing the performance of the method based on the final polyhedron, instead of the intermediate EGI.

Finally, a new approach for Attitude Determination (AD) from light curves has been developed. It models the probability distribution of the state as a Gaussian mixture, and assumes that shape and surface reflectance are known. Each kernel is updated with an Unscented Kalman Filter (UKF), while the mixture size is adapted at each step using a non-linearity measure of the light curve model. Simulations show that this filter, dubbed Adaptive Gaussian Mixtures Unscented Kalman Filter (AGMUKF) for AD, outperforms the traditional UKF. Furthermore, results suggest that the AGMUKF for AD has the potential to match the state-of-the-art Particle Filter (PF) performance-wise, but at a lower computational cost thanks to its ability to adapt the mixture size.

This thesis is closed by a demonstration that the contributions of this thesis to the field of Object Characterization can potentially constitute building blocks of derived use cases, such as health monitoring of operational satellites.

Kurzfassung

Der Bestand an Objekten in der Erdumlaufbahn ist stetig und in immer schnellerem Tempo gewachsen. Wenn sich dieser Trend nicht umkehrt, wird die Nutzung des erdnahen Weltraums irgendwann nicht mehr möglich sein. Ein umfassendes Bild davon, wie viele Objekte es gibt, wo sie sich befinden und welche Eigenschaften sie haben, ist der Schlüssel zur Entwicklung von Strategien zur Eindämmung dieses Problems. Die vorliegende Arbeit befasst sich mit dem letztgenannten Punkt: Es wird untersucht, welche Objekteigenschaften auf welche Weise anhand von Zeitreihen von Helligkeitsmessungen (auch bekannt als Lichtkurven) geschätzt werden können. Eine erste Literaturrecherche ergab drei Hauptforschungsbereiche, zu denen diese Arbeit einen wissenschaftlichen Beitrag leisten kann.

Erstens wurde eine neue Bildsynthese-Methode für die Simulation von Lichtkurven nicht-konvexer Objekte entwickelt. Unter der Annahme, dass das Objekt ein Polyeder mit flachen Facetten ist, vergleicht es diese Facetten miteinander, um festzustellen, welche davon beobachtbar und beleuchtet sind. Wenn das Objekt aus wenigen, großen Facetten besteht, wird gezeigt, dass diese neue Methode den traditionellen Ray-Tracing-Ansatz sowohl in Bezug auf die Genauigkeit als auch auf die Geschwindigkeit übertrifft.

Als Nächstes wurde ein bestehendes zweistufiges Verfahren analysiert, das die Form eines Objekts aus seiner Lichtkurve und den Lageephemeriden schätzt. Diese Methode 1) extrahiert ein Extended Gaussian Image (EGI) mit Hilfe der gewichteten Methode der kleinsten Quadrate und rekonstruiert daraus 2) ein konvexes Polyeder mit Hilfe der Minkowski-Minimierung. Der Beitrag dieser Arbeit zeigt, dass dichte EGIs auch dann verwendet werden können, wenn der erste Schritt gemäß der linearen Kontrolltheorie nicht beobachtbar ist; leitet einen analytischen Ausdruck für die Hessian der Volumenbeschränkung des zweiten Schritts ab; und führt die neuartige Homothetizität ein, ein Kriterium, das es ermöglicht, die Leistung der Methode auf der Grundlage des endgültigen Polyeders anstelle des Zwischen-EGI zu bewerten.

Schließlich wurde ein neuer Ansatz für die Lagebestimmung (AD, *attitude determination*) aus Lichtkurven entwickelt. Er modelliert die Wahrscheinlichkeitsverteilung des Zustands als Gaußsche Mischung und setzt voraus, dass Form und Reflexionsgrad der Oberfläche bekannt sind. Jeder Kernel wird mit einem Unscented Kalman-Filter (UKF) aktualisiert, während die Größe der Mischung bei jedem Schritt anhand eines Nichtlinearitätsmaßes des Lichtkurvenmodells angepasst wird. Simulationen zeigen, dass dieser Filter, der für AD Adaptive Gaußsche Mischung Unscented Kalman-Filter (AGMUKF) genannt wird, den traditionellen UKF übertrifft. Darüber hinaus deuten die Ergebnisse darauf hin, dass der AGMUKF für AD das Potenzial hat, mit dem modernen Partikelfilter (PF) leistungsmäßig gleichzuziehen, aber dank seiner Fähigkeit, die Mischungsgröße anzupassen, zu geringeren

Rechenkosten.

Diese Arbeit schließt mit der Demonstration, dass die Beiträge dieser Arbeit zum Bereich der Objektcharakterisierung potenziell Bausteine für abgeleitete Anwendungsfälle darstellen können, wie z.B. die Gesundheitsüberwachung von operativen Satelliten.

Acknowledgements

First and foremost, I want to thank my academic supervisor Prof. Roger Förstner for his guidance and constructive feedback throughout all the stages of my Ph.D., despite the global pandemic. Thank you for your valuable insights and for the opportunity to pursue a doctorate under your supervision.

I also want to thank Prof. Carolin Frueh, for the effort put into assessing this work and for her invaluable feedback.

And thank you Prof. Gerdts, for having guided and organized the entire assessment process.

Next, my most sincere gratitude goes to my Airbus supervisor Dr. Jens Utzmann, who, during our three years working together, has spared no effort to produce the ideal environment so that I could focus on my research, and worry about nothing else. Thank you for your relentless support and your contagious enthusiasm.

My gratitude extends to my fellow Ph.D. student and office mate Guido Pedone. It has been a pleasure to work along his side, exchanging ideas, sharing know-how, and supporting each other during all this time. I have had a lot of fun.

A nivell personal, no tinc paraules per agrair el suport incondicional de la meva dona, que es va llançar a l'aventura de marxar de Barcelona per venir a viure a Alemanya. No seria on sóc si no fos per la seva valentia, la seva generositat i la seva fortalesa.

També vull agrair profundament als meus pares per haver-me encoratjat sempre a perseguir els meus somnis, i per ensenyar-me que la felicitat rau en fer de la millor manera allò que hom es proposi. Juntament amb la meva germana, sempre m'han fet costat i hi he pogut comptar. Per això, de nou, moltes gràcies!

Vull acabar donant les gràcies, també, a la resta de la meva família i amics, que han fet que els dies passin volant al llarg de tota aquesta aventura.

Contents

List of Tables	xiii
List of Figures	xv
List of Acronyms	xvii
List of Symbols	xxi
1. Introduction	1
1.1. Object Characterization	1
1.2. Research Questions	3
1.3. Specific Research Goals	3
1.4. Procedure and structure	4
2. Literature Review	5
2.1. Light Curve Simulation	5
2.2. Object Characterization Breakdown	7
2.2.1. Orbit Ephemeris	8
2.2.2. Attitude	8
2.2.3. Shape Characterization	10
2.2.4. Surface Materials Characterization	11
2.2.5. Simultaneous Characterization	12
2.2.6. Abstract Characterization	12
2.2.7. Summary	13
2.3. Observability of Object Characteristics	13
2.4. Object Characterization Derived Use Cases	15
2.5. The Scientific Contribution of This Thesis	16
3. Light Curve Model	19
3.1. Problem Definition	20
3.1.1. Observation Geometry	20
3.1.2. Object Shape and Size Representation	21
3.1.3. Attitude	21
3.1.4. Bidirectional Reflectance Distribution Function	23
3.1.5. The Path of Light	26
3.1.6. The Sensor	27
3.2. Motivation for a New Solution to the Shadows Problem	28
3.3. The Shadow Projection Algorithm	29
3.3.1. Projection onto a Plane	30
3.3.2. Lifting from the Plane	31

3.3.3. Polygon Clipping	32
3.3.4. Algorithm	33
3.4. Evaluation	37
3.5. Conclusion	39
4. Shape Determination	41
4.1. Problem Definition	41
4.1.1. Estimating the EGI from a Light Curve	42
4.1.2. Generating a Convex Polyhedron from an EGI	45
4.2. Homotheticity Distance	46
4.3. Derivatives of the Volume over the Supports	47
4.4. Observability w.r.t. the EGI Size	54
4.5. Simulation Scenarios	55
4.5.1. Test Shapes	55
4.5.2. Scenario Geometry	56
4.5.3. Light Curve Measurement	56
4.5.4. Observability Conditions	59
4.5.5. Noise Mitigation	59
4.5.6. Hessian of the Volume	59
4.6. Results and Discussion	59
4.6.1. Cube	59
4.6.2. Icosahedron	63
4.6.3. Cylinder	64
4.6.4. Winged Cuboid	65
4.6.5. General Discussion	66
4.7. Potential Extensions	67
4.7.1. Albedo Variegation	67
4.7.2. Non-Convexity	68
4.8. Conclusion	68
5. Attitude Determination	71
5.1. Problem Definition	72
5.1.1. Non-Linear Sequential Estimation	72
5.1.2. Attitude Dynamic Model	73
5.1.3. Light Curve Measurement Model	74
5.2. Motivation for a New Attitude Determination Method	75
5.2.1. The Non-Linearity Index	76
5.2.2. Non-Linearity Analysis of the Attitude Determination Problem	77
5.3. The Adaptive Gaussian Mixtures Unscented Kalman Filter	78
5.3.1. Probability Distribution	80
5.3.2. Filter Initialization	81
5.3.3. Propagation Step	81
5.3.4. Measurement Estimation	82
5.3.5. Refinement Step	83
5.3.6. Update step	84
5.3.7. Coarsening step	84
5.3.8. Filter Reset	85

5.4. Simulation Scenarios and Results	86
5.4.1. AGMUKF <i>vs.</i> UKF	86
5.4.2. Highly non-linear case	96
5.5. Conclusion	102
6. Use Cases	105
6.1. Object Characterization Pipeline	105
6.2. Frequency Analysis	107
6.3. Diagnosis of an Operational Satellite	109
6.4. Other Use Cases	112
7. Conclusion and Outlook	115
7.1. Summary	115
7.2. Outlook	116
7.3. Final Word	118
Bibliography	119
A. Framework	129
A.1. Special Perturbations Orbit Determination and Orbit Analysis Toolkit	129
A.2. Airbus Robotic Telescope	130
B. Summary of OC Methods from Light Curves in the Literature	131
C. Quaternion Algebra	133
D. Extended Attitude Concepts	137
D.1. Attitude Representations	137
D.1.1. Attitude as a Quaternion	137
D.1.2. Attitude as an Euler Axis-Angle Pair	138
D.1.3. Attitude as a Direction Cosine Matrix	138
D.1.4. Attitude as Generalized Rodrigues Parameters	138
D.2. Attitude Quaternion Derivative	139
D.3. Attitude Quaternion Given Two Vectors	140
D.4. Generalized Rodrigues Parameters Derivations	141
D.4.1. Taylor Expansion of the GRP	141
D.4.2. Change of the Reference Quaternion	141
E. Three Dimensional Objects	145
E.1. Extended Gaussian Images	145
E.2. Sphere Tessellation	146
F. The Unscented Transform	151

List of Tables

3.1. Shadow Projection test shape	38
4.1. Rank of the Gramian at different light curve and EGI sampling densities . . .	60
B.1. Summary of AD techniques	131
B.2. Summary of shape characterization techniques	131
B.3. Summary of material characterization techniques	132
B.4. Summary of combined characterization techniques	132
B.5. Summary of abstract characterization techniques	132

List of Figures

3.1.	Observation geometry of the Sun-object-observer triangle	20
3.2.	The observation geometry for a surface element	24
3.3.	2D illustration of the shadow projection algorithm	36
3.4.	Shadow Projection test shape	38
3.5.	Illumination scenario computed using Shadow Projection	40
4.1.	Change of volume of a polyhedron as a function of one of its supports	50
4.2.	Zoom on 2 nd order contribution to the volume versus support change	52
4.3.	Test shapes	56
4.4.	EGIs of the test shapes, based on a $d = 5$ quad-cube tessellation	57
4.5.	Simulated light curves	58
4.6.	Cube results, averaged over 64 Monte Carlo runs	60
4.7.	Original (black) versus recovered (red) shape for the cube	61
4.8.	Cube results, 64 Monte Carlo runs, $N_t = 1000$, different Hessian strategies	62
4.9.	Icosahedron results, averaged over 64 Monte Carlo runs	63
4.10.	Original (black) versus recovered (red) shape for the icosahedron	64
4.11.	Cylinder results, averaged over 64 Monte Carlo runs	65
4.12.	Original (black) versus recovered (red) shape for the cylinder.	65
4.13.	Original (black) versus recovered (red) shape for the winged cuboid	66
5.1.	NLI of one UKF run, based on the scenario from Section 5.4.1	78
5.2.	PDF of \hat{z}_k for a UKF run, based on the scenario from Section 5.4.1	79
5.3.	Diagram of the AGMUKF algorithm for AD	80
5.4.	Simulated object shape	87
5.5.	One noisy light curve of the Wetterer et al. (2009) scenario	89
5.6.	One UKF run on the Wetterer et al. (2009) scenario	90
5.7.	One AGMUKF run on the Wetterer et al. (2009) scenario	91
5.8.	LL over 100 MC UKF runs for the Wetterer et al. (2009) scenario	94
5.9.	LL over 100 MC AGMUKF runs for the Wetterer et al. (2009) scenario	94
5.10.	PDF of \hat{z}_k for a AGMUKF run, based on the Wetterer et al. (2009) scenario	95
5.11.	Kernels for one AGMUKF run of the Wetterer et al. (2009) scenario	96
5.12.	One noisy light curve of the Linares et al. (2014a) scenario	97
5.13.	One AGMUKF run on the Linares et al. (2014a) scenario	98
5.14.	LL over a single MC AGMUKF run for the Linares et al. (2014a) scenario	100
5.15.	NLI of one AGMUKF run for the Linares et al. (2014a) scenario	100
5.16.	Kernels for one AGMUKF run of the Linares et al. (2014a) scenario	101
6.1.	Flow-chart of the OC capabilities analysed in this thesis	106
6.2.	Image of the ISO satellite taken with ART.	107
6.3.	Light curve of the ISO satellite	108

List of Figures

6.4.	Lomb–Scargle periodogram of the detrended ISO light curve	108
6.5.	Phase plot of the detrended ISO light curve	108
6.6.	One noisy light curve of the diagnosis scenario	110
6.7.	LL over a single AGMUKF run for the diagnosis scenario	112
A.1.	The Airbus Robotic Telescope	130
E.1.	The quad-cube tessellation	148
E.2.	Spherical triangles involved in the quad-cube tessellation coordinate system .	148

List of Acronyms

- AD** Attitude Determination. 4, 8–10, 13, 17, 71, 73, 75, 77, 78, 80, 82, 86, 96, 102, 103, 105, 109, 111, 112, 116, 117, 129
- ADR** Active Debris Removal. 1, 2, 4, 15, 16, 18, 112, 113, 117
- AGMUKF** Adaptive Gaussian Mixtures Unscented Kalman Filter. 4, 17, 71, 72, 78, 80, 82, 85–88, 92, 93, 96, 99–103, 107, 109, 111, 112, 116, 117
- AIJ** AstroImageJ. 108
- ART** the Airbus Robotic Telescope. 56, 107, 109, 129, 130
- AS** Ashikhmin-Shirley. xxiii, 24, 75
- BFGS** Broyden–Fletcher–Goldfarb–Shanno. 59, 62
- BRDF** Bidirectional Reflecion Distribution Function. xxii, 6, 14, 23–25, 55, 75, 87, 109
- CA** Collision Avoidance. 15, 113, 117
- CAD** Computer Aided Design. 5, 6
- CCD** Charge-Coupled Device. xxi, xxii, 6, 16, 130
- CNN** Convolutional Neural Network. 9, 11, 13, 131, 132
- CT** Cook-Torrance. xxiii, 24, 25, 55, 75
- DCM** Direction Cosine Matrix. xxiii, 30, 137, 138
- DoF** Degree of Freedom. 13, 73, 137
- EGI** Extended Gaussian Image. xxv, 3, 4, 11, 12, 14, 17, 41–47, 54, 55, 59, 61, 63–68, 107, 109, 115–117, 131, 145–147, 149
- EKF** Extended Kalman Filter. 10, 12, 17, 71, 73, 78, 116, 131
- EM** Expection-Maximizaiton. 12, 101
- FFT** Fast Fourier Transform. 8
- FG** Fragmentation Analysis. 15
- FIM** Fisher Information Matrix. 13, 14
- FSPL** Free-Space Path Loss. 19

List of Acronyms

- GCRF** Geocentric Celestial Reference Frame. 21, 73, 111
- GEO** Geosynchronous Earth Orbit. 14, 109
- GM** Gaussian Mixture. xxii, xxiv, 10, 12, 17, 71, 78, 80, 81, 83, 84, 92, 99, 101, 116, 117, 132
- GPU** Graphics Processing Unit. 6, 29, 117
- GRP** Generalized Rodrigues Parameters. xxi–xxiii, 73, 74, 81, 85, 86, 92, 99, 138
- LEO** Low Earth Orbit. 14
- LL** Log-Likelihood. 92, 93, 101
- LoS** Line of Sight. 28
- MC** Monte Carlo. 60–62, 64, 66, 77, 88, 92, 93, 96, 100–102
- MEO** Mid Earth Orbit. 56
- ML** Machine Learning. 11–13, 16, 109, 112, 118, 132
- NLI** Non-Linearity Index. xxv, 76–78, 80, 83, 99, 101–103, 116
- NNLS** Non-Negative linear Least Squares. 43, 44, 54, 68
- OC** Object Characterization. 1–5, 7, 8, 13–16, 18, 19, 24, 27, 29, 105, 107, 109, 112, 113, 115–118
- OD** Orbit Determination. 2, 8, 13, 18, 71, 76, 78, 101, 102, 112, 113, 116, 117, 129
- PAB** Phase Angle Bisector. xxiv, 21, 24, 109
- PDF** Probability Density Function. 12, 17, 18, 71–73, 76–78, 80–85, 92, 93, 99, 101, 109, 111, 116
- PF** Particle Filter. 4, 10–12, 17, 18, 41, 71, 73, 78, 96, 102, 103, 116, 117, 131, 132
- PHD** Probability Hypothesis Density. 10
- PMF** Probability Mass Function. 80
- RE** Re-entry Analysis. 15, 113, 117
- RSO** Resident Space Object. xxiii, 1–3, 6–13, 15, 20, 68, 72, 102, 105, 109, 115, 117, 129
- SI** Interational System. 111
- SLR** Satellite Laser Ranging. 8, 131
- SNR** Signal to Noise Ratio. 56, 59, 61–64, 66, 68, 69

- SPOOK** the Special Perturbations Orbit determination and Orbit analysis toolKit. 129, 130
- SRUKF** Square Root Unscented Kalman Filter. 86
- SSA** Space Situational Awareness. 1
- SST** Space Surveillance and Tracking. 1–4, 10, 15, 16, 105, 109, 113, 115–118, 129
- TLE** Two-Line Element. 13, 86
- UKF** Unscented Kalman Filter. 10, 12, 13, 16, 17, 71, 74, 77, 78, 80–82, 84–88, 92, 93, 96, 102, 107, 116, 117, 131, 132
- USKF** Unscented-Schmidt Kalman Filter. 14
- UT** Unscented Transform. xxiv, xxv, 4, 71, 76–78, 82, 86, 88, 96, 102
- UTC** Coordinated Universal Time. 56, 111
- WLS** Weighted Least Squares. 9

List of Symbols

This is a list of mathematical symbols used throughout this thesis. In general, each symbol has a unique meaning. However, some conflicts arise with most commonly used symbols. For these cases, this list contains the most generally used meaning. When abuse of notation is inevitable because finding an alternative, non-conflicting symbol would overcomplicate mathematical derivations, minority meanings are specified explicitly in the text. These alternative interpretations are typically contained in one or a few consecutive paragraphs, giving little room for confusion. For example, v represents *velocity*. However, in some places it can be used as a generic vector: *given any vector v , its Euclidean norm is $\|v\|$* ; after this statement, v means again *velocity*.

A denotes area.

a is the denominator parameter in the quaternion to Generalized Rodrigues Parameters (GRP) conversion.

\mathcal{A} denotes area that is both visible by an observer, and illuminated by the Sun.

\mathcal{A}_r is the projection of \mathcal{A} on the plane perpendicular to u_r .

\mathbf{a} represents albedo-area product.

B represents the body frame. As a super index, it is the frame in which the vector is represented.

B_λ represents a wavelength bandwidth.

c is the speed of light in vacuum.

C_a is the number of electrons counted by a Charge-Coupled Device (CCD) sensor over some finite time.

d is the number of divisions used to discretize each side of a quad-cube tessellation.

δ is the fraction of diffuse reflection, a.k.a. diffuse coefficient.

d_r is the absolute distance between the object and the observer, a.k.a. range.

d_s is the absolute distance between the object and the Sun.

e is the rotation axis of the Euler angle-axis representation of attitude. It can also be the edge of a polyhedron.

E_a is the irradiance from a distant object, received by an observer.

$E_{a,0}$ is the irradiance from the Vega star, received by an observer.

$E_{a,\lambda}$ is the spectral irradiance from a distant object, received by an observer.

List of Symbols

- $E_{r,\lambda}$ is the spectral irradiance from a distant surface, received by an observer.
- $E_{s,\lambda}$ is the spectral irradiance received from the Sun.
- f is the scale parameter in the quaternion to Generalized Rodrigues Parameters (GRP) conversion; it can also represent the dynamic model of a sequential estimation problem.
- $F^{(k)}$ is the k -th facet of a polyhedron.
- F_0 is the specular reflectance of a material.
- F_a is the electron flux detected on a Charge-Coupled Device (CCD) sensor.
- $F_{a,\lambda}$ is the spectral electron flux detected on a Charge-Coupled Device (CCD) sensor.
- f_r represents a Bidirectional Reflectance Distribution Function (BRDF).
- H is homotheticity distance.
- h is Planck's constant. It is used to represent the measurement model of a sequential estimation problem, too.
- \mathfrak{h} represents the support(s) of a polyhedron face(s).
- \bar{H} is relative homotheticity distance.
- I represents the inertial frame. As a super index, it is the frame in which the vector is represented.
- $\Im(\cdot)$ returns the imaginary part of a complex number, or the hyper-imaginary part (as a 3-dimensional vector) of a quaternion.
- \mathcal{I}_l is the map that converts the indices of the facets of a polyhedron as seen by its l -th facet, to the general facet indexing system \mathfrak{L} .
- I_G is the inertia tensor of a rigid solid.
- I_p is a matrix that transforms a 3-dimensional vector into a 2-dimensional one by cropping the third component.
- i, j, k, l are arbitrary indices.
- $\mathbf{i}, \mathbf{j}, \mathbf{k}$ are the first, second and third hyper-imaginary numbers, respectively.
- k_B is Boltzmann's constant.
- \mathcal{L} is the light curve vector.
- \mathfrak{L} is the subset of natural numbers used to index the facets of a polyhedron.
- L_r is the reflected radiance.
- $L_{r,\lambda}$ is the reflected spectral radiance.
- $L_{s,\lambda}$ is the spectral radiance emitted by the Sun.
- M is the number of kernels in a Gaussian Mixture (GM).
- m is used as a generic natural number in some derivations; in particular, it is the number of observable states.

- \mathbf{m} is the facet slope parameter of the Cook-Torrance (CT) model.
- m_A is apparent magnitude.
- m_I is instrumental magnitude.
- N is a generic natural number.
- n is used as a generic natural number in some derivations; in particular it is the number of hidden states. In some context, it can be a vector normal to a polyhedron facet.
- \mathcal{N} is the Normal distribution.
- normals* is an array of column normal vectors.
- N_a is the number of cells of a spheric tessellation.
- N_f is the number of facets of a polyhedron.
- N_{MC} is the number of Monte Carlo iterations.
- N_t is the number of measurements in a light curve.
- n_u is the first anisotropy index of the Ashikhmin-Shirley (AS) model.
- n_v is the second anisotropy index of the Ashikhmin-Shirley (AS) model.
- O represents a Resident Space Object (RSO) as a closed 3-dimensional volume. It can also denote the Landau notation.
- P is the covariance of a statistical distribution.
- \mathcal{P} is a polyhedron.
- $p_{a,f}$ is the Generalized Rodrigues Parameters (GRP) representation of attitude. When it is clear by context, sometimes it is represented simply as p .
- p_r is a probability value.
- Q_k is the process noise covariance matrix at time step k .
- q is a quaternion expressed as a 4-dimensional vector.
- q_i, q_j, q_k are the first, second and third hyper-imaginary parts of quaternion \mathbf{q} , respectively.
- q_r is the real part of quaternion \mathbf{q} . it is a short-hand notation for $\Re(\mathbf{q})$.
- \mathbf{q} is a quaternion.
- \mathbf{q}_A^B is an attitude quaternion that transforms a generic vector expressed in frame A to frame B .
- $\tilde{\mathbf{q}}$ is a reference quaternion.
- \mathfrak{q} is the hyper-imaginary part of quaternion \mathbf{q} . It is a short-hand notation for $\Im(\mathbf{q})$.
- $\Re(\cdot)$ returns the real part of a complex number or a quaternion.
- R_A^B is a Direction Cosine Matrix (DCM) that transforms a generic vector expressed in frame A to frame B .

List of Symbols

- R_d is the bidirectional diffuse reflectance.
- R_s is the bidirectional specular reflectance.
- R_k is the measurement noise covariance matrix at time step k .
- S is a surface.
- \mathcal{S} is the scattering matrix.
- $\hat{\mathcal{S}}$ is the scattering matrix without albedo contribution.
- S_s is the surface of the Sun.
- T is torque.
- t symbolizes time.
- \mathcal{T} is an operator that scales and translates a polyhedron.
- t_e is the integration time during which an optical sensor is exposed, a.k.a. exposure time.
- T_s is the temperature of the Sun.
- u_a is a unit vector along the normal of a telescope aperture.
- u_b is the unit vector along the bisection of u_s and u_r , a.k.a. Phase Angle Bisector (PAB).
- u_n is a unit vector along the normal of a surface element.
- u_r is the unit vector that points from object to the observer.
- u_s is the unit vector that points from the object to the Sun.
- u_u is a unit vector along the first anisotropy direction of a surface element.
- u_v is a unit vector along the second anisotropy direction of a surface element.
- V is the volume of a 3-dimensional object.
- v is a generic vector.
- $v_i^{(k)}$ is the i -th vertex of the k -th facet of a polyhedron.
- w represents the weights of a Gaussian Mixture (GM).
- X is a random variable that represents a hidden state.
- x is a realization of X , but can also represent the first axis of a reference frame.
- \mathcal{X} is a sigma cloud in the Unscented Transform (UT).
- y is the second axis of a reference frame.
- Z is a random variable that represents an observable state.
- z is a realization of Z , but can also represent the third axis of a reference frame.
- α is a regularization parameter, or the spread parameter of the Unscented Transform (UT), depending on the context.

β is the vertical coordinate of the quad-cube tessellation, or a parameter of the Unscented Transform (UT), depending on the context.

γ is the horizontal of the quad-cube tessellation.

ε represents a small number in various contexts.

Θ is the hyper-complex part of the left-hand side quaternion multiplication matrix.

θ_a is the angle between aperture normal and line of sight.

θ_b is the polar coordinate of u_b in facet coordinate system.

θ_r is the polar coordinate of u_r in facet coordinate system.

θ_s is the polar coordinate of u_s in facet coordinate system.

λ can symbolize wavelength, or one of the parameters of the Unscented Transform (UT).

μ is the mean of a statistical distribution.

Ξ is the hyper-complex part of the right-hand side quaternion multiplication matrix.

π is the ratio of the circumference of a circle over its diameter.

π_u represents the projection of a 3-dimensional vector onto a plane orthogonal to direction u .

π_{uv} represents the lifting of a point projected onto the plane orthogonal to u , to another plane orthogonal to v .

ρ is the diffuse reflectance of a material.

ϱ is the closeness function used to construct sphere tessellations for Extended Gaussian Images (EGIs).

τ is the overall transmissivity of an optical system.

$\Phi_{a,\lambda}$ is the spectral flux of light that enters a telescope.

ϕ represents the Non-Linearity Index (NLI).

ϕ_b is the azimuth coordinate of u_b in facet coordinate system.

ϕ_r is the azimuth coordinate of u_r in facet coordinate system.

ϕ_s is the azimuth coordinate of u_s in facet coordinate system.

φ is azimuth or longitude in a spherical coordinate system.

φ_b is the angle between u_s and u_r , a.k.a. phase angle.

ψ is the rotation angle of the Euler angle-axis representation of attitude.

ω is angular velocity.

$\omega_{A/B}^C$ is angular velocity of frame A relative to frame B , expressed in frame C .

Ω_r is the solid angle of the object as seen by the observer.

Ω_s is the solid angle of the illumination source as seen by the object.

List of Symbols

δO represents the 3-dimensional outer surface of O .

δP represents the set of facets of polyhedron P .

ϑ is elevation or latitude in a spherical coordinate system.

\mathbb{A} is the set of unit quaternions with non-negative real part.

\mathbb{H} represents quaternion algebra.

\mathbb{N} represents the entire set of Natural numbers.

$\mathbb{N}_{>0}$ represents the set of strictly positive Natural numbers, *i.e.* $(0, \infty)$.

\mathbb{R} represents the entire set of Real numbers.

$\mathbb{R}_{\geq 0}$ represents the set of non-negative Real numbers, *i.e.* $[0, \infty)$.

$\mathbb{R}_{>0}$ represents the set of strictly positive Real numbers, *i.e.* $(0, \infty)$.

$|\cdot|$ returns the absolute value of a number (including quaternions).

$\|\cdot\|$ returns the Euclidean norm of a vector.

$[\cdot]_{\times}$ is the cross product matrix.

$[\cdot]_{\odot}^{+}$ is the left-hand side quaternion multiplication matrix.

$[\cdot]_{\odot}^{-}$ is the right-hand side quaternion multiplication matrix.

\angle returns the angle between two vectors.

\times represents the cross product between two 3-dimensional vectors.

\odot represents quaternion multiplication using Hamilton convention.

lift is an operator that lifts a point from a projection plane to a higher dimensional space, given a triangle in that space, as well as its projection.

vset is an operator that returns the set of vertices of a polygon, ordered around its normal direction.

1. Introduction

This thesis is a scientific contribution to the field of Object Characterization (OC) within the scope of Space Surveillance and Tracking (SST). The aim of SST systems is to *observe* the population of objects that orbit the Earth, a.k.a. Resident Space Objects (RSOs), and then exploit this information in the form of services and applications (*i.e.* use cases) (ESA, 2022). In turn, OC consists in the estimation of different characteristics of these same RSOs. There is a wide range of such characteristics, from physical properties of the system to abstract attributes linked to its functionality. Examples of the former are shape, attitude and orbit; for an artificial satellite, operational status, mission goal and ownership are examples of the latter. This dissertation contributes to science via the analysis, improvement, and original development of methods and algorithms used to estimate characteristics of RSOs. In particular, it researches the problems of attitude and shape determination from light curves (*i.e.* time-series of telescope measurements of the brightness of a RSO), as well as light curve simulation itself.

This chapter introduces the topic of this thesis in further detail. Section 1.1 below motivates and justifies this work, highlighting the key role of OC in preserving the usability of the space environment around our planet. Next, Section 1.2 distils these motivations into the research questions that guide this thesis. This is followed by Section 1.3, which refines these questions into a set of specific goals. Finally, Section 1.4 sets out the structure for the remaining of this document.

1.1. Object Characterization

With the recent proliferation of mega-constellations of artificial objects, the number of RSOs is expected to increase substantially in the years to come (S. D. O. ESA, 2019). If no actions are taken, the most profitable orbit regimes will soon become impracticable. This would impact space-derived services that are pillars to the current global economy, such as communications, navigation, weather forecasting and climate change monitoring (OECD, 2019).

The first and most immediate challenge associated to this problem is the lack of exact knowledge on how many RSOs orbit the Earth, and of which nature (Aglietti, 2020). Without this key piece of information, further mitigation activities such as Active Debris Removal (ADR) or precise collision risk assessment are not possible. SST is the branch of Space Situational Awareness (SSA) that aims to catalogue RSOs, including their up-to-date orbital ephemerides and any other characteristics (*e.g.* bus type or origin) (SatCen, 2023). As a service-enabling technology, SST has become a valuable product in the space industry, and it is essential for efficient space traffic management.

1. Introduction

To characterize RSOs, SST providers use measurement techniques that do not require the participation of the object being observed—*i.e.* optical sensors or telescopes and Radar/Lidar-based systems. Such measurements contain information not only regarding the instantaneous location of the object and its momentum—from which its orbit can be inferred—but also about other characteristics, such as its orientation, size, shape and reflective surface properties. Therefore, by extension these measurements can even give clues regarding more abstract properties, such as the capabilities of its propulsive systems, the power output of its solar panels, or even its origin or the goals of its mission. Space¹ Object Characterization (OC) is the sub-field of SST that studies which characteristics can be extracted from these measurements and how. Although orbit ephemerides can be considered as object characteristics themselves, Orbit Determination (OD) is often treated as a sub-field on its own, so that OC typically deals with any observable characteristics other than the object’s orbit.

The characterization of RSOs has seen an increase of scientific interest during the last two decades. This can be justified naively with the tenet: *the more is known about all the objects orbiting the Earth, the better*. In particular, OC has the potential to provide information with measurable and significant value. For example, during the design of an ADR mission that targets the removal of a piece of space debris, knowledge on its attitude state or its shape reduces the uncertainty associated to the mission success, which leads to a narrower constraining of the risks, reducing the part of the mission budget dedicated to contingencies. For a more thorough exploration of OC-enabled use cases, the reader may refer to Section 2.4 of the Literature Review Chapter.

Despite recent advancements, OC as a field of science still has potential for growth. This dissertation looks into the problem of Object Characterization based on use-cases, and proposes improvements and new solutions to classical OC problems.

In particular, this thesis focuses on OC from light curves. These are time series of brightness measurements obtained from optical telescopes, for the common scenario when the image obtained is non-resolved: the object appears as a point², so that its features are not directly recognizable. Figure 6.2 is one such image of the ISO satellite. Although it does not tell much about the object characteristics, the evolution of measured light intensity (and potentially polarization and/or spectral content) along a series of such consecutive images can tell much more (see Figure 6.3). To illustrate, a sphere and a cube with equal surface properties, and equal conditions (same orbit, same cross section area at observation time, etc.) would produce the same indistinguishable point on a telescope image. However, if both rotate at the same rate, and a light curve is recorded, while the signal from the sphere would be flat, only influenced by the relative movement of the Sun, object and observer, the cube would produce a periodic light curve where every full rotation around itself would correspond to four periods. Light-curve-based OC exploits this intuition in a formal way to extract quantifiable characteristics of the object.

¹The *Space* qualifier is omitted from here onward, since it is redundant in the context of this thesis—*i.e.* only Object Characterization (OC) is used for the rest of this document.

²Or a streak, if there is relative movement between the object and the field of view.

1.2. Research Questions

The aim of this research work is to enhance the understanding on what is achievable in terms of Object Characterization from light curves and derived use cases.

To guide these aims, the following specific research questions have been posed:

- Which Resident Space Object characteristics can be deduced from light curves and how?
- Which observability conditions are necessary to do so?
- Which SST use cases does direct Object Characterization enable?

These questions are but a summary to which the reader can refer to for a brief understanding of the essence and overall scope of this thesis. However, they are too broad to serve as specific objectives. This is, they need to be broken down through a state of the art review of the field. Therefore, the specific objectives identified in the Literature Review Chapter are summarized in the following section.

1.3. Specific Research Goals

This is a list of specific research goals of this thesis, which have been derived from the three research questions above through the Literature Review in Chapter 2. Each of these goals corresponds to particular scientific contribution of this thesis. They have been grouped in four main sub-groups, according to the specific OC sub-field.

- Light curve modelling
 - Optimize the self-shadowing problem, inherent in light curve modelling, for the particular case of a non-convex object whose surface can be described with a small number of flat facets.
- Shape determination from light curves using the Extended Gaussian Image (EGI) plus Minkowski minimization algorithm
 - Propose a new metric that compares how close two polyhedrons are from being mutually homothetic.
 - Assess the performance of the algorithm for different variants and tuning parameters.
 - Derive the Hessian of the volume of a polyhedron relative to its supports, so that it can be used within Minkowski minimization.
 - Investigate the influence of the non-negativity constraint on the observability of the density of the EGI tessellation w.r.t. the number of samples in the light curve, including sensor noise.
- Sequential attitude estimation from light curves
 - Characterize the non-linearity of the sequential attitude estimation problem from light curves.

1. Introduction

- Adapt the Adaptive Gaussian Mixtures Unscented Kalman Filter (AGMUKF) to the attitude estimation problem.
- Assess its potential as a competitor to the state-of-the-art Particle Filter (PF).
- Use cases
 - Demonstrate that the other contributions of this thesis have the potential to improve OC-enabled use cases such as ADR and health monitoring of operational satellites.

These goals are discussed again in more depth 1) at the end of the Literature Review Chapter, where they have been derived; 2) at the end of this thesis, as part of the Conclusion and Outlook Chapter; and 3) individually in their respective core chapters of this thesis (see the following section).

1.4. Procedure and structure

The content of the rest of this thesis is summarized here. First, Chapter 2 investigates the literature on the topic of the three research questions from Section 1.2, and sets the specific and exact research goals of this thesis (summarized above in Section 1.3). Next, Chapter 3 describes the light curve model, with the aim to lay out the concepts, foundations and mathematical tools necessary to then develop Object Characterization (OC). Its main contribution is the development of a new rendering algorithm for light curves. Chapter 4 revisits, verifies and expands the light-curve-based shape determination algorithm based on EGIs and Minkowski minimization. Next, Chapter 5 develops the new AGMUKF for Attitude Determination (AD), a non-linear sequential solver that estimates attitude from light curves. It is followed by Chapter 6, which wraps the results of the previous chapters by focusing on the third research question—it explains how the developments of this thesis can be exploited to enable or enhance SST-related use cases. This dissertation is closed by Chapter 7 with a summary of the new contributions of this thesis to science, as well as the recommended work for future research.

The reader will find the bibliography sources used throughout this document at the end, followed by the Appendices. Appendix A summarizes the software and hardware framework that has been used to support the scientific work of this thesis. Then, Appendix B contains a quick reference of the methods highlighted in the Literature Review Chapter. Afterwards, Appendix C contains necessary properties of quaternion algebra, while Appendix D describes attitude representations and includes attitude-related derivations. Next, Appendix E defines the concept of EGI and describes the quad-cube tessellation of the sphere. Finally Appendix F describes the Unscented Transform.

2. Literature Review

Having established the goals of this thesis, this chapter explores the available literature on the state-of-the-art of space Object Characterization (OC) from light curves. First, Section 2.1 explores the state-of-the-art of light curve simulation, a process that is key to properly understand how OC itself works. Then, the research objectives themselves are studied. This is done from the point of view of the three research questions. Therefore, each of the three sections that follow (Sections 2.2 to 2.4) focus on each question, in the order they have been presented in Section 1.2. Finally, Section 2.5 closes the Literature Review with the identification of niches in the current literature of light-curve-based OC—*i.e.* where this thesis can contribute. As such, this last section justifies the specific research goals of this work from Section 1.3.

2.1. Light Curve Simulation

The first point that is relevant in this literature review is how light curves are simulated. OC is the *inverse problem* to light curve simulation or modelling, which is analogously named the *forward problem*. As such, understanding how object characteristics affect the outcome of optical measurements—joined together in the form of light curves—is a first step towards answering the first research question of this dissertation.

Furthermore, light curve simulation is an excellent tool to validate OC algorithms: it provides the ground truth to the problem, allowing the researcher to cheaply and quickly explore varied sets of scenarios, as well as to perform Monte Carlo simulations.

The problem of light curve simulation is well understood in general, and several works have dealt with it. One possible approach is to simulate the whole observation chain: an optical telescope image is simulated, which can then be processed to extract the brightness measurement of the object of interest. This is the case of *e.g.* Watanabe et al. (2017)—they use the software Blender to generate the image of the object, plus the modelling tool SOAPY to perturb it with realistic sensor and atmospheric effects. Another example is the high-fidelity image generation tool developed by Lockheed Martin and the Rochester Institute of Technology (Bennett et al., 2014), where they use Computer Aided Design (CAD) models, and adapt the radiometric model DIRSIG (Schott et al., 1999) to incorporate texture and polarization properties of reflected light. An alternative option is to skip the image generation step, and directly estimate the total amount of light received at the sensor. This is *e.g.* the case of Willison (2016). Another comprehensive light curve simulation pipeline that also skips image generation can be found in the doctoral dissertation of Fan (2020). There are several parts that need to be addressed when considering a light curve model.

- The first factor is the observation geometry or the relative position of the three agents involved: the Sun (illumination source), the object (upon whose surfaces the light

2. Literature Review

reflects) and the observer (which receives the reflected light and records its properties). This is a celestial mechanics problem where the orbits of the object (and possibly the observer, if not ground-based) around the Earth, as well as that of the Earth around the Sun, need to be known. This is well established knowledge in the field of physics, and is widely covered by any modern or classical textbook on the topic, such as Vallado et al. (2013).

- The next factor is the attitude state of the object. Attitude dynamics and kinematics are textbook material, too (*e.g.* Titterton et al. (2004)). Environmental torques that may perturb the attitude of a RSO have been explored (Wetterer et al., 2014), as well as how they couple with orbital dynamics (Früh et al., 2014).
- The next point relevant to light curve simulation is how the object shape is modelled. The field of CAD is mature enough—three dimensional objects are often modelled as a mesh of flat facets connected along their edges and vertices. (Sarkar, 2017)
- Then, there is the topic of modelling the reflection phenomena of light bouncing off the surfaces of the object. Although the physical and quantum principles of reflection and refraction are well understood (Hecht, 2002), such fundamental laws are unfit to be implemented on a system where computation resources are realistic—it would require modelling every atom in the material, which is unfeasible and intractable with a small, finite set of parameters. Thus, Nicodemus (1965) introduced what is known as the Bidirectional Reflectance Distribution Function (BRDF), a family of functions that can be used to simplify this problem, while still preserving the ability to model arbitrarily complex reflection patterns. There are plenty of BRDF models available in the literature. The simplest one uses the diffuse Lambertian model (Hecht, 2002). In order of complexity, then there are those that consider specular and diffuse effects (Cook et al., 1982), and even include anisotropic components (Ashikhmin et al., 2000). These are still tunable with a handful of parameters. Even more complex BRDFs can model polarization effects, *e.g.* using Mueller matrices (Pasqual et al., 2017), or with generalized Gaussian distributions (Renhorn et al., 2015). Finally, even more detail can be achieved when using look-up tables that empirically define the BRDF as a function of its parameters (Willison et al., 2015). Henderson (2014) presents a comprehensive comparison of different BRDF models for light curve simulation.
- Another factor to consider is the sensor model. This is, how the telescope captures the light and outputs it as an image (a two-dimensional array of pixels). For the common case of Charge-Coupled Device (CCD) detectors, these type of models are well understood. The Handbook of CCD Astronomy is one such example (Tobergte et al., 2013). It details how the light is detected using a CCD telescope, and which are the main perturbation effects. Früh (2015) or Coder et al. (2017) are but two examples where CCD models are used in the scope of characterizing RSOs from light curves.
- The final item that should be taken into account is the problem of rendering. This is, to determine the path of the light when traversing the shape geometry of the object, so that each photon is perturbed by the surfaces it interacted with. This is a task that can be extremely expensive—high fidelity approaches require advanced ray tracing techniques and cutting edge Graphics Processing Units (GPUs). In the field of light curve simulation, a wide range of approaches has been used. Piergentili et al. (2017)

use the simplest approach possible, which is to neglect self-shadowing for non-convex objects. They argue that this leads to the peaks in the light curve differing from reality the most, but that this is not a critical hindrance in their OC application. Früh et al. (2014) consider self-shadowing effects by tessellating the surface of the object, and then using ray tracing to see which rays originated in each cell of the object can reach the illumination source. They conclude that the thinner the tessellation, the higher the precision but the more expensive the computation. They introduce the concept of shadow map, that can be used as an averaged solution to the ray tracing output, which is more efficient to compute. Ray tracing is used in Fan (2020). Another approach is the one developed by Wang et al. (2018). It consists in using objects determined by small triangular meshes, so that each triangle can be approximated as a vertex. Then, each of these vertices is checked against the other facets, to determine if the facet it represents is hidden or not. Finally, the most detailed approach is to use actual rendering software, such as Blender in Watanabe et al. (2017). The disadvantage of this approach is the additional complexity of seamlessly integrating external tools to OC frameworks.

This is a summary of the aspects to be considered in a light curve model. Each of these points is covered in Chapter 3, which establishes the light curve model used within this dissertation.

2.2. Object Characterization Breakdown

This section breaks the first research question (*which Resident Space Object characteristics can be deduced from light curves and how?*) down into individual characteristics of a RSO. A good starting point is the work of Cox et al. (2016), which introduces the Space Object Ontology as an attempt to describe all RSO characteristics in a single framework. Some included characteristics are:

- whether the RSO is natural or artificial,
- functions and capabilities,
- orbit and manoeuvres,
- attitude,
- object features (shape, size and mass),
- object identifier,
- involved entities (*e.g.* owner, operator, etc.),

among others. Object Characterization (OC) includes, thus, any piece of information that can be associated to an object in any way. To approach this first research question, first it is necessary to classify and organize the characteristics themselves.

The first grouping that can be defined is whether the characteristic represents a physical or an abstract property of the object. A physical characteristic is such that represents an aspect of the physical nature of the object, such as its position in space and time, its attitude, or the composition of the object materials. An abstract property, instead, is a piece

2. Literature Review

of information not related to a physical aspect, such as the object identification according to some naming convention, the mission objective or the identity of the owner of the object. Measurements are inherently physical. Therefore, intuitively it shall be easier to obtain physical characteristics of the object from them. Abstract characteristics, on the contrary, may require further effort or subsequent analysis of the physical ones.

The rest of this section gives an overview of which characteristics have been estimated from light curves in the past, and how. The first and most obvious physical characteristic of a RSO is its position and velocity in time—*i.e.* its orbit. Other physical characteristics that typically influence the light curves directly are the attitude of the object, its size, shape and reflective properties of its surfaces. Orbit and attitude are a time-variant status, while the rest are of a more permanent nature—while they might change over time due to *e.g.* degradation, they are unlikely to change over short periods.

2.2.1. Orbit Ephemeris

Orbit Determination (OD) (velocity and position or orbital elements) is itself a well established and consolidated field, covered by textbooks such as Vallado et al. (2013) or Wiesel (2010). Although there is still new research going on around this topic (Schiemenz et al., 2020), this thesis focuses on object characteristics other than the orbit of the object, since these fields are comparatively less developed.

2.2.2. Attitude

Chronologically, initial OC research focused on full or partial Attitude Determination (AD). Full determination refers to estimating the entire attitude state and angular rates of an object, while partial refers to obtaining less information: *e.g.* only the rotation period, but not the direction of the spin axis.

AD methods can be further divided into two types, those that require a-priori knowledge on the shape, size and reflective properties of the object, and those that do not (Hall et al., 2014).

Partial, shape-independent methods

Starting with the shape-independent¹, partial AD methods, their state is relatively mature in the literature. Linder et al. (2015) present a compendium of six different methods to extract the apparent period of a light curve, both for cases when the light curve is equispaced and when it is not. Most of the methods they assess are based on the Fourier transform, namely the Fast Fourier Transform (FFT), various types of periodogram and Welch’s method. They additionally study the epoch folding method, which uses different candidate periods to fit an artificial pulse to the folded light curve (Larsson, 1996). To validate these methods, Linder et al. (2015) use apparent periods obtained from objects with retro-reflectors using Satellite Laser Ranging (SLR). Their results surmise that, although none of these methods has 100 %

¹Here shape-independent is abusing the notation to imply independence not only from shape, but also from reflective properties and size—*i.e.* from any non-dynamic property of the object.

success rate (*e.g.* fractions of the actual period may be obtained), cross-checking between them often allows correct extraction of apparent periods for stably spinning objects.

Once they have established the spin period of a light curve, if any, Linder et al. (2015) additionally obtain phase plots of the light curve—*i.e.* the light curve is plotted against the residual of time divided by period. A phase plot can be thought of as an intermediate partial measurement and/or characteristic of the object, since it contains the apparent rotation period, but might encode information pertinent to other characteristics, such as shape or reflective properties.

Full, shape-independent methods

There have been some attempts to extract full attitude state from the object’s light curve, without knowledge on its shape, size or reflective properties. The earliest work applied to artificial objects is that of Hall et al. (2006). They use the epoch method (Magnusson, 1986), originally designed to determine the spin inertial axis and period of asteroids rotating around a principle axis. It exploits the fact that the period observed on the light curve of a rotating object differs slightly from the inertial rotation period, as a function of the relative movement between the object, the observer and the Sun, as well as the orientation of the spin axis. It uses Weighted Least Squares (WLS) to fit the spin axis and sidereal period to variations over time of the apparent period, extracted by piece-wise Fourier analysis of the light curve. This method, however, is severely limited on the observability conditions of the object. For on-ground observers, it often provides two equally probable distinct solutions for the spin axis (Somers, 2011).

Hall et al. (2014) propose a potential expansion to the epoch method, which could deal with precessing objects—*i.e.* RSOs with axial symmetry that rotate around a non-principal axis. They propose a double-Fourier analysis of the light curve, in which the main and precessing frequencies could be identified. There has been no implementation of such case, neither on simulated nor on real measurements, however. Another partial, more recent AD approach based on Convolutional Neural Networks (CNNs) has been proposed by Badura et al. (2022) to provide attitude status classification (stable, rotating, tumbling, etc.).

Full, shape-dependent methods

Finally, there is plenty of literature regarding techniques to extract the full attitude state from passive measurements, when there is *a priori* information on the object shape, size and reflective properties.

Davis et al. (1957) were some of the first to try to estimate the attitude of an artificial spinning satellite with known cylindrical shape. They proposed observing the satellite in question with an optical telescope, so that specular reflections would be caught. Specular reflections constrict the orientation of the satellite to a subset of all possible orientations, as a function of the observer-object-Sun triangle. From enough observations, given the known shape and reflective properties of the object, it is possible to determine the spinning axis of said object. Their approach was further expanded to account for diffuse reflections (Giese, 1963) or to exploit maximum/minimum ratio of observed brightness (Williams, 1969), which simplified the acquisition of the measurement data. These methods, however, are particular

2. Literature Review

to only cylinder shapes, a small subset of all the possible shapes a RSO could have (e.g. cuboid, winged or more complex).

Another family of methods is the sequential estimation of RSO attitude given a known, but arbitrary shape/reflective model. These all need an initial estimation of attitude, although the most advanced versions can deal with high initial uncertainty. Wetterer et al. (2009) pioneer the use of the Unscented Kalman Filter (UKF) to estimate the attitude and the angular rates of RSOs from light curves. Holzinger et al. (2012), Holzinger et al. (2014), and Linares et al. (2014a) improve upon the inability of the Extended Kalman Filter (EKF) and the UKF to capture the high-non-linearities of this problem by using the Particle Filter (PF) instead. While Linares et al. (2014a) estimate the attitude of uncontrolled objects, Holzinger et al. (2012) and Holzinger et al. (2014) do so for agile, controlled objects with shape model uncertainty. Coder et al. (2015) improve the computational tractability of the agile scenario with the Marginalized Particle Filter: a Singer process noise model to describe unknown manoeuvring torques justifies the linearization of attitude rates. Coder et al. (2017) improve this approach with an extended noise model of the measurement process, which allows to initialize the filter with bigger state uncertainty. Real attitude ephemeris of the Hubble Space Telescope have been used to validate this implementation (Coder et al., 2018). Regarding uncontrolled objects, Bernard et al. (2018) compare the performance of the EKF against a bootstrapped Particle Filter. They identify that the EKF diverges because of the extreme non-linearity of the measurement model. Du et al. (2018) show that attitude estimation of uncontrolled objects can be improved w.r.t. the PF approach by using an Unscented Particle Filter: particles become kernels, which are propagated with a bank of UKFs. Finally, Burton et al. (2021) define pseudo-measurements as incident-reflecting directions pairs, and use them to explore two extensions to the EKF for AD. In one, they filter all unfeasible pseudo-measurements out, before feeding the surviving ones to an EKF. In the other one, they use a Gaussian Mixture (GM) to represent the (attitude) state space by wrapping the EKF core with a Probability Hypothesis Density (PHD).

A different alternative is that proposed by Piergentili et al. (2017). They treat the attitude characterisation problem as a batch optimization one. They employ heuristic optimization to find the attitude state at a given instant, such that the corresponding simulated light curve best fits the measured one in the weighted least squares sense. They assume known shape and reflective parameters.

2.2.3. Shape Characterization

The way the shape/size and reflective properties of an object affect its light curve are often strongly coupled at any range of the electromagnetic spectrum. For example, a small-area, high-albedo plate could reflect as much Sun light as a wider, yet low-albedo one in the same orbit. Thus, unless specific information about the object is available, which could help the SST agents discern these two factors, their effects on measurements are often indistinguishable from one another. Often, *shape characterization* actually refers to the characterization of some abstract quantity that contains both the effects of shape and reflective properties, but from which neither can be extracted individually. The rest of this subsection describes techniques available in the literature to extract shape (sometimes coupled with surface properties) from light curves.

The first family of shape-recovery methods is the one based on the method developed by Kaasalainen et al. (2001). They recover asteroid shapes in a two-step process: first, they extract an Extended Gaussian Image (EGI) from light curve data; then, they recover the shape itself using the work of Little (1985), which exploits Minkowski’s theorem (Minkowski, 1897). Calef et al. (2006) port this process to the case of artificial objects, exploring the benefits of ridge regularization. Bradley et al. (2014) apply this method to simulated non-convex objects. More recently, Friedman et al. (2019) have further explored the observability of this method based on the Gramian of the measurement model. Furthermore Fan et al. (2019) and Fan et al. (2020) extend the method to use two consecutive light curves, an initialization one where full observability is necessary, plus a refinement one that does not need to contain all the shape information. They prove that this modification is more robust than the one-big-light-curve approach when measurement noise is present. Finally, Fan et al. (2021) further modify the method to incorporate an arbitrary number of light curves. After initializing the filter with the two-light-curves method, the additional information of each new light curve is fused with the EGI estimate via a PF scheme. In any variant, this method is unable to obtain non-convex shapes, returning instead convex-equivalent ones—*i.e.* the unique convex shape that would produce the same light curve.

An attitude-independent alternative that estimates average area-albedo of the cross section of a stable RSO is proposed in Payne et al. (2017). They introduce a set of light curve based methods, by which the average magnitude of a phase plot of the object of interest is compared against a catalogue of well-known objects. The difference is used to return the albedo-area both for simple-shaped objects (where the phase plot has a single peak), and for objects with distinct recognizable features (*e.g.* specular glints from solar panels).

Fulcoly et al. (2012) plant the seed for the classification-based approach to shape determination. They study a simulated set of light curves for a few selected shapes, and detect that the lower bound intensity versus phase angle curve is strongly correlated to the shape, regardless of size, attitude and surface reflective properties. This leads to the development of a shape-classification method based on comparing this lower bound of a test light curve, against the lower bound of each simulated set, on a shape by shape basis. It’s success rate is of $\sim 70\%$ for stable objects, and $\sim 80\%$ for rotating ones. A generalized approach based on Machine Learning (ML) is introduced by Furfaro et al. (2019), by which a CNN is trained to classify light curves among a set of known shapes. Even if the observation scenario is different than that of Fulcoly et al. (2012), the success rate is increased to more than 90% . Wu et al. (2016) implement another gradient-descent neural network to classify RSOs by shape, this time with the infrared signature taken by a space observer as input.

2.2.4. Surface Materials Characterization

The reflective properties of the surfaces of RSOs influence any electromagnetic-based measurements (*i.e.* light curves included). Thus, it should be possible to infer these properties from them. However, as indicated in the previous subsection, reflectivity, shape and size are strongly coupled, and their effects on passive measurements are not so easily decoupled. In the literature, some approaches have been proposed to determine the material composition of RSO surfaces. Stryjewski et al. (2010) postulate that polarized

2. Literature Review

light curves should contain independent information on material composition, since neither shape nor size should influence the polarization of the light upon reflecting off said surfaces. They suggest the construction of RSO classifiers based on a symmetry observed on different material diattenuation coefficients. This lead eventually to the work of Dianetti et al. (2019), where multiple-model adaptive estimation is used to identify RSO surface materials from polarized light curves. The dissertation of Dianetti (2020) expands on this to exploit spectroscopic light curves, *a.k.a.* differential photometry, to characterize surface materials.

Other works demonstrate the potential of using differential photometry to characterize RSO materials: Reyes (2018) catalogues the spectroscopic signature of frequently-used spacecraft materials, for future use in characterization techniques, while Zigo et al. (2019) classify objects between three different dominant materials, based on their differential photometry signature.

In general, material characterization literature suggests that polarized (Beamer et al., 2018; Cegarra Polo et al., 2019) and spectroscopic (Reyes et al., 2021) light curves can potentially be used for ML-based approaches to the material classification problem.

2.2.5. Simultaneous Characterization

This subsection deals with those methods and techniques in the literature that aim at estimating more than one *type* of physical characteristic from light curves at the same time. For example, Henderson et al. (2012) and Henderson (2014) use a batch-UKF and a bootstrapped-PF to simultaneously estimate position, attitude, albedo and size of a cuboid. On the same line, Linares et al. (2014b) fuse a light curve using a plain UKF to estimate position, attitude, mass and area of the RSO. Linares et al. (2012) present a similar UKF-based approach that estimates attitude and position together with an EGI parametrization of shape. Linares et al. (2014c) do the same, but use a multiple-model adaptive algorithm to capture the shape, instead. Linares et al. (2018) expand the state with facet normals and albedo-area products to perform a simultaneous attitude-shape estimation, albeit with known angular rates. They model the particles of a PF as a Hamiltonian system in probability space to improve filter efficiency by providing high-probability samples. Furthermore, after a certain number of iterations, Lasso regularization is used to preserve only those facet normals that contribute the most, leaving the rest out of the state vector. Finally Qin et al. (2018) estimate position and attitude using a classical PF, while Yun et al. (2020) additionally estimate surface reflectivity, too. Furthermore, the latter improve the computational cost of the PF by approximating the Probability Density Function (PDF) represented by the particle cloud of the filter with a Gaussian Mixture (GM), using the Expectation-Maximization (EM) algorithm. Then, a conventional EKF or UKF measurement update for each GM kernel substitutes the costlier PF update.

2.2.6. Abstract Characterization

This final subsection summarizes other characterization approaches in the literature, whose goal is to obtain more abstract characteristics. Most of the *more abstract* characterization

of objects comes in the form of classification, since these kind of characteristics are not so easily represented by finite subsets of N -dimensional real numbers. For example, Furfaro et al. (2016) propose an Ontology-based classification algorithm based on CNNs that learn to identify object classes from their light curves. They demonstrate this approach by training a network to identify the status and type of a RSO: active, inactive, space debris or fragment. Another variant of this approach is validated in Linares et al. (2016b), where objects are classified as rocket bodies, payloads or debris. Furfaro et al. (2018) compare this approach against other classification techniques, showing that CNNs surpass previous ML implementations.

Another subgroup of abstract OC is manoeuvre detection and characterization. This branch of OC rarely uses direct passive measurements, such as Radar data or light curves. Instead, it often relies on processed results in the form of orbital ephemerides. For example, Kelecy et al. (2007) propose a Delta-V manoeuvre detection method based on the identification of orbital-energy discontinuities extracted from Two-Line Element (TLE) data. They achieve impulsive manoeuvre characterization accuracy down to the cm s^{-1} . More recent approaches are based on ML. While Linares et al. (2017) combine an UKF for Orbit Determination (OD) plus reinforcement learning, Abay et al. (2018) do so with generative adversarial networks.

The manoeuvres that may be detected using direct measurements are those that modify the attitude of the RSO, since this has a direct effect on its light curve. In fact, the Attitude Determination (AD) methods developed in Holzinger et al. (2014) and Coder et al. (2017), mentioned above, are a good example of this application. These works estimate attitude under the assumption that the RSO is subject to control torques, modelled as a Markov chain of exponentially time-correlated acceleration states. Thus, angular accelerations are part of the output of their filters, providing an effective estimation of attitude manoeuvres.

2.2.7. Summary

To sum up, the current section gives a literature-review-based answer to the first question of this dissertation: *Which RSO characteristics can be deduced from light curves and how?*

For quick reference, the reader may refer to Appendix B, where the content of this section is summarized into a set of compact tables.

2.3. Observability of Object Characteristics

This next section reviews the literature that assesses how observable particular object characteristics are, in different scenarios. This is, it addresses the second research question: *Which observability conditions are necessary to do so?* In this context, *observability* is to be understood as the measure on how well a given characteristic (hidden state) can be inferred from a measurement sample (observable state).

First, Hinks et al. (2013) study the observability of the attitude and reflective parameters (albedo) from one single light curve measurement using its Fisher Information Matrix (FIM). They show how, for the attitude case, the FIM has rank 1. Thus, since the attitude state has 3 Degrees of Freedom (DoFs), for any attitude state there exists a plane along which

2. Literature Review

it can move without altering the measurement. They assume the Ashikhmin-Shirley BRDF (Ashikhmin et al., 2000). They conclude that the anisotropy direction of any surface of the object, along with its specular reflectance, is not observable unless the surface is observed exactly from the glint direction relative to the incoming Sun light. Linares et al. (2016a) build upon this concept to analyse other reflective models and prismatic shapes. They observe a marked decrease in attitude observability when the object shape tends from an elongated prism to a cylinder, for the case when the Sun and the observer lie on the object-centred plane perpendicular to the axial axis. Dianetti et al. (2020) use the FIM to study the observability of attitude from light curves, but this time using multiple spectral channels. Their results show that attitude is not more observable than with a single monochromatic light curve, except for corner-case configurations—the gain obtained from having the information spread over multiple channels is countered by the increased noise per channel.

Next, Früh (2015) proposes the use of the Gramian of the linearized estimation problem, considering position, attitude, shape and reflective properties as the hidden state, and optical measurements as the observations. This allows to consider the joint observability from a chain of measurements, in contrast to Hinks et al. (2013). Furthermore, she modifies the classical Gramian from control theory to incorporate the effect of measurement noise. Astrometric measurements (right ascension and declination on the sky) are analysed separately from photometric ones (light curves). Früh (2015) concludes that shape, attitude and reflective properties are not directly observable from astrometric measurements: even if the orbital dynamic model incorporates non-conservative forces and torques due to these characteristics, their effects are too coupled and only an abstract hidden state that summarizes them is truly observable. This same issue is expected when using light curves. Friedman et al. (2018) expand the Gramian approach to assess how long the observations should be, and how many astrometric measurements are necessary, to be able to extract reflective properties for different Low Earth Orbit (LEO) and Geosynchronous Earth Orbit (GEO) cases. The Gramian is used by Friedman et al. (2019) to assess the observability of the convex-equivalent shape, when using the EGI plus Minkowski method (see Table B.2) to extract it from the light curve. They demonstrate how observability analysis can be used to choose an efficient observation window and frame rate. They conclude that the light curve must be at least as long as the number of cells of EGI used (see Chapter 4). Subbarao et al. (2019) compare different reflective models in terms of observability for the OC case, too, but using the observability matrix, instead.

One particularly interesting approach, following the line of observability analysis, is that of Dianetti et al. (2018). They propose an observability-driven filter that estimates the object orbit and other characteristics using astrometric and photometric measurements. Whenever the Gramian/FIM of the *e.g.* attitude state becomes positive-definite (least singular value above some threshold), this state can be incorporated into the filter state, so that a more realistic model can be used. A similar approach within an Unscented-Schmidt Kalman Filter (USKF) to estimate position, mass, attitude and shape parameters is presented in Richardson et al. (2018), where different parameters are toggled between *considered* (the uncertainty of the parameter is taken into account, but it is not updated) and *estimated* (the parameter is updated), depending on their FIM score.

The following list sums up the answer to the second research question, *which observability*

conditions are necessary to do so?, as per the above literature review:

- Simultaneous shape and attitude characterization from light curves is strongly coupled in most cases, making it difficult to discern between the two, if neither is known a-priori. (Früh, 2015)
- Observability analysis can be used to design resource-efficient measurement scenarios. (Friedman et al., 2017)
- Multi-model approaches exist to decide which parameters to estimate depending on the observability assessment for each particular scenario. (Dianetti et al., 2018; Richardson et al., 2018)

2.4. Object Characterization Derived Use Cases

This section of the Literature Review tackles the third research question: *which SST use cases does direct Object Characterization enable?*

The most evident and straightforward use cases that SST offers are clearly defined in the Service Portfolio of the EU SST Support Framework (EU SST, 2021). They identify three main use cases:

- **Collision Avoidance (CA)** warns operators of oncoming high risk events;
- **Re-entry Analysis (RE)** monitors high risk atmosphere reentry events of man-made objects; and
- **Fragmentation Analysis (FG)** catalogues and analyses destructive events, when multiple new objects result from a collision between RSOs, or any other type of fragmentation event.

OC has been identified as a key asset to improve the quality of these services. For example, Friedman et al. (2022) state that shape and reflective properties determination is essential for precise orbit propagation. It follows that the more accurate the model of the RSO orbit becomes, the more reliable will collision warnings be. The same is true for atmospheric reentry events.

Another field where OC can contribute to better quality services is the domain of **Active Debris Removal (ADR)**. Liou (2013) identify ground-based Radar and optical OC as one of the main challenges faced by the ADR sector. They justify this statement on the fact that knowledge of the attitude states of space debris is necessary to design any ADR mission. On the same line, Juillard et al. (2021) state that the chaser agent in an ADR mission must be able to estimate the attitude of its target. In fact, object characterization capabilities (mainly for shape and attitude/angular momentum) are a common by-design feature of ADR agents. For example, in Lv et al. (2022) they are used by the chaser spacecraft during its approach to the target to autonomously decide which is the best-suited ADR method: harpoon or net. The need for ADR-oriented OC capabilities is made even more evident in the doctoral dissertation of Jankovic (2022), a work that proposes an ADR ontology designed to incorporate object characteristics of different origin, and digest them for optimal ADR decision-making.

2. Literature Review

Finally, another potential OC-enabled use case is **Health Monitoring** of operational satellites. Ryan et al. (2007) propose to use light curves to identify unexpected changes in the characteristics of operated satellites, which might be indicators of health issues (*e.g.* changes in surface properties might indicate degradation due to radiation exposure or debris impacts). Enomoto et al. (2016) propose an UKF shape and attitude estimation approach based on light curves, and using multiple-model estimation, to assess the health status of satellite components (*e.g.* antennae direction or position of solar panel). Likewise, from the ML point of view, Liu et al. (2021) conclude that neural networks could be trained to identify satellite health issues from light curves. Matsushita et al. (2019) further highlight the importance of OC for health monitoring purposes.

To summarize,

- characterization of attitude, shape and reflective properties positively impacts the accuracy of orbital predictions. This, in turn, improves the reliability of collision and reentry warnings, common use cases of the SST pipeline; (Friedman et al., 2022)
- ADR strongly benefits from a-priori knowledge of the characteristics of objects targeted for removal; (Liou, 2013)
- OC techniques have been identified as potential data sources for satellite health monitoring pipelines. (Enomoto et al., 2016)

2.5. The Scientific Contribution of This Thesis

Having explored the available literature on the three research questions (Section 1.2 of the Introduction) it is clear that they are answered in extensive detail by previous works. This means that the goals of this thesis need to be further particularized to fill the niches therein. Given the wide range of characteristics and methods available, in order not to dilute the contribution of this dissertation, the focus has been set on shape and attitude determination using panchromatic (one single, wide channel across the visual spectrum) light curves obtained from a CCD-equipped, ground-based optical telescope. This section of the literature review justifies this decision by showing that it can lead toward original scientific contribution in the field of Object Characterization.

The first niche has been identified in the light curve simulation field. In particular, regarding the rendering problem and self-shadowing on non-convex objects. The current approach that implements a solution for self-shadowing without relying on external software, in the state of the art, is ray tracing. The main goal of ray tracing, however, is to generate images as two-dimensional arrays of pixels. Rays are propagated from each pixel, through the observation scenario until they reach the light source or they dissipate (*e.g.* by bouncing a predetermined number of times on object surfaces) (Shirley et al., 2003). This approach has several drawbacks. Overall this can be significantly expensive to compute, especially if a fine resolution is desired, necessary for good light curve simulation accuracy (Früh et al., 2014). Additionally, being the light curve a single scalar measurement (not a 2D array), the effort to first calculate each pixel contribution through ray tracing, and then add all the contributions together, seems like a long way to achieve the goal. The other method presented above is that of Wang et al. (2018). This method is still an approximation, because each facet of the

object can only be shadowed or not as a whole. In their use case facets are small, so that this hypothesis does not introduce significant errors. However, this method loses accuracy when flat facets are of comparable size to the overall object characteristic dimension. Sometimes, it is interesting to simulate light curves of such simple geometries, where the object is defined by a few, big flat surfaces. There is no state-of-the-art method that can handle these cases with good accuracy and efficient computation times, for the reasons exposed above. This thesis introduces an alternative approach to the self-shadowing problem that addresses this niche.

The second niche that has been identified lies in the EGI plus Minkowski minimization method used to estimate convex-equivalent shapes from light curves (Table B.2). Although the method was introduced 20 years ago in the asteroid community (Kaasalainen et al., 2001), and is based on even older numerical methods (Little, 1985), it has arisen recent interest in the community. Several expansions have been proposed (Fan et al., 2020; Fan et al., 2021) that improve its performance with realistic light curves. Furthermore, its observability conditions have been explored (Friedman et al., 2019). However, there is some room for improvement.

- First, there is no established quantitative criteria to assess the performance of the method, but which can be applied directly on the obtained polyhedron shape. The state of the art is to assess how well the partial EGI result of the actual shape of the object is matched, but the actual final polyhedron is assessed qualitatively only (Calef et al., 2006; Fan et al., 2019). This thesis addresses this by proposing one such quantitative criterion.
- Second, Little (1985) suggests that the Minkowski minimization step of this method can be improved. He states that its convergence speed might benefit from exploiting analytical expressions for the gradient and Hessian of the volume of a polyhedron with its supports. These expressions have not been found in the literature, so this point is explored here.
- Third, the observability of the method has been explored using traditional, linearized (Gramian-based) techniques (Friedman et al., 2019). However, the nature of the problem does not allow negative areas, which means the solution space is constrained, hence the conclusions of Gramian-based observability analysis need to be reviewed. This is done in this thesis.

The third niche has been found for the sequential estimation approach to Attitude Determination from light curves. Two major methods exist. On the one hand, there is either the UKF- or the EKF-based approach, which is computationally efficient but is shown to be unfit to deal with the high non-linearity and weak-observability inherent to the problem (Wetterer et al., 2009; Yun et al., 2020). Although weak-observability can be mitigated with observability-driven implementations (Dianetti et al., 2018; Richardson et al., 2018), the non-linearity is still a problem. On the other hand, there is the family of PF-based approaches (Linares et al., 2014a), which capture non-linear, multi-modal PDFs well, but tend to be computationally expensive. Some improvements have been proposed to mitigate this issue, including the use of GMs (Yun et al., 2020) or a Hamiltonian Markov Chain (Linares et al., 2018). However, this dissertation proposes a third alternative, inspired on the Adaptive Gaussian Mixtures Unscented Kalman Filter

2. Literature Review

introduced in Schiemenz et al. (2020) for Orbit Determination. It is an alternative solution that can model complex, multi-modal PDFs, like the PF, but has the potential to be computationally cheaper, because it adapts the number of kernels (similar to the PF particles) based on an online assessment of the non-linearity present in the system at each time step, coupled with the novel Gaussian splitting libraries of Vittaldev et al. (2016).

Finally, regarding the third question on OC-enabled use cases, this thesis briefly explores the potential of these two methods toward ADR and health monitoring. Collision avoidance is left for further work, because it requires coupled modelling of attitude and shape with orbital dynamics. This falls out of the scope of this dissertation.

3. Light Curve Model

A light curve is the time series of brightness measurements taken by a telescope when observing a target object. Each measurement is a 1-dimensional value that is a function of the observed light intensity. This measurement is mainly conditioned by the following elements.

- The **observation geometry** refers to the relative position between the Sun, the object and the observer. This influences the measurement in two key aspects:
 - first, the observation geometry defines the distance travelled by the light until it is measured—the longer the distance, the higher the Free-Space Path Loss (FSPL), which is proportional to the inverse squared distance;
 - second, it determines which is the phase angle, the angle between the incident (Sun to object) and reflection (object to observer) directions—this is a key factor that determines how much light is reflected out of the surfaces of the object.
- Next, the **object characteristics** determine how the reflection happens locally (around the object). This is,
 - the **size** of the object influences how many photons interact with it;
 - the **shape** establishes the local geometry of the reflection;
 - the **surface material properties** define the spectral intensity (and polarization) of the light reflected from each surface element of the object; and
 - the **attitude** state determines the incident and reflection directions from the point of view of the object.
- Another factor that influences the measured intensity is the **sensor**. Its aperture size, focal length, quantum efficiency and overall transmissivity affect the final reading.
- Finally, other environmental factors can have an effect on measured intensity, such as **atmospheric attenuation** (in case of ground-based observers) or **Sun irradiance fluctuations**.¹

This chapter starts with Section 3.1, which describes the models used to solve the *Forward Problem*, a.k.a. Light Curve Simulation. It establishes the necessary relations between the variables of the problem (summarized above) and the measurement. These relationships are key to later study the OC problem, or *Inverse Problem*, which tries to determine a sub-set of this factors (*e.g.* shape or attitude) from the measurement. Next, Section 3.2

¹These two factors are not taken into account within the scope of this thesis. They are complex issues that would overcomplicate this work and distract from the goals at hand. Their effects are roughly modelled by artificial noise, when necessary.

3. Light Curve Model

identifies the need for an efficient rendering algorithm when simulating objects with simple, polyhedron geometries. Section 3.3 follows with what constitutes the scientific contribution of this chapter. This is the Shadow Projection algorithm, which proposes a new solution to said rendering problem by projecting shadows between triangular facets. This new algorithm is then tested in Section 3.4. Finally, the chapter concludes with a summary in Section 3.5.

This chapter is adapted from and expands on Vallverdú Cabrera et al. (2021a, Section 3), the first paper published as part of the doctoral research of this thesis. As such, parts of Vallverdú Cabrera et al. (2021a) are included herein.

3.1. Problem Definition

This section lays out the equations from the literature, and derived from the assumptions taken in this thesis when necessary, that are required to model the optical measurement of a RSO. Additionally, it establishes most of the definitions used throughout this thesis in coming chapters, as well as nomenclature and mathematical conventions.

This section is a prelude to the first research question, *Which Resident Space Object characteristics can be deduced from light curves [...]?*. This is, all the characteristics that can influence an optical measurement are discussed: only those that influence it are susceptible to be estimated in the inverse problem. Therefore, for the most part, this section can be seen by the reader as the equivalent of the *Fundamentals* of this thesis. All the theory necessary to build up to the research described in the subsequent sections and chapters is introduced here.

3.1.1. Observation Geometry

In this section the three agents whose interaction produces the measurement are established. These are the Sun, which acts as the illumination source; the object, which reflects the light coming from the Sun; and the observer, which receives the light coming from the object and registers its intensity.

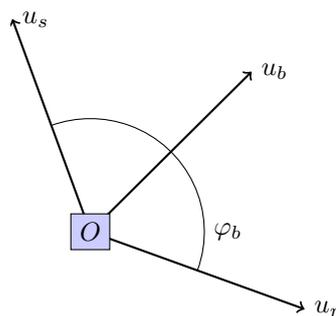


Figure 3.1.: Observation geometry projected on the plane formed by the Sun, the object and the observer (Vallverdú Cabrera et al., 2021a)

The complete observation geometry is shown in Figure 3.1. It represents the plane formed

by the Sun, the object (O) and the observer. The vectors are all unitary: u_s is the incident direction, which points away from the object and toward the Sun; u_r , the reflected direction, toward the observer. The vector u_b is the bisector of the other two,

$$u_b = \frac{u_s + u_r}{\|u_s\| + \|u_r\|} \quad (3.1)$$

a.k.a. Phase Angle Bisector (PAB). The angle between u_s and u_r is known as the phase angle, $\varphi_b = \angle(u_s, u_r)$.

These, together with the distance from object to Sun d_s , and from observer to object d_r , a.k.a. range, are the parameters that define the path travelled by the light throughout the whole process.

The main hypothesis that constrains the observation geometry for the case at hand is that the dimensions of the object can be neglected compared to the distances involved (d_s and d_r). This means that the light rays from the Sun are parallel in the vicinity of the object, and the same happens for the reflected ones at the observer location. In other words, this dissertation focuses on non-resolved objects—*i.e.* objects that are too far away from the observer to appear as anything more than a single point.

3.1.2. Object Shape and Size Representation

The next factors to be discussed are the shape and size of an object, and how they are represented throughout this work.

A tractable way to represent a three-dimensional object is to approximate it by an Euclidean region bound by a closed set of triangular surfaces, called facets (Wang et al., 2018). This is, an object $O \subset \mathbb{R}^3$ is delimited by N_f facets with three vertices each, connected with each other along their edges. Thus, its boundary δO is a closed surface made up by a set of facets:

$$\delta O = \left\{ F^{(k)} \mid k \in 1 \dots N_f \right\}. \quad (3.2)$$

Each facet $F^{(k)} \subset \delta O$ is uniquely defined with the set of its three vertices $v_i^{(k)}$,

$$\text{vset} \left(F^{(k)} \right) = \left\{ v_1^{(k)}, v_2^{(k)}, v_3^{(k)} \right\}, \quad (3.3)$$

where the operator `vset` returns the vertex set that uniquely defines a polygon, ordered as per the right-hand rule around the facet normal.

3.1.3. Attitude

On the one hand, the observation geometry (object orbit, observer and Sun locations) is given in the inertial frame I ; throughout this work, the inertial frame refers to the Geocentric Celestial Reference Frame (GCRF) as defined in Kaplan (2006). On the other hand, the geometry of the object (its shape and size) is defined in the body frame B —*i.e.* any frame attached to the object. The object attitude links the observation and object geometries: it is used to determine how the object itself is oriented relative to the Sun and the observer.

3. Light Curve Model

In particular, the u_s and u_r directions, which are naturally known in inertial frame, need to be transformed to body frame.

This section establishes the convention followed throughout the rest of this document to represent attitude. Quaternions, whose algebra is symbolized as \mathbb{H} (see Appendix C) are the main attitude representation used in this thesis. Different conventions exist in the literature to represent and manipulate quaternions (Solà, 2017). Thus, this section states which ones have been adopted in this work.

In particular, the quaternion product used throughout this thesis follows the Hamilton convention, as defined in Equation C.7. As stated in Appendix C, this work uses the symbol \odot to represent said Hamiltonian product of quaternions, following the notation used in Markley et al. (2014).

Attitude in this thesis is represented by the unit quaternion $\mathbf{q}_B^I \in \mathbb{A}$ that transforms vector in body frame $v^B \in \mathbb{R}^3$ to inertial frame $v^I \in \mathbb{R}^3$ representation. This is

$$\mathbf{v}^I = (\mathbf{q}_B^I)^* \odot \mathbf{v}^B \odot \mathbf{q}_B^I. \quad (3.4)$$

\mathbb{A} is the subset of \mathbb{H} of unit quaternions with positive real part (see Appendix D).

The convention used here is that bold symbols represent quaternions (*e.g.* $\mathbf{q} \in \mathbb{H}$), while the normal-font counterpart (*e.g.* $q \in \mathbb{R}^4$) is its vector representation. To illustrate:

$$\mathbf{q} = q_r + q_i \mathbf{i} + q_j \mathbf{j} + q_k \mathbf{k} \Rightarrow q = \begin{bmatrix} q_r \\ q_i \\ q_j \\ q_k \end{bmatrix}. \quad (3.5)$$

Following this same convention, whenever a symbol that has been previously identified as a three-dimensional vector (*e.g.* $v \in \mathbb{R}^3$) is shown in bold font (*e.g.* $\mathbf{v} \in \mathbb{H}$), it represents the quaternion whose real part is 0 and whose three hyper-imaginary components (\mathbf{i} , \mathbf{j} and \mathbf{k}) hold the three dimensions of the vector. This is

$$v = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \Rightarrow \mathbf{v} = v_1 \mathbf{i} + v_2 \mathbf{j} + v_3 \mathbf{k}. \quad (3.6)$$

Furthermore, the $*$ super-index indicates quaternion conjugate. Refer to Appendices C and D for a detailed description of these conventions, as well as a comprehensive summary of quaternion algebra \mathbb{H} .

If the results of this thesis are to be compared to other works that use different quaternion conventions, these two final notes need to be considered:

- First, in this work, attitude quaternions represent the rotation between coordinate frames, opposite to the typical approach on computer graphics where they represent vector rotations inside a single frame (Solà, 2017). This choice is reflected by the fact that in Equation 3.4 the conjugate quaternion multiplies on the left, while the non-conjugate one is on the right. This has the consequence that quaternion composition goes from left to right (opposite to *e.g.* direction cosine or rotation matrices). This is,

the quaternion that transforms from frame B to frame I can be composed by the one that goes from B to intermediate frame C , and the one from C to I as follows:

$$\mathbf{q}_B^I = \mathbf{q}_B^C \odot \mathbf{q}_C^I. \quad (3.7)$$

- Second, the vector representation of a quaternion (q in non-bold font) has the real part in the first position, while the hyper-imaginary parts \mathbf{i} , \mathbf{j} and \mathbf{k} occupy the second, third and fourth. This is opposite to some literature where the real part is placed at the end of the vector representation.

With the quaternion representation of attitude, kinematics are defined by the simple equation

$$\dot{\mathbf{q}}_B^I = -\frac{1}{2} \boldsymbol{\omega}_{B/I}^B \odot \mathbf{q}_B^I \quad (3.8)$$

where the dot indicates derivative over time, and $\boldsymbol{\omega}_{B/I}^B$ is the instantaneous angular velocity of frame B relative to frame I , expressed in B coordinates. See Appendix D for the derivation of this equation.

Attitude dynamics in Newtonian physics can be described by the conservation of angular momentum of a constant-mass, rigid system as (Wertz et al., 2011)

$$I_G^B \dot{\boldsymbol{\omega}}_{B/I}^B = T^B - \boldsymbol{\omega}_{B/I}^B \times \left(I_G^B \boldsymbol{\omega}_{B/I}^B \right), \quad (3.9)$$

where $T \in \mathbb{R}^3$ is the net torque applied to the object, while $I_G \in \mathbb{R}^{3 \times 3}$ is the tensor of inertia of the object. Without loss of generality, the rest of this thesis assumes that the body frame is aligned with the principal axes of the object—*i.e.* I_G^B is diagonal. The symbol \times indicates cross product.

3.1.4. Bidirectional Reflectance Distribution Function

This section deals with the object's surface properties, and how these interact with the light. This effect is modelled with the Bidirectional Reflectance Distribution Function (BRDF), defined by Nicodemus (1965) as

$$f_r(\lambda, u_s, u_r) = \frac{dL_{r,\lambda}(\lambda, u_r)}{dE_{s,\lambda}(\lambda, u_s)}. \quad (3.10)$$

The BRDF represented by f_r gives the ratio, in inverse steradians or sr^{-1} , between the reflected spectral radiance $dL_{r,\lambda}$ that leaves a surface element dS in the direction u_r , in $\text{W}/\text{sr}/\text{m}^2$, over the spectral irradiance $dE_{s,\lambda}$ incident on the surface element in the direction opposite to u_s , in W/m^2 . It is dependent on the wavelength of light, λ , hence the *spectral* qualification.² The differential form is commonly used to allow for contributions to the total $L_{r,\lambda}$ other than $dE_{s,\lambda}$.

Figure 3.2 shows the same observation geometry from Figure 3.1, but from the point of view of a surface element of the object with anisotropic directions u_u and u_v and normal u_n . The

²The λ subscript indicates derivative over wavelength. From here onward, the explicit dependency on λ will be dropped from spectral quantities, as it will be implicit in the subindex—e.g. $L_{s,\lambda}(u_s) \equiv L_{s,\lambda}(\lambda, u_s)$.

3. Light Curve Model

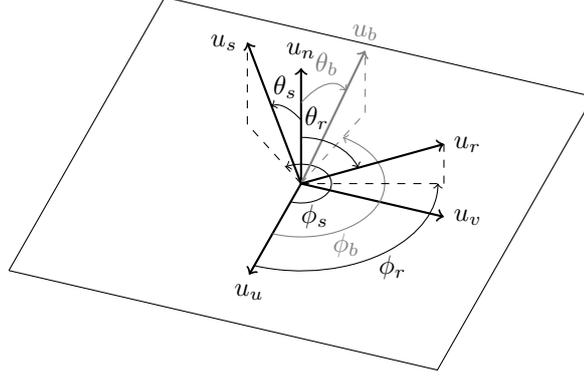


Figure 3.2.: The observation geometry for a surface element (Vallverdú Cabrera et al., 2021a)

angle pairs (θ_s, ϕ_s) , (θ_r, ϕ_r) and (θ_b, ϕ_b) are the spherical coordinates (polar, azimuth) of the incident, reflected and PAB directions respectively, as seen by the *facet coordinate system*, whose base is formed by $[u_u \ u_v \ u_n]$.

Several BRDFs models have been used to simulate light curves in the literature (see Section 2.1). Henderson (2014) concludes that the Ashikhmin-Shirley (AS) (Ashikhmin et al., 2000) is the best suited for anisotropic materials. Subbarao et al. (2019) identify the Cook-Torrance (CT) (Cook et al., 1982) as one of the BRDFs whose observability conditions are better. These two are often used in other OC-related literature, such as Wetterer et al. (2009) and Linares et al. (2014a). Furthermore, they conserve energy, which is a necessary condition for physically realistic modelling (Wetterer et al., 2014). Thus, they are the ones that have been used as main models within this thesis.

Below follow the equations for these two methods (Wetterer et al., 2014). For both methods, the BRDF is

$$f_r(\lambda, u_s, u_r) = \mathfrak{d}R_d + (1 - \mathfrak{d})R_s, \quad (3.11)$$

where R_d and R_s are the bidirectional diffuse and specular reflectances, respectively, and \mathfrak{d} is the fraction of the total BRDF that R_d represents. In the most general case, these three parameters may depend on wavelength λ , while only R_d and R_s depend on u_s and u_r . The dependence of R_d and R_s on λ is implemented through the diffuse $\rho(\lambda)$ and specular $F_0(\lambda)$ reflectances of the material.

Ashikhmin-Shirley. The bidirectional radiances of the AS model are

$$R_d = \frac{28\rho}{23\pi}(1 - sF_0) \left(1 - \left(1 - \frac{u_n^T u_s}{2}\right)^5\right) \left(1 - \left(1 - \frac{u_n^T u_r}{2}\right)^5\right) \quad (3.12)$$

and

$$R_s = \frac{\sqrt{(n_u + 1)(n_v + 1)}}{8\pi} \frac{F_{AS}}{c \max(u_n^T u_s, u_n^T u_r)} (\cos \theta_b)^{n_u \cos^2 \phi_b + n_v \sin^2 \phi_b}, \quad (3.13)$$

where

$$F_{AS} = F_0 + \left(\frac{1}{s} - F_0\right) (1 - c)^5 \quad (3.14)$$

is the Fresnel approximation used, and

$$c = u_r^T u_b. \quad (3.15)$$

The parameters $n_v \geq 1$ and $n_u \geq 1$ are called exponential factors and govern the anisotropy of the surface element, as well as the shape of the reflection hemisphere of the BRDF. This is, for higher values of n_v , the bidirectional specular reflectance along direction u_v allows less energy dispersion away from the plane of pure specular reflection; n_u has the same effect along direction u_u .

Cook-Torrance with Lambert diffuse component. For the CT, the bidirectional reflectances are³

$$R_d = \frac{\rho}{\pi} \quad (3.16)$$

and

$$R_s = \frac{DGF_{CT}}{4(u_n^T u_r)(u_n^T u_s)}. \quad (3.17)$$

The terms in the specular component are the slope distribution function,

$$D = \frac{\exp\left(-\frac{\tan^2 \theta_b}{\mathbf{m}^2}\right)}{\pi \mathbf{m}^2 \cos^4 \theta_b}; \quad (3.18)$$

the geometrical attenuation factor,

$$G = \min\left(1, \frac{2(u_n^T u_b)(u_n^T u_r)}{c}, \frac{2(u_n^T u_b)(u_n^T u_s)}{c}\right); \quad (3.19)$$

and the Fresnel equation,

$$F_{CT} = \frac{1}{2} \frac{(g-c)^2}{(g+c)^2} \left(1 + \frac{(c(g+c)-1)^2}{(c(g-c)+1)^2}\right), \quad (3.20)$$

where c is the one defined in Equation 3.15,

$$g^2 = n_i^2 + c^2 - 1, \quad (3.21)$$

while

$$n_i = \frac{1 + \sqrt{F_0}}{1 - \sqrt{F_0}} \quad (3.22)$$

is the index of refraction. Here there is no anisotropy, and the shape of the BRDF is instead controlled by the facet slope parameter $\mathbf{m} \in [0, 1]$. There is a rough equivalence between the two methods at $n_v = n_u = 2/\mathbf{m}^2$. Thus, for lower values of \mathbf{m} , the bidirectional

³Some discrepancy exists between the CT version presented here, the one from Wetterer et al. (2014) (also used *e.g.* in Holzinger et al. (2014)), with the original formulation in Cook et al. (1982). Namely, R_s is not divided by 4 there. Furthermore, other sources differ slightly in other terms. For example, in Coder et al. (2018) or in Du et al. (2018), R_s is divided by π , instead of by 4. These are all small differences, which only change how the specular component varies w.r.t. the diffuse one. The shape of the specular component itself, as well as the properties of the function (non-linearity), are unaffected. For consistency, in this thesis the version of Wetterer et al. (2014) is used.

3. Light Curve Model

specular reflectance allows less energy dispersion away from the plane of pure specular reflection.

The irradiance is computed as the total radiance that is incident on the surface element orthogonally (Nicodemus et al., 1977). Thus, it can be obtained as

$$dE_{s,\lambda}(u_s) = L_{s,\lambda}(u_s) \cos \theta_s d\Omega_s, \quad (3.23)$$

where $L_{s,\lambda}$ is the incident radiance, $d\Omega_s$ is the solid angle element in the illumination direction and $\theta_s = \angle(u_n, u_s)$ is defined as in Figure 3.2.

3.1.5. The Path of Light

This section computes the irradiance received at the location of a distant observer due to the Sun light reflected by the object, as per the hypotheses that constrain the model used in this thesis.

First, the illumination source (the Sun) is assumed to emit black-body radiation, then $L_{s,\lambda}$ follows Planck's law (Planck et al., 1914):

$$L_{s,\lambda} = B_\lambda(T_s) = \frac{2hc^2}{\lambda^5} \left(\exp\left(\frac{hc}{\lambda k_B T_s}\right) - 1 \right)^{-1}, \quad (3.24)$$

where T_s is the temperature of the illumination source; h , Planck's constant; c , the speed of light in vacuum; and k_B , Boltzmann's constant.

Then, assuming a single illumination source, the total spectral radiance reflected by a surface element can be computed by integrating along the area of the illumination source S_s :

$$L_{r,\lambda}(u_s, u_r) = \int_{S_s} dL_{r,\lambda}(\lambda, u_r) = \int_{S_s} f_r(\lambda, u_s, u_r) L_{s,\lambda}(u_s) \cos \theta_s d\Omega_s. \quad (3.25)$$

If the surface element is in the vicinity of the Earth, and the single illumination source is that of the Sun, then the total incident solid angle becomes on average $\Omega_s \sim 2.72 \times 10^{-4}$ sr. This value is sufficiently small for θ_s and u_s to be considered constant along $d\Omega_s$. Thus, the above integral for the scenario studied in this thesis simplifies to

$$L_{r,\lambda}(u_s, u_r) \simeq f_r(\lambda, u_s, u_r) \cos \theta_s E_{s,\lambda}^\perp(u_s), \quad (3.26)$$

where $E_{s,\lambda}^\perp(u_s) = \Omega_s L_{s,\lambda}(u_s)$ is the irradiance a surface element orthogonal to u_s would receive.

Next, the spectral irradiance produced by the surface element dS , as received from a distant observer located along the direction of u_r , is

$$dE_{r,\lambda}^\perp(u_s, u_r) = L_{r,\lambda}(u_s, u_r) d\Omega_r, \quad (3.27)$$

where $d\Omega_r$ is the solid angle element as seen at the observer location. This is, if the observer is located at a distance d_r from the surface element in question,

$$d\Omega_r = \frac{\cos(\theta_r) dS}{d_r^2}. \quad (3.28)$$

Then, the spectral irradiance received at the observer is obtained by integrating Equation 3.27 along the visible part of each facet. This is

$$E_{a,\lambda}(u_s, u_r) = \sum_{k=1}^{N_f} \int_{\mathcal{A}^{(k)}} \frac{\cos(\theta_r^{(k)})}{d_r^2} L_{r,\lambda}^{(k)}(u_s, u_r) dS, \quad (3.29)$$

where $\mathcal{A}^{(k)}$ is the area of the portion of facet k that is both illuminated by the Sun, and visible from the observer direction. Again, assuming the object is orders of magnitude smaller than the distance between object and observer, in this dissertation this integral has been simplified to

$$E_{a,\lambda}(u_s, u_r) = \sum_{k=1}^{N_f} \frac{\mathcal{A}^{(k)}(u_s, u_r) \cos(\theta_r^{(k)})}{d_r^2} L_{r,\lambda}^{(k)}(u_s, u_r), \quad (3.30)$$

in W/m^3 .

The total radiance reflected over a certain bandwidth B_λ is obtained by integrating the spectral one over wavelength, which yields

$$L_r = \int_{B_\lambda} L_{r,\lambda} d\lambda, \quad (3.31)$$

in $\text{W}/\text{sr}/\text{m}^2$. Likewise, the total irradiance received at observer location will be

$$E_a = \int_{B_\lambda} E_{a,\lambda} d\lambda = \sum_{k=1}^{N_f} \frac{\mathcal{A}^{(k)}(u_s, u_r) \cos(\theta_r^{(k)})}{d_r^2} L_r^{(k)}(u_s, u_r), \quad (3.32)$$

in W/m^2

Finally, taking into account the above assumptions, apparent magnitude—the logarithmic measure of an object brightness, as seen from an atmosphere-less Earth (Tassoul, 2004)—can be computed as

$$m_A = -2.5 \log_{10} \left(\frac{E_a}{E_{a,0}} \right), \quad (3.33)$$

where $E_{a,0}$ is the irradiance received from the Vega star at the observer location (assuming no obstructions of any kind).

Except for the computation of $\mathcal{A}^{(k)}$, which is covered in its dedicated section (3.3), this completes the set of assumptions needed by this thesis to compute the power per unit surface of the Sun light reflected out of the object and received at the observer.

3.1.6. The Sensor

This section deals with the model of the sensor—the optical telescope—, whose task is to capture this light and register its intensity, so that it can be used by the (computer) OC algorithms studied and developed in the following chapters.

3. Light Curve Model

This work assumes that the aperture is small enough for $E_{a,\lambda}$ to be constant along its surface. Being the aperture area A_a and normal u_a , this assumption yields the spectral flux that reaches its focal plane:

$$\Phi_{a,\lambda} = A_a \cos(\theta_a) E_{a,\lambda}, \quad (3.34)$$

in W m^{-1} , where $\theta_a = \angle(u_a, -u_r)$ is the Line of Sight (LoS) angle.

The probability of an incident photon of wavelength λ being detected as an electron is modelled as the transmissivity $\tau(\lambda)$ for the entire telescope pipeline—it includes quantum efficiency and other losses (Tobergte et al., 2013). Thus, the spectral electron flux is

$$F_{a,\lambda} = \Phi_{a,\lambda} \tau(\lambda) \frac{\lambda}{hc} \quad (3.35)$$

in $\text{e}^- \text{s}^{-1} \text{m}^{-1}$. Next, this thesis assumes the measurement can be obtained by integrating over some finite bandwidth B_λ . Therefore, the total electron flux in $\text{e}^- \text{s}^{-1}$ used here is

$$F_a = \int_{B_\lambda} F_{a,\lambda} d\lambda. \quad (3.36)$$

It is a function of the received total irradiance E_a (Equation 3.32), although it is not necessary proportional to it.

Telescope measurements are often expressed in logarithmic scale as instrumental magnitude, which similar to the apparent magnitude m_A (Equation 3.33), is defined as (Tobergte et al., 2013)

$$m_I = -2.5 \log_{10} F_a. \quad (3.37)$$

The last assumption made in this section is that the telescope does not compute the electron flux, but instead it counts the number of electrons accumulated over the integration or exposure time t_e . Thus, the flux is approximated as

$$F_a \simeq \frac{C_a}{t_e}, \quad (3.38)$$

being C_a the number of electrons e^- counted during the window t_e .

3.2. Motivation for a New Solution to the Shadows Problem

This section deals with the problem of finding $\mathcal{A}^{(k)}$, the illuminated and visible portion of the area of each facet k . It has been omitted on purpose in the previous section, because it constitutes the main scientific contribution of this chapter.

When computing $\mathcal{A}^{(k)}$, it is necessary to distinguish between convex and non-convex objects. In the particular case for convex objects, the computation of $\mathcal{A}^{(k)}$ is straightforward (Piergentili et al., 2017):

$$\mathcal{A}^{(k)} = \delta(\theta_s^{(k)}, \theta_r^{(k)}) A_k \quad (3.39)$$

where A_k is the area of the facet and

$$\delta(\theta_s^{(k)}, \theta_r^{(k)}) = \begin{cases} 1 & \cos \theta_s^{(k)} > 0 \text{ and } \cos \theta_r^{(k)} > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (3.40)$$

In words, only the facets that look toward the Sun and the observer at the same time are fully accounted for, and the rest are ignored.

For non-convex objects, however, this computation complicates greatly, because facets can shadow each other. The current state of the art uses approximate solutions to solve the shadows problem for light curve simulation. Ray tracing is expensive to compute, especially without dedicated hardware—GPUs. Thus, it yields low accuracy because either resolution must be sacrificed in favour of computational tractability, or an approximate solution has to be used, such as the hidden vertex removal (Früh et al., 2014). Alternatively, Wang et al. (2018) check facets overlap in both illumination and reflected directions, but assuming the facets are small enough to be accounted for as points. This is a good approximation if the object is defined as a set of many small (triangular) facets. However, it is not efficient for simple geometries, where a few, big facets suffice to define the object: these facets need to be divided into smaller ones until the main assumption of the method can be accepted.

Simulating accurate light curves of simplified non-convex objects is beneficial for fast validation of OC algorithms, especially because resolving the shadows problem is typically the bottleneck of the whole light curve simulation pipeline—the rest of the model is made of straightforward analytical equations (see the previous section).

Therefore, this dissertation introduces the new *shadow projection* algorithm (see Algorithm 1 in the following section) with the aim to approximate a solution to this problem, that fills the niche where the state-of-the-art methods mentioned above lag behind: simple objects with a few flat facets. Simulating these kind of simplified objects accurately and inexpensively contributes to faster prototyping of OC algorithms.

3.3. The Shadow Projection Algorithm

The Shadow Projection algorithm was qualitatively introduced in Vallverdú Cabrera et al. (2021a) for the first time, the first work published within the scope of this thesis. This section gives a fully detailed and formal description of the algorithm, instead.

This algorithm is constrained by the following two hypotheses:

1. light rays only bounce once—*i.e.* they *die out* upon their second reflection event;
2. the facets that define O can only intersect with each other at their borders—*i.e.* edges or vertices.

The output of the shadow projection algorithm is directly the product $\mathcal{A}_r^{(k)} = \mathcal{A}^{(k)} \cos(\theta_r^{(k)})$. The value of $\mathcal{A}^{(k)}$ itself is not necessary, since it never appears alone in the reflection formulas above.

Before addressing the algorithm itself, some auxiliary operations, such as frame transformations and other geometrical manipulations, need to be introduced.

3.3.1. Projection onto a Plane

The first operation needed is the one that projects any point $x \in \mathbb{R}^3$ onto a plane with its own coordinate system. In particular, given $u \in \mathbb{R}^3$, $\|u\| = 1$, the projection plane is

$$U = \{x \in \mathbb{R}^3 \mid u^T x = 0\} . \quad (3.41)$$

This is, U contains the origin. To define the coordinate system associated to U , one needs two independent, and preferably orthogonal, vectors $e_1, e_2 \in U$ that act as a base. The goal is to express the projection of x on U as a linear combination of e_1 and e_2 . One can achieve this by proposing a frame transformation from the original frame onto some frame whose base is $[e_1 \ e_2 \ u]$, and then losing the u component.

To do so, one can use the quaternion that transforms between two different frames, when the representations of one vector are known in both (See Appendix D.3). Given $u = [u_1 \ u_2 \ u_3]^T$, one such frame transformation quaternion can be obtained using Equation D.16, with $v = u$, and $w = [0 \ 0 \ 1]^T$. In other words, one finds a quaternion that transforms a vector from whatever original frame where the normal of U is u , to a destiny frame where the normal of U is parallel to the third base axis. This results in

$$q_u = \frac{1}{\sqrt{u_1^2 + u_2^2 + (1 + u_3)^2}} \begin{bmatrix} 1 + u_3 \\ -u_2 \\ u_1 \\ 0 \end{bmatrix} . \quad (3.42)$$

This is,

$$\hat{x} = (q_u)^* x q_u . \quad (3.43)$$

After the transformation, the third coordinate of x can be eliminated, yielding

$$y = I_p \hat{x} , \quad (3.44)$$

where

$$I_p = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} . \quad (3.45)$$

The projection map that returns y as a function x is defined as $\pi_u : \mathbb{R}^3 \rightarrow \mathbb{R}^2$. This is, $y = \pi_u(x)$. Note that this is a linear map, since

$$y = \pi_u(x) = (I_p R_u) x , \quad (3.46)$$

where $R_u = R(q_u)$ is the Direction Cosine Matrix (DCM) form of the transformation quaternion q_u (See Appendix C).

Furthermore, it is important to remark that $x = \pi_u^{-1}(y)$ does not exist, since the solution is not a unique point in \mathbb{R}^3 , but a whole line orthogonal to U . In the following section, an alternative to obtain a pseudo inverse of π_u is proposed.

3.3.2. Lifting from the Plane

The second operation that is needed by the shadow projection algorithm is lifting a projected point (as in the previous subsection) back to some other coordinate system, including 3- or higher-dimensional ones. As said above, π_u^{-1} is not defined. However, let $V \subset \mathbb{R}^N$, for $N \geq 2$, be another plane that is not orthogonal to U . If one demands that the lifted point is in V , then one can define a unique lifting operation if one knows its inverse. This inverse is the linear map that transforms from whatever coordinate frame V is expressed in (let its points be defined as $z \in V$) to U coordinates (y in the previous section). Let this map be $\pi_{vu} : V \rightarrow \mathbb{R}^2$. For the case of $V \subset \mathbb{R}^3$, this is

$$\pi_{vu}(z) = \pi_u(x = z). \quad (3.47)$$

The goal is to find $\pi_{uv} = \pi_{vu}^{-1}$. Since both the domain and the image of π_{vu} are an infinite plane, π_{vu}^{-1} does exist if π_{vu} is linear and bijective—*i.e.* U and V are not orthogonal. The method proposed here to solve this problem exploits the fact that any plane can be fully defined by three of its points that form a non-degenerate triangle (non-zero area). This approach matches naturally with the fact that the object facets are all triangles.

Let $z_a, z_b, z_c \in V$ be three such points. Since the triangle T_z formed by these three points (as $\text{vset}(T_z) = \{z_a, z_b, z_c\}$) is non-degenerate, any two of its edges can be used as a (not necessarily orthogonal) base of the plane V . This is, any point $z \in V$ can be expressed as a linear combination of one of the triangle's vertices and the two edges that exit from it. This is

$$z = z_a + t(z_b - z_a) + s(z_c - z_a), \quad (3.48)$$

where coefficients $t, s \in \mathbb{R}$ are the coordinates when using the edges as base vectors.

The fact that π_{vu} is a linear map means that there exists $M \in \mathbb{R}^{2 \times N}$ such that

$$y = \pi_{vu}(z) = Mz. \quad (3.49)$$

For the $V \subset \mathbb{R}^3$ particular case, M can be obtained from Equation 3.47 as $M = I_p R_u$. Thus, Equation 3.48 can be transformed into

$$y = M(z_a + t(z_b - z_a) + s(z_c - z_a)) = y_a + t(y_b - y_a) + s(y_c - y_a). \quad (3.50)$$

Effectively, the coordinates (t, s) that represent some point in plane V as a linear combination of two edges of the triangle T_z defined in V , also represent the projection of that point onto U as a function of the corresponding edges of the triangle T_y (as $\text{vset}(T_y) = \{y_a, y_b, y_c\}$).

This means that given $z_a, z_b, z_c \in V$, π_{uv} and some point y , one can find t and s by solving Equation 3.50. This yields

$$t = \frac{(y - y_a) \times (y_c - y_a)}{(y_b - y_a) \times (y_c - y_a)}, \quad (3.51a)$$

$$s = -\frac{(y - y_a) \times (y_b - y_a)}{(y_b - y_a) \times (y_c - y_a)}. \quad (3.51b)$$

Finally, one can obtain z from Equation 3.48. From here onward, the lifting operation of point y from the plane represented by triangle T_y to the one represented by triangle T_z shall be referred to as

$$z = \text{lift}(y; T_y, T_z). \quad (3.52)$$

3.3.3. Polygon Clipping

Polygon clipping (a.k.a. polygon boolean operations) is the last key point of the Shadow Projection Algorithm. In particular, polygon intersection and subtraction are needed.

Triangle Intersection Detection

The simplest operation needed is to detect whether two triangles overlap—*i.e.* whether their intersection is not null. If the triangles are defined as a set of ordered vertices in \mathbb{R}^2 , this is a simple task that can be accomplished using the algorithm proposed in *e.g.* Mol (2022).

Along general lines, this triangle intersection algorithm checks whether any of the edges of one triangle intersect with the edges of the other. If any do intersect, it means that the triangles overlap. Refer to Mol (2022) for the exact logic.

This algorithm is interesting to use when one wants to know whether two triangles intersect, but may not be interested in calculating the actual intersection, which is a slower process.

Polygon Intersection and Subtraction Calculation

The other necessary processes are, given two polygons $P^{(1)}, P^{(2)} \subset \mathbb{R}^2$, to actually compute their intersection,

$$P^{(1 \wedge 2)} := P^{(1)} \wedge P^{(2)}, \quad (3.53)$$

and any of the possible subtractions,

$$P^{(1 \setminus 2)} := P^{(1)} \setminus P^{(2)} \quad (3.54)$$

and

$$P^{(2 \setminus 1)} := P^{(2)} \setminus P^{(1)}. \quad (3.55)$$

This can be achieved efficiently with *e.g.* the algorithms proposed by Vatti (1992).

The general principle of the Vatti (1992) algorithm is to scan the two operand polygons over the areas formed between horizontal lines that pass through consecutive vertices (in the vertical sense) of any of the operand polygons. Edge intersections are registered, and used to form the vertices of the output polygon. Different vertex generation rules are used depending on the operation (union, intersection or subtraction) being performed. A fully-detailed explanation can be found in Vatti (1992).

This algorithm can efficiently handle any type of 2-dimensional polygons, including holed, composed and/or self-intersecting ones. This is its most interesting feature for the case at hand. As explained below, it is possible that, over the course of the shadow projection algorithm, several polygons are subtracted from a single triangle. With every subtraction, the original triangle can acquire dents, holes and even become split in smaller pieces. Thus, a clipping algorithm that can handle these type of polygons is necessary.

3.3.4. Algorithm

At the start of the algorithm, each facet $F^{(k)}$ is projected onto the planes U_r and U_s , which are defined by the equations $r^T u_r = 0$ and $r^T u_s = 0$, respectively. Let them be called observer and Sun planes, respectively. To do so, it suffices to transform the vertex sets that define each facet as follows:

$$\text{vset} \left(F_r^{(k)} \right) = \left\{ \pi_{u_r}(v) \mid v \in \text{vset} \left(F^{(k)} \right) \right\} \quad (3.56)$$

and

$$\text{vset} \left(F_s^{(k)} \right) = \left\{ \pi_{u_s}(v) \mid v \in \text{vset} \left(F^{(k)} \right) \right\}. \quad (3.57)$$

This step is similar to the projections used in Wang et al. (2018).

In the next step, the set $\{F_r^{(k)}\}$ is copied into what will be the set of output visible portions of each facet:

$$V^{(k)} = F_r^{(k)}. \quad (3.58)$$

Then, the projections on the observer plane are compared against each other. For each pair j and k , the algorithm finds the polygon intersection $F_r^{(j \wedge k)} = F_r^{(j)} \wedge F_r^{(k)}$. If the intersection is not null, it needs to find which of the two facets is in front of the other, from the point of view of the observer. To do so, first it selects one point of $F_r^{(j \wedge k)}$. The algorithm uses the arithmetic mean of its vertices ($\bar{v}_r^{(j \wedge k)}$). Since the intersection of two triangles will always be a convex polygon, $\bar{v}_r^{(j \wedge k)}$ is guaranteed to be inside the polygon defined by $F_r^{(j \wedge k)}$. By extension, $\bar{v}_r^{(j \wedge k)}$ is inside the triangles $F_r^{(j)}$ and $F_r^{(k)}$, too. Then, the algorithm recovers the points $\bar{v}^{(j)}$ and $\bar{v}^{(k)}$ that are inside the triangles $F^{(j)}$ and $F^{(k)}$, such that $\pi_{u_r}(\bar{v}^{(j)}) = \pi_{u_s}(\bar{v}^{(k)}) = \bar{v}_r^{(j \wedge k)}$. It does so by using the lifting operation described in the previous subsection, together with the triangles $F^{(j)}$ and $F^{(k)}$.

At this point, one can compute the relative distance of any point from the observer as

$$\text{dist}(x) = -x^T u_r. \quad (3.59)$$

This does not mean that $\text{dist}(x) = 0$ implies that the observer is exactly at x ; but rather that if $\text{dist}(x) > \text{dist}(x')$, then x is farther from the observer than x' . One can apply this to $\bar{v}^{(j)}$ and $\bar{v}^{(k)}$ to assess which facet covers which: if $\text{dist}(\bar{v}^{(j)}) > \text{dist}(\bar{v}^{(k)})$, then facet j is behind facet k , from the point of view of the observer, and vice-versa.

Let f be the index of the facet that is farther, while c is that of the one that is closer. Then, it is necessary to subtract (using polygon boolean operations) the closest facet $F^{(c)}$ from the visible portion of the farthest one $V^{(f)}$:

$$V^{(f)} \leftarrow V^{(f)} \setminus F^{(c)}. \quad (3.60)$$

This cycle is repeated for all combinations of $j \in 1 \dots N_f - 1$ and $k \in j + 1 \dots N_f$. However, as facets become covered, some visible portions $V^{(k)}$ may become null. Whenever two facets j and k to be compared have $V^{(j)} = V^{(k)} = \emptyset$, the pair is skipped. However, note that if only one is obscured, the comparison still needs to be done, since it could still happen that the obscured one is covering the other.

3. Light Curve Model

At the end of this cycle, the set $\{V^{(k)}\}$ contains all the polygons that represent the portions of each facet that are visible from the observer, already projected on the observer plane. Note that these polygons may not be triangles anymore, or even convex. In fact, some could even be composed holed polygons. This is not an issue, since they can still be represented by a set of vertices: two non-intersecting triangles can be represented as a single polygon by connecting one vertex of each triangle using two overlapping edges, while a holed polygon can be converted into a non-convex one by joining an inner and an outer vertex in the same way (Vatti, 1992).

At this point, the area of each polygon $V^{(k)}$ is the area of facet k orthogonal to the observer direction that is visible from the observer. However, it is still necessary to calculate how much of these areas are actually illuminated by the Sun, too. To do so, the facets are compared again to each other, but this time using their projections on the Sun plane, $F_s^{(j)}$, $F_s^{(k)}$. Again, for every pair where $F_s^{(j \wedge k)} = F_s^{(j)} \wedge F_s^{(k)} \neq \emptyset$ and $V^{(j)}, V^{(k)} \neq \emptyset$, the algorithm performs the same lifting operation on a common point to assess which facet covers which, from the Sun point of view. However, in this second cycle there is the additional difficulty that the facet $F_s^{(c)}$ cannot be directly subtracted from $V^{(f)}$, because they exist in different projections. Instead, the whole intersection polygon $F_r^{(j \wedge k)}$ is lifted (vertex by vertex) directly onto U_r . Since it is a shadow cast on facet f , the triangle pair used for the projection is $F_s^{(f)}$ and $F_r^{(f)}$. This is, the lifted polygon $S^{(c)}$ is defined by

$$\text{vset} \left(S^{(c)} \right) = \left\{ \text{lift} \left(v; F_s^{(f)}, F_r^{(f)} \right) \mid v \in \text{vset} \left(F_s^{(j \wedge k)} \right) \right\}. \quad (3.61)$$

It can now be subtracted from $V^{(f)}$:

$$V^{(f)} \leftarrow V^{(f)} \setminus S^{(c)}. \quad (3.62)$$

Finally, given $\text{vset}(V^{(k)}) = \{v_1, \dots, v_{N_k}\}$ one can compute the total projected area of each facet that is both visible by the observer and illuminated by the Sun

$$\mathcal{A}_r^{(k)} = \cos(\theta_r^{(k)}) \mathcal{A}^{(k)} = \left\| \sum_{j=1}^{N_k} v_j \times v_{(j \bmod N_k) + 1} \right\|, \quad (3.63)$$

where N_k here is the number of vertices of $V^{(k)}$, while $x \bmod y$ returns the modulus of x divided by y .

Algorithm 1 summarizes the shadow projection algorithm.

Figure 3.3 gives a graphical explanation of the algorithm applied to an object composed by the three facets $F^{(1)}$ (blue), $F^{(2)}$ (brown) and $F^{(3)}$ (lilac), which extend indefinitely along the direction perpendicular to the plane. These facets and their projections are represented by solid lines of the respective colour—and a dot, in the degenerate case of $F_s^{(2)}$. The dashed lines are the projection planes, symbolized by U_s and U_r , with normals u_s and u_r , respectively. Thick grey overlays represent shadows; and the also grey arrows *track* the shadow projection between facets and from U_s to U_r (the shadow lifting operation). The projections are drawn with an offset to their respective planes, to indicate the distance w.r.t. the illumination source or the observer. This notion is key to decide which of the overlapping facets is shadowing the other. (Vallverdu Cabrera et al., 2021a)

Algorithm 1 Shadow projection

Input: $u_s, u_r, \{F^{(k)} \mid k \in 1 \dots N_f\}$ **Output:** $\{\mathcal{A}_r^{(k)} = \cos \theta_r^{(k)} \mathcal{A}^{(k)} \mid k \in 1 \dots N_f\}$

```

1: for  $k \leftarrow 1, N_f$  do
2:    $F_r^{(k)} \leftarrow$  project  $F^{(k)}$  onto  $U_r : r^T u_r = 0$  (Equation 3.56)
3:    $F_s^{(k)} \leftarrow$  project  $F^{(k)}$  onto  $U_s : r^T u_s = 0$  (Equation 3.57)
4:    $V^{(k)} \leftarrow F_r^{(k)}$  if  $\delta(\theta_s^{(k)}, \theta_r^{(k)}) = 1$  else  $\emptyset$ 
5: end for
6: for all  $j, k \in \mathbb{N} \mid 1 \leq j < k \leq N_f$  do
7:   if  $V^{(j)} = V^{(k)} = \emptyset$  then cycle
8:   if  $F_r^{(j)} \wedge F_r^{(k)} = \emptyset$  (Mol, 2022) then cycle
9:    $F_r^{(j \wedge k)} \leftarrow F_r^{(j)} \wedge F_r^{(k)}$  (Vatti, 1992)
10:   $f, c \leftarrow$  FINDCLOSEST( $j, k, r$ )
11:   $V^{(f)} \leftarrow V^{(j)} \setminus F_r^{(c)}$  (Vatti, 1992)
12: end for
13: for all  $j, k \in \mathbb{N} \mid 1 \leq j < k \leq N_f$  do
14:   if  $V^{(j)} = V^{(k)} = \emptyset$  then cycle
15:   if  $F_s^{(j)} \wedge F_s^{(k)} = \emptyset$  (Mol, 2022) then cycle
16:    $F_s^{(j \wedge k)} \leftarrow F_s^{(j)} \wedge F_s^{(k)}$  (Vatti, 1992)
17:    $f, c \leftarrow$  FINDCLOSEST( $j, k, i$ )
18:    $S^{(c)} \leftarrow$  lift  $F_s^{(j \wedge k)}$  onto  $U_r$  (Equation 3.61)
19:    $V^{(f)} \leftarrow V^{(j)} \setminus S^{(c)}$  (Vatti, 1992)
20: end for
21:  $\mathcal{A}_r^{(k)} \leftarrow$  Equation 3.63 applied on  $V^{(k)}$  for  $k \leftarrow 1, N_f$ 

22: function FINDCLOSEST( $a, b, d$ )
23:    $\bar{v}_d^{(a \wedge b)} \leftarrow$  average of vset( $F_d^{(a \wedge b)}$ )
24:    $\bar{v}^{(k)} =$  lift( $\bar{v}_d^{(a \wedge b)}; F_d^{(k)}, F^{(k)}$ ) for all  $k \in \{a, b\}$ 
25:    $f, c \leftarrow a, b$  if  $\text{dist}(\bar{v}^{(a)}) > \text{dist}(\bar{v}^{(b)})$  else  $f, c \leftarrow b, a$  (Equation 3.59)
26:   return  $f, c$ 
27: end function

```

3. Light Curve Model

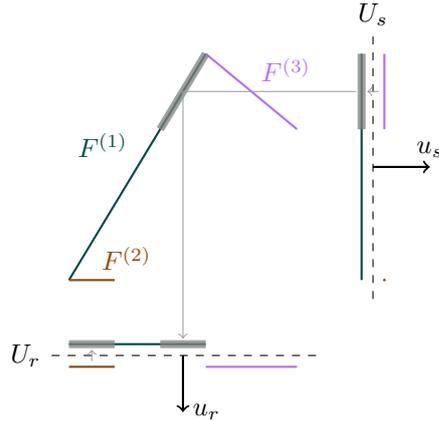


Figure 3.3.: 2D illustration of the shadow projection algorithm (Vallverdú Cabrera et al., 2021a)

Following the steps of the shadow projection algorithm in the case of Figure 3.3, first $F^{(1)}$ is detected to be partially covered by $F^{(2)}$ in the u_r direction. In the second part, the algorithm detects that $F^{(3)}$ blocks some of the incident light from reaching another portion of $F^{(1)}$, and this shadow is transported all the way to the U_r plane. Thus, $V^{(1)}$ ends up being the non-shadowed part of $F_r^{(1)}$, as shown in the image. In this example, $V^{(2)} = V^{(3)} = \emptyset$ due to Step 4 of Algorithm 1—*i.e.* $F^{(2)}$ is parallel to the light rays, while $F^{(3)}$ receives the light from the opposite side from which the observer sees it.

In terms of speed, this algorithm needs to check twice (once in each projection) the intersection of the 2-combinations of the N_f set of facets. This is, it needs to perform $N_f(N_f - 1)$ polygon intersection operations plus at most this many polygon subtractions (one per non-empty intersection) and half as many $i \rightarrow r$ projections (one per non-empty intersection on U_s). Thus, for big N_f , the computation time scales with $O(N_f^2)$.⁴ Furthermore, all iterative steps, except the subtraction operation, which occurs only seldom, could be parallelized. (Vallverdú Cabrera et al., 2021a)

One final advantage of this algorithm is that areas visible from the observer point of view, regardless of the illumination condition, can be already extracted at Step 4. This is useful to account for radiated light, more significant in the infrared spectrum. (Vallverdú Cabrera et al., 2021a)

Compared to the alternatives in the literature, both ray tracing Früh et al. (2014) and hidden vertex removal Wang et al. (2018) methods provide an approximate solution to the self-shadowing problem (given the 1-bounce hypothesis). Instead, the algorithm proposed here gives an exact solution. It does not discretize the surfaces in pixels (Früh et al., 2014), nor treats the object facets as vertices when checking if they are shadowed or not (Wang et al., 2018). Instead, the exact shadowed portions of each facet are subtracted, down to the numerical precision of the machine used. Where this method stands out the most, is when the object is defined by a few, big facets. In this case, ray tracing would still need

⁴The term $O(g(x))$ is the Landau notation that indicates that, if $g : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$, there exist $\delta, M \in \mathbb{R}_{>0}$ such that $|O(g(x))| \leq Mg(x)$ for any $0 < |x| < \delta$.

to discretize the surfaces thinly to reach acceptable accuracy, and Wang et al. (2018) would commit huge errors by assuming that facets can be treated as points. The Shadow Projection algorithm, however, solves such a case with a few iterations, because only a few facets need to be compared against each other, and gives an exact solution.

3.4. Evaluation

The Shadow Projection algorithm proposed in the previous section has been compared against the Möller et al. (1997) ray tracing approach used *e.g.* in Fan (2020). Table 3.1 contains the configuration of the shape used to test these two algorithms. Table 3.1a contains the list of vertices, while Table 3.1b contains the sets of three vertex indices that define each of the 14 facets of the shape. The shape is in fact a cube, but where one of its faces is concave, forming an inverted pyramid whose vertex is at the origin—these are the last 4 facets in Table 3.1b. Figure 3.4 shows the front and rear views of this shape.

In the test, the following observation (u_r) and Sun (u_s) vectors have been used:

$$u_r = [\cos \theta_s \cos \phi_s \quad \cos \theta_s \sin \phi_s \quad \sin \theta_s]^T, \quad (3.64a)$$

$$u_s = [\cos \theta_r \cos \phi_r \quad \cos \theta_r \sin \phi_r \quad \sin \theta_r]^T, \quad (3.64b)$$

where θ_r and θ_s are the elevations relative to the x - y plane of the body frame, while ϕ_r and ϕ_s are the azimuths around z from x to y .

First, both algorithms have been run against all the possible combinations of $\theta_r, \theta_s \in \{-90^\circ, -60^\circ, -30^\circ, 0^\circ, 30^\circ, 60^\circ, 90^\circ\}$ and $\phi_r, \phi_s \in \{0^\circ, 30^\circ, 60^\circ, 90^\circ, 120^\circ, 150^\circ, 180^\circ\}$. There are 7 values per parameter, so this results in a total of $N_t = 7^4 = 2401$ tests. In the particular case of the ray tracing algorithm, the pixels used have a resolution of 0.01×0.01 units squared. This is, one face of the cube is exactly covered by $200 \times 200 = 4 \cdot 10^4$ pixels. To compare the two algorithms, their results have been evaluated using the L1-norm of the facet-by-facet area error, averaged over all tests:

$$\epsilon = \frac{1}{N_t} \sum_{i=1}^{N_t} \sum_{k=1}^{N_f} \left| \mathcal{A}_{r,RT}^{(k)} - \mathcal{A}_{r,SP}^{(k)} \right| (\theta_r^{(i)}, \phi_r^{(i)}, \theta_s^{(i)}, \phi_s^{(i)}), \quad (3.65)$$

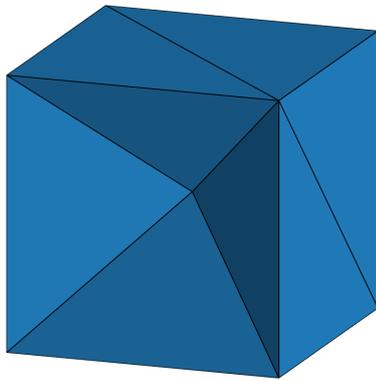
where $N_f = 14$ is the number of facets, and the RT (ray tracing) and SP (shadow projection) subscripts indicate the algorithm used to calculate the projected area ($\mathcal{A}_r^{(k)}$). Each i indicates the unique combination of angles that corresponds to test i . The result for this test case is $\epsilon \simeq 0.4$ square units. For a shape whose characteristic orthogonal projection is of 4 square units (the area of one of the faces of the cube), this is a 10% average error on recovered area. Meanwhile, the ray tracing algorithm needed 46 times the computation time of the shadow projection algorithm, on average.

These results show how the shadow projection algorithm introduced in this dissertation is clearly superior to the classical ray tracing alternative, when the shape is simply defined by just a few facets. At 46 times the computational cost of the shadow projection algorithm, the ray tracing one was able to produce results with an average 10% error. Of course,

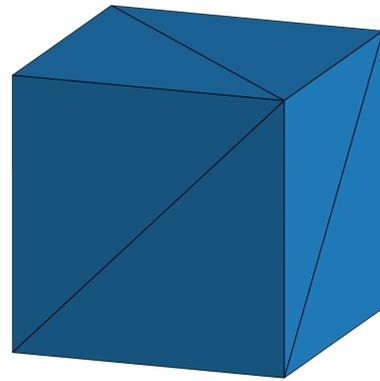
3. Light Curve Model

Table 3.1.: Shadow Projection test shape

(a) Vertices				(b) Facets			
Vertex	x	y	z	Facet	Vertex 1	Vertex 2	Vertex 3
1	-1.0	-1.0	-1.0	1	1	3	4
2	+1.0	-1.0	-1.0	2	1	4	2
3	-1.0	+1.0	-1.0	3	5	6	7
4	+1.0	+1.0	-1.0	4	6	8	7
5	-1.0	-1.0	+1.0	5	1	5	7
6	+1.0	-1.0	+1.0	6	1	7	3
7	-1.0	+1.0	+1.0	7	2	4	8
8	+1.0	+1.0	+1.0	8	2	8	6
9	+0.0	+0.0	+0.0	9	1	2	6
				10	1	6	5
				11	3	7	9
				12	7	8	9
				13	8	4	9
				14	4	3	9



(a) Front



(b) Back

Figure 3.4.: Shadow Projection test shape

increasing the ray tracing algorithm resolution would improve upon this error quadratically, but the computational time needed would also increase quadratically.

To close this section, Figure 3.5 shows the particular result of the shadow projection algorithm for the case of $\theta_r = 60^\circ$, $\phi_r = 80^\circ$, $\theta_s = 20^\circ$ and $\phi_s = 145^\circ$. The surfaces shaded in green are the parts directly visible from the observer location (regardless of illumination), while the orange ones are those visible from the observer and illuminated at the same time. The green vector indicates the observer direction $u_r(\theta_r, \phi_r)$, while the orange one is the Sun direction $u_s(\theta_s, \phi_s)$. In particular, Figures 3.5c and 3.5d are the views from the observer and Sun points of view (*i.e.* orthogonal to u_r and u_s), respectively. This is why the corresponding vectors do not appear in these two figures.⁵

Examples of the full light curve simulation pipeline are shown in the following chapters, in the role of input data. Since they do not constitute further innovation, it has been deemed unnecessary to include these examples in the current chapter.

3.5. Conclusion

This chapter has presented the light curve model that is used in the remaining of this thesis. It has introduced the *foreground*, or necessary theory on which the rest of the research exposed below is based.

The scientific contribution of this chapter is the Shadow Projection algorithm. This algorithm solves the self-shadowing problem of the rendering step of a light curve. This is, for non-convex objects, it calculates which parts of an object block the Sun light from reaching which other parts, and then it finds which parts block the observer from seeing which other parts. Compared to previous methods, the Shadow Projection algorithm provides an exact solution, not having to compromise at some finite spatial resolution. It excels when the object shape can be described as a polyhedron with a few, big triangular facets. In this case, its computational speed is greater than traditional methods used to solve the self-shadowing problem for light curve simulation.

The Shadow Projection algorithm has been tested against the traditional ray tracing approach. For a simple, non-convex test shape, it has been proven to run ~ 46 times faster than the ray tracing alternative, for the case when ray tracing could only achieve 10% projected area error.

⁵Note that the result of the Shadow Projection algorithm is only the projected areas on observer plane $\{\mathcal{A}_r^{(k)}\}$. To obtain this figure, the lifting function described above was used to transport the intermediate results $\{V^{(k)}\}$ from Algorithm 1 back to the three-dimensional facets of the object.

3. Light Curve Model

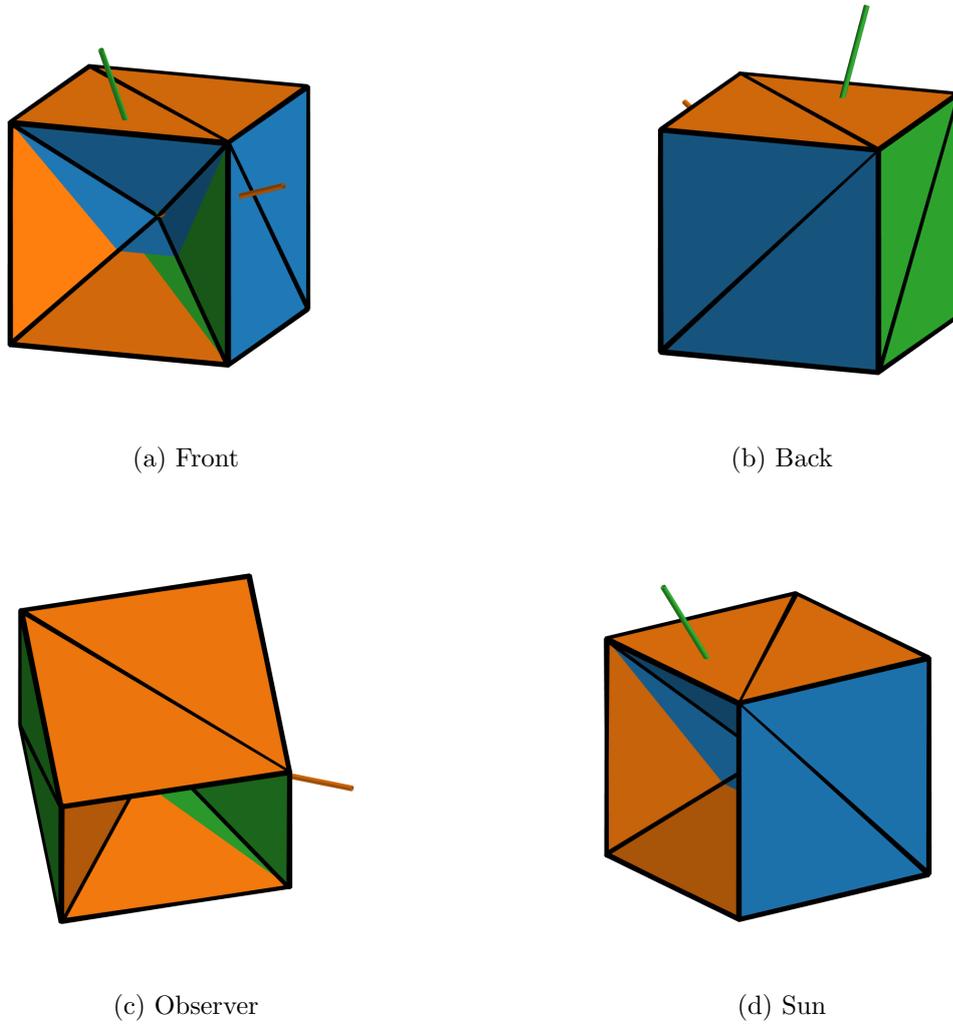


Figure 3.5.: Illumination scenario for $\theta_r = 60^\circ$, $\phi_r = 80^\circ$, $\theta_s = 20^\circ$ and $\phi_s = 145^\circ$, computed using Shadow Projection

4. Shape Determination

This chapter studies the light-curve-based shape determination algorithm introduced in Kaasalainen et al. (2001) discussed in Section 2.2.3, which estimates an Extended Gaussian Image (EGI) from a light curve, and then uses Minkowski minimization to obtain a unique convex shape in the form of a polyhedron.

Other than the light curve, this method needs the attitude of the object as an input, which may not be available for all objects, especially if dealing with *e.g.* space debris or non-collaborative satellites. Furthermore, it can only estimate convex shapes. This is, if the object observed is not convex, a convex-equivalent shape is estimated instead—*i.e.* the convex shape that would produce the closest light curve to the measured one.

Nonetheless, this method is still relevant as a potential building block to more complex approaches that may exploit it directly, which would benefit from the insights gained studying it. An example is the attitude-shape sequential estimation technique from Linares et al. (2018), where a similar approach is embedded in a Marginalized PF.

Section 4.1 that follows describes the core algorithm itself, based on the literature. Next, Section 4.2 proposes a new metric to assess how well the final shape recovered by the algorithm resembles the true shape. Section 4.3 follows with the derivation of a potential improvement to the numerical solution to Minkowski minimization: the derivatives of the volume of a polyhedron relative to its supports. It is followed by Section 4.4, which revisits the observability of the EGI as a function of its tessellation density. Section 4.5 describes then the multiple scenarios used to assess these potential improvements, and Section 4.6 presents and discusses the subsequent results. Section 4.7 proposes some ways forward to overcome some of the shortcomings of this method. The chapter closes with Section 4.8, with a summary of the findings herein.

This chapter is adapted from and expands on Vallverdú Cabrera et al. (2021b), the second paper published as part of the doctoral research of this thesis. As such, parts of Vallverdú Cabrera et al. (2021b) are included herein.

4.1. Problem Definition

This section explains the shape determination algorithm that recovers a convex shape from a light curve, mainly based on the algorithm introduced by Kaasalainen et al. (2001). The process is divided in two steps:

1. first, an Extended Gaussian Image (EGI) is estimated from the input light curve and attitude profile;

4. Shape Determination

2. second, a shape, formally represented by a convex polyhedron in 3D, is recovered from the EGI.

The following definition establishes the nomenclature used for a polyhedron for the remaining of this chapter:

Definition 1. Let \mathcal{P} be a **polyhedron** in \mathbb{R}^3 whose outer closed surface is bound by N_f flat facets. Let the facets be indexed within the subspace of the natural numbers

$$\mathcal{L} = \{1 \dots N_f\}. \quad (4.1)$$

Each facet can then be referred to as $F^{(l)} \subset \delta\mathcal{P}$, being $\delta\mathcal{P}$ the surface of \mathcal{P} . Each facet has area A_l and a unit normal vector n_l .

4.1.1. Estimating the EGI from a Light Curve

The Extended Gaussian Image (EGI) is a parametrization of a convex shape. Its formal definition (Horn, 1984) has been included in Appendix E.1. In particular, it can represent a convex polyhedron (i.e. the particular case of a finite number of flat facets); the EGI of (the surface of) a convex polyhedron \mathcal{P} is (Horn, 1984; Little, 1985):

Definition 2. Given the polyhedron \mathcal{P} from Definition 1, the **Extended Gaussian Image** of (the surface of) \mathcal{P} can be expressed as the set of its normal-area pairs:

$$\text{EGI}_{\mathcal{P}} = \{A_i n_i \mid i \in 1 \dots N_f\}. \quad (4.2)$$

As long as

$$\sum_{i=1}^{N_f} A_i n_i = 0, \quad (4.3)$$

the EGI uniquely represents a convex polyhedron (Little, 1985)¹.

The EGI is estimated from a diverse-enough (Friedman et al., 2019) light curve by solving the minimization problem

$$\min_{\mathbf{a}} J = \left\| \mathcal{L} - \hat{\mathcal{S}} \mathbf{a} \right\|^2. \quad (4.4)$$

where $\mathcal{L} \in \mathbb{R}^{N_t}$ contains all the light curve measurements taken at N_t different times; $\mathbf{a} \in N_a$, the vector of albedo-area products of an EGI with N_a predefined directions; and $\hat{\mathcal{S}} \in \mathbb{R}^{N_t \times N_a}$, the reflection matrix.

The EGI needs to be predefined with N_a directions, using a tessellation scheme like *e.g.* the quad-cube used by Friedman et al. (2019). This scheme is further developed in Appendix E.2.

Each element $\hat{\mathcal{S}}_{ij}$ of the reflection matrix contains all the parameters of the light curve simulation model from Chapter 3 except for albedo ρ and area A , which are contained in $\mathbf{a}_j = \rho_j A_j$. This method assumes that all the albedo values of each element of the EGI are equal, $\rho_j = \rho$. Then, \mathbf{a} can still be used as the *areas* of an EGI, because isotropic scaling

¹See Appendix E.

does not invalidate the zero-residual condition imposed by Equation 4.3. Thus, its equation is

$$\hat{\mathcal{S}}_{ij} = \frac{1}{\rho} \delta(\theta_s^{(j)}, \theta_r^{(j)}) \frac{A_a}{d_r^2} \cos \theta_a^{(j)} \cos \theta_r^{(j)} \cos \theta_s^{(j)} \int_{B_\lambda} f_r^{(j)}(\lambda, u_s, u_r) \tau(\lambda) E_{s,\lambda}^\perp(u_s) \frac{\lambda}{hc} d\lambda. \quad (4.5)$$

This minimization must be subject to two constraints that bound \mathbf{a} .

- First, it makes no physical sense that either areas or albedos are negative. Thus, the area-albedo product must be always positive. This is, $\mathbf{a}_j > 0$ for all $j \in 1 \dots N_a$.
- Second, for the EGI to uniquely represent a convex polyhedron, the condition $\sum_{i=1}^{N_f} A_i n_i = 0$ (See Equation 4.3) needs to be satisfied.

These two points are discussed below, followed by a note on noise-mitigation techniques within the scope of the light curve to EGI step.

Positivity of \mathbf{a} . This regression should be bound so that $\mathbf{a} > 0$, because a negative albedo-area has no physical meaning. There are different approaches by which this condition can be achieved.

- Kaasalainen et al. (2001) suggest, for the application of this method on asteroid shape recovery, to substitute \mathbf{a}_j by $\exp(b_j)$, and then optimize for $b = [b_1 \dots b_{N_a}]^T$. This renders the least squares non-linear, which then they solve using the conjugate gradient method. Optimizing for unlimited b naturally guarantees positive \mathbf{a} . Furthermore, the exponential function being bijective grants that a unique minimum of J in b corresponds to the same minimum in \mathbf{a} .
- For use with artificial satellites, Calef et al. (2006) and Fan et al. (2020) use constrained optimization directly on \mathbf{a} .

The current work employs the Non-Negative linear Least Squares (NNLS) method from Lawson et al. (1995), which aligns with the second option.

Convexity of the Recovered EGI. If the Recovered EGI is to be translated into a unique convex polyhedron, the convexity condition expressed by Equation 4.3 has to be observed. However, the optimal solution to Equation 4.4 does not necessarily comply with this condition. The possible reasons may be noise in the measurement (Fan et al., 2019; Fan et al., 2020) or uncertainty in the model. The latter may come from different sources:

- First, the assumptions used to construct $\hat{\mathcal{S}}$ (*i.e.* reflective model, attitude, etc.) may not match the observed object closely enough.
- Second, the tessellation scheme used to generate the EGI may miss important facets from the observed object, which causes deviations on the convexity condition.
- Third, albedo variegation may cause the recovered EGI not to represent a convex shape, because the varying albedos of the different facets of the object deform the \mathbf{a} vector out of convexity (Kaasalainen et al., 2001).

4. Shape Determination

- Finally, noise and biases in the light curve can cause the convexity residual to be non-zero.

In the implementation evaluated here, deviations from convexity are corrected by using the simple least squares solution suggested by Little (1985). This is, being \mathbf{a} the result of the unconstrained problem from Equation 4.4 under NNLS, a second least squares is performed afterwards to find an alternative solution \mathbf{a}^* that does comply with Equation 4.3. This is

$$\begin{aligned} \min_{\mathbf{a}^*} \quad & J = \|\mathbf{a} - \mathbf{a}^*\|^2, \\ \text{subj. to} \quad & \sum_{j=1}^{N_a} n_j \mathbf{a}^* = 0. \end{aligned} \quad (4.6)$$

Putting all the normal vectors into a matrix

$$\mathbf{n} = [n_1 \quad \cdots \quad n_{N_a}], \quad (4.7)$$

this has the analytical solution

$$\hat{\mathbf{a}}^* = \hat{\mathbf{a}} - \mathbf{n}^T (\mathbf{n} \mathbf{n}^T)^{-1} \mathbf{n} \hat{\mathbf{a}}, \quad (4.8)$$

In words, \mathbf{a}^* is the closest approximation to the unconstrained albedo-areas solution \mathbf{a} in the least squares sense that, at the same time, lies within the subspace where Equation 4.3 is true. Although forcing this constraint ensures that the obtained EGI corresponds to a unique convex shape, it may add unwanted biases to the final solution.

Other solutions exist to enforce this condition. Kaasalainen et al. (2001) use regularization directly on Equation 4.4 to enforce this condition, while Fan et al. (2021) use a second least squares approach similar to the one presented above, but that instead seeks to minimize $\|\sum_{j=1}^{N_a} n_j \mathbf{a}^*\|^2$ subject to $0 < \mathbf{a}_j^* - \mathbf{a}_j < D_j$, where D_j is an upper bound for the change in \mathbf{a} . Both approaches seek the same goal as the one used by Little (1985), which is preferred in the current work because of its analytical solution.

Stability against Noise. To guarantee enough detail on the recovered shape, it is necessary to sample the EGI on the unit sphere with a tessellation scheme that is dense enough (Friedman et al., 2019). However, the computational cost of the Minkowski minimization problem increases with the number of directions of the EGI.

Nonetheless, due to noise in the signal and numerical errors, in practice EGIs recovered from light curves may present a relatively high number of small, non-zero values. These destabilize the Minkowski minimization process (Calef et al., 2006), and can be mitigated in different ways. Fan et al. (2020) do so by artificially zeroing all albedo-area values below a certain threshold. Furthermore, they deal with noisy light curves by employing two separate consecutive light curves (Fan et al., 2020) or by wrapping the shape inversion process in a particle-filter-based estimator using importance resampling (Fan et al., 2021). Calef et al. (2006) use Tikhonov (a.k.a. ridge) regularization,

$$\min_{\mathbf{a}} J = \left\| \mathcal{L} - \hat{\mathcal{S}} \mathbf{a} \right\|^2 + (\alpha \|\mathbf{a}\|)^2, \quad (4.9)$$

and suggest a cross-validation scheme to choose the best value for α . Additionally, this dissertation proposes Lasso regularization,

$$\min_{\mathbf{a}} J = \left\| \mathcal{L} - \hat{\mathbf{S}}\mathbf{a} \right\|^2 + \alpha \sum_{j=1}^{N_a} |\mathbf{a}_j|, \quad (4.10)$$

to help reduce noise sensitivity, because of its sharpening properties²—artificial objects are likely to have less, bigger facets than *e.g.* asteroids, with sharper edges.

4.1.2. Generating a Convex Polyhedron from an EGI

The second step is called Minkowski minimization, and recovers a convex polyhedron from a finite EGI like the one defined by Equation 4.2.

As explained by Little (1985), if the list of supports and corresponding unitary normals of a convex polyhedron are known, the vertices, edges and facets of the same can be recovered. The list of supports describes the location of all the planes that bind the polyhedron to be recovered:

Definition 3. *Given the polyhedron \mathcal{P} from Definition 1, the **support** of the l^{th} facet of \mathcal{P} is defined as*

$$\mathbf{h}_l = n_l^T x, \forall x \in F^{(l)}, \quad (4.11)$$

where n_l is the unitary out-facing normal of the facet $F^{(l)}$. This is, the support of a facet is the minimum distance between the origin and the plane that contains said facet. (Little, 1985)

The supports are then used to obtain the location of the vertices of the dual of the target polyhedron.

Definition 4. *The dual of a convex polyhedron \mathcal{P} is some other polyhedron \mathcal{P}' such that its vertices are the set of $\{n_l/\mathbf{h}_l\}$, where n_l and \mathbf{h}_l are the normal and support of each facet l of \mathcal{P} —i.e. each facet of \mathcal{P} is mapped onto a vertex of \mathcal{P}' . In the same way, the dual of \mathcal{P}' is \mathcal{P} —i.e. each vertex of \mathcal{P} is mapped onto a facet of \mathcal{P}' . (Little, 1985)*

Next, the target polyhedron is obtained by performing a convex hull of the locations described by the supports as vertices, and then calculating the dual of such hull (Little, 1985). The adjacency information of \mathcal{P} (i.e. which vertices belong to each facet, and which facet is adjacent to each vertex) can be obtained by interchanging facets and vertices in the adjacency information of said convex hull.

The convex hull operation on a cloud of points can be efficiently obtained with the *quickhull* algorithm described in Barber et al. (1996). This produces the facets of said convex hull as a set of lists of vertices that belong to each facet. Simple iteration through this set provides the inverse one: the set of lists of facets that touch each vertex.

Overall, this whole operation can be denoted as the function $\mathcal{P}(\mathbf{h})$, which *obtains the vertices of a polyhedron and its adjacency information from its list of supports $\mathbf{h} \in \mathbb{R}^{N_f}$.*

²More on the results section.

4. Shape Determination

Minkowski's theorem states that only one unique convex polyhedron exists for each EGI. Since one can recover the polyhedron if one knows its supports, it must be possible to recover the supports from the EGI. Little (1985) describes this process. As explained therein, the Brunn-Minkowski theorem states that given two polyhedrons \mathcal{P} and \mathcal{Q} with the same number of facets oriented along the same set of normal vectors n_i , the product $\mathfrak{h}_{\mathcal{P}}^T A_{\mathcal{Q}}$ is minimum at $\mathfrak{h}_{\mathcal{P}}$ if \mathcal{P} and \mathcal{Q} are homothetic, given a constant³ volume of \mathcal{P} .

Definition 5. *Two 3D polyhedrons \mathcal{P} and \mathcal{Q} are **homothetic** if there exist $s \in \mathbb{R}_{>0}$ and $t \in \mathbb{R}^3$ such that*

$$\mathcal{P} = \{x \mid x = sy + t, y \in \mathcal{Q}\} , \quad (4.12)$$

This is, two polyhedrons are homothetic if there exists a scaling and translation transformation that makes one equal to the other. (Little, 1985)

Thus, he proposes to solve the following minimization problem:

$$\begin{aligned} \min_{\mathfrak{h}} \quad & J = \mathfrak{h}^T A_{\mathcal{Q}} , \\ \text{subj. to} \quad & V(\mathcal{P}(\mathfrak{h})) = 1 , \\ & \mathfrak{h} > 0 , \end{aligned} \quad (4.13)$$

where $A_{\mathcal{Q}}$ is the vector of areas from the EGI. Then, the optimization result $\mathfrak{h}_{\mathcal{P}}$ is the vector of the supports of a unit volume polyhedron (\mathcal{P}) that is homothetic to \mathcal{Q} , the one uniquely defined by the areas of the EGI. The function $V(\mathcal{P})$ returns the volume of the polyhedron \mathcal{P} (Little, 1985):

$$V(\mathcal{P}) = \frac{1}{3} \mathfrak{h}_l^T A . \quad (4.14)$$

This is a problem of constrained convex minimization, and can be solved in various ways. Little (1985) uses a gradient-based approach by which he first initializes the supports vector \mathfrak{h} to unity. Then, he

1. constructs the polyhedron \mathcal{P} from \mathfrak{h} (as explained above),
2. computes the volume V from \mathcal{P} and its gradient ∇V with \mathfrak{h} ,
3. scales \mathfrak{h} by $V^{-1/3}$ to enforce the unit volume constraint (as per Equation 4.19), and finally
4. modifies \mathfrak{h} by a step parallel to ∇V projected onto the hyperplane of $V(\mathfrak{h}) = 1$.

This is repeated until \mathfrak{h} does not vary significantly.

4.2. Homotheticity Distance

The first scientific contribution of this chapter is to introduce a systematic way to evaluate the performance of the EGI plus Minkowski minimization method for shape determination, against any range of scenario parameters and implementation approaches.

³since any constant value works, the simplest possible is used: unity.

Previous studies (Calef et al., 2006; Fan et al., 2019) on this shape determination method assess its performance by matching how close the EGI result is (e.g. in root-mean squared sense) to the known EGI of the test polyhedron. However, this leaves the EGI to polyhedron step out of the assessment.

To assess the success of each light curve inversion from end to end (from light curve to polyhedron), this section introduces a new technique to evaluate how close two distinct polyhedrons are, based on the homotheticity concept (Definition 5). In particular, the proposed criterion checks how far from homotheticity they are.

To do so, first it is necessary to define the remnant volume operation:

Definition 6. *the **remnant volume operation** $\Delta V(\mathcal{P}, \mathcal{Q})$ between two polyhedrons \mathcal{P} and \mathcal{Q} is the volume of the exclusive logical or (\vee) between \mathcal{P} and \mathcal{Q} :*

$$\Delta V(\mathcal{P}, \mathcal{Q}) := V(\mathcal{P} \vee \mathcal{Q}) = V(\mathcal{P}) + V(\mathcal{Q}) - 2V(\mathcal{P} \wedge \mathcal{Q}), \quad (4.15)$$

where \wedge represents the polyhedron clipping intersection and $V(\mathcal{P})$ is the function that returns the volume of \mathcal{P} .

Given two convex polyhedrons defined by their limiting planes, the intersection can be easily computed e.g. by clipping and capping (Ahn et al., 2008).

Then, the proposed criterion is:

Definition 7. *The **homotheticity distance** is*

$$H(\mathcal{P}, \mathcal{Q}) = \min_{s,t} \Delta V(\mathcal{P}, \mathcal{T}(\mathcal{Q}, s, t)), \quad (4.16)$$

where $\mathcal{T}(\mathcal{Q}, s, t)$ scales \mathcal{Q} isotropically by the factor s and translates it by the vector t , as in Equation 4.12. In words, the homotheticity distance is the volume remnant of \mathcal{P} and \mathcal{Q}^* , where \mathcal{Q}^* is scaled and translated from \mathcal{Q} so that $\Delta V(\mathcal{P}, \mathcal{Q}^*)$ is minimum.

The homotheticity distance of two homothetic polyhedrons is 0; the higher it becomes, the farther the two tested polyhedrons will be from homotheticity. In this work, Equation 4.16 is solved with either Powell’s (Powell, 1964) or L-BFGS-B (Zhu et al., 1997) minimization algorithms.

To be able to compare different target shapes, $H(\mathcal{P}, \mathcal{Q})$ is normalized by the volume of \mathcal{P} ,

$$\bar{H}(\mathcal{P}, \mathcal{Q}) := \frac{H(\mathcal{P}, \mathcal{Q})}{V(\mathcal{P})}. \quad (4.17)$$

4.3. Derivatives of the Volume over the Supports

The second contribution of this chapter is to explore potential improvements on the Minkowski minimization part of the shape determination algorithm—i.e. from EGI to polyhedron.

The method proposed by Little (1985) has two points of potential improvement. First, the gradient calculated in Little (1985) is erroneous, since he differentiates V with \mathfrak{h} from

4. Shape Determination

Equation 4.14, but neglecting that A_l does depend on \mathfrak{h} . This is not critical, because the resulting gradient is a third of the actual one, but still parallel to it. Thus the only consequence is that the method takes more time to converge. Second, as stated by Little himself, this method does not exploit the Hessian information of $V(\mathfrak{h})$, which can be used to speed the convergence of the method.

In Kaasalainen et al. (2001) they use a similar method of gradient descent projected onto the constraint hyperplane, but with a correctly computed gradient. Thus, the first improvement point is covered.

This is not so for the second point. Following the recommendation from Little (1985), the present work incorporates the Hessian information, to improve the convergence speed of the method. In particular, this dissertation uses the trust-region interior point method described in Byrd et al. (1999), which employs both the gradient and Hessian of the objective function and the constraints. The objective function has a straightforward solution. Being $J = \mathfrak{h}^T A_Q = \sum_{i=1}^{N_f} \mathfrak{h}_i A_i$:

$$\frac{\partial J}{\partial \mathfrak{h}_i} = A_i, \quad (4.18a)$$

$$\frac{\partial^2 J}{\partial \mathfrak{h}_i \partial \mathfrak{h}_j} = 0. \quad (4.18b)$$

Note that the A_i in Equation 4.13 are the area values from vector A_Q , not the final areas of $\mathcal{P}(\mathfrak{h})$. Thus, these are fixed and do not depend on \mathfrak{h} .

For the constraint, the derivatives of V with \mathfrak{h} are deduced below.

To do so, first of all, it is necessary to establish the dependence between the volume of a polyhedron and its supports. The following is a version of Little, 1985, Eq. 3.1 ($A = \sum_{i,j} c_{ij} h_i h_j$) expanded to the volume. It will be necessary to find the Hessian with its supports later, as well as to establish the nomenclature and conventions used.

Theorem 1. *Given the polyhedron from Definition 1, there exists a three-dimensional array $C \in \mathbb{R}^{N_f \times N_f \times N_f}$ such that the volume V of a polyhedron, as a function of the supports of its facets, \mathfrak{h} , is*

$$V = \frac{1}{3} \sum_{k=1}^{N_f} \sum_{j=1}^{N_f} \sum_{i=1}^{N_f} \mathfrak{h}_i \mathfrak{h}_j \mathfrak{h}_k C_{ijk}, \quad (4.19)$$

for $i, j, k \in \mathfrak{L}$.

Proof. Let any facet of the polyhedron, identified with index l , have $M \geq 3$ edges, which connect it to M adjacent facets. The first step is to re-index the involved facets using the alternative indexing subspace $\mathfrak{M}_l = \{0 \dots M\}$. This is, given $m \in \mathfrak{M}_l$:

- $m = 0$ corresponds to facet l ;
- $m = 1$ corresponds to an arbitrary facet that is adjacent to facet $m = 0$ along one of its edges;
- $m > 1$ correspond to the other edge-adjacent facets to facet 0, sorted counter-clockwise relative to facet 1 (positive sense around the outer normal of facet 0).

4.3. Derivatives of the Volume over the Supports

This gives rise to the map $\mathcal{I}_l : \mathfrak{M}_l \rightarrow \mathfrak{L}$ such that $\mathcal{I}_l(m)$ gives the corresponding index of facet m onto the \mathfrak{L} indexing set of \mathcal{P} . Naturally, $\mathcal{I}_l(0) = l$. Thus, the edges of facet 0 can be indexed with $m \in 1..M$, where e_m corresponds to the edge with facet m . Likewise, each vertex of facet 0, denoted v_m , corresponds to the intersection between e_m and $e_{(m \bmod M)+1}$. For the sake of brevity, from here onward, for any $a \in \mathbb{N}_{>0}$, $a < M$, $m + a$ implies $(m \bmod M) + a$ whenever $m + a > M$.

With these definitions, one can express the signed area of facet 0 as

$$A = \frac{1}{2} \sum_{m=1}^M n_0^T (v_m \times v_{m+1})^T, \quad (4.20)$$

where $\{v_m\}$ are the vertices of the facet. The operator \times denotes the cross product. Note that each addend of the above equation is in fact the contribution of the area of the triangle formed between the two involved vertices and the projection of the origin on the plane of facet 0. Some of these contributions may be negative, if the origin is projected outside the facet, thus the area itself can be negative.

Now, each vertex m can be defined as a function of the polyhedron supports and facet outer normals as the solution of the following equation (composed by three instances of Equation 4.11):

$$\begin{bmatrix} n_0^T \\ n_m^T \\ n_{m+1}^T \end{bmatrix} v_m = \begin{bmatrix} \mathfrak{h}_0 \\ \mathfrak{h}_m \\ \mathfrak{h}_{m+1} \end{bmatrix}. \quad (4.21)$$

In a more compact way, defining $\mathbf{n}_m := [n_0 \ n_m \ n_{m+1}]^T$ and $\mathcal{H}_m := [\mathfrak{h}_0 \ \mathfrak{h}_m \ \mathfrak{h}_{m+1}]^T$, this can be solved as

$$v_m = \mathbf{n}_m^{-1} \mathcal{H}_m. \quad (4.22)$$

As long as the three involved facets are in different planes, this system always has a unique solution⁴.

Substituting into Equation 4.20, after some algebraic manipulation, one obtains

$$A = \frac{1}{2} \sum_{m=1}^M \mathcal{H}_{m+1}^T \mathbf{n}_{m+1}^{-1T} [n_0]_{\times} \mathbf{n}_m^{-1} \mathcal{H}_m, \quad (4.23)$$

where $[\cdot]_{\times}$ is the cross-product matrix, such that $[a]_{\times} b$ is equivalent to $a \times b$.

Substituting Equation 4.23 into Equation 4.14 gives

$$V = \frac{1}{3} \sum_{l=1}^L \mathfrak{h}_l \left[\frac{1}{2} \sum_{m=1}^M \mathcal{H}_{m+1}^T \mathbf{n}_{m+1}^{-1T} [n_0]_{\times} \mathbf{n}_m^{-1} \mathcal{H}_m \right]_l. \quad (4.24)$$

The $[\cdot]_l$ notation in this context implies the indices inside the brackets are expressed in the \mathfrak{M}_l system, while elements outside the bracket are indexed in \mathfrak{L} .

From here onward, it is trivial to prove that a three-dimensional array C such as the one described above can be constructed by first filling it with zeros, and then adding the

⁴This implies that any polyhedron with adjacent parallel facets needs to have these facets merged into a single polygon before it can be used in this context.

4. Shape Determination

contribution of each $\left[\frac{1}{2}\mathbf{n}_{m+1}^{-1T}[n_0]_{\times}\mathbf{n}_m^{-1}\right]_l$ sub-matrix into the corresponding C_{ljk} location, for $j = \{\mathcal{I}_l(0), \mathcal{I}_l(m+1), \mathcal{I}_l(m+2)\}$ and $k = \{\mathcal{I}_l(0), \mathcal{I}_l(m), \mathcal{I}_l(m+1)\}$.

It is thus proven that the claimed array C exists. □

Regarding the gradient of the volume of a polyhedron relative to its supports, the proper equation is known (Kaasalainen et al., 2001). However, its proof is used to then deduce the Hessian. Therefore:

Theorem 2. *Given a polyhedron \mathcal{P} in \mathbb{R}^3 with N_f facets, the derivative of its volume V with the support of its l^{th} facet is equal to the area A_l of said facet:*

$$\frac{\partial V}{\partial h_l} = A_l. \quad (4.25)$$

Proof. Let \mathcal{P}' be a second polyhedron, which is almost equal to \mathcal{P} , with the only difference that its l^{th} support is varied by $h'_l = h_l + \epsilon$, with $\epsilon \ll h_l$.

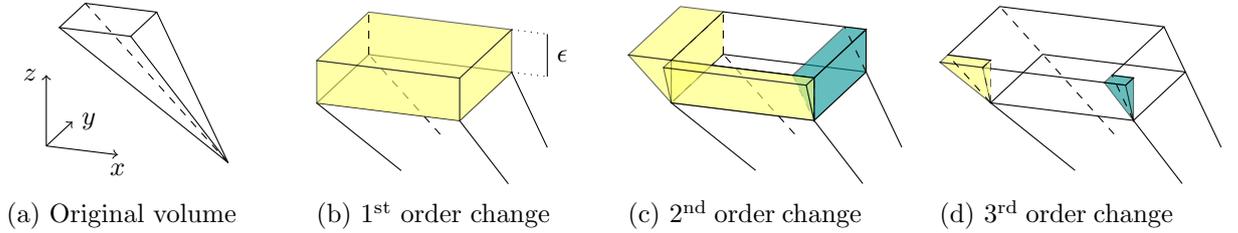


Figure 4.1.: Change of volume of a polyhedron as a function of one of its supports

Figure 4.1 illustrates how such a polyhedron \mathcal{P}' can be obtained from another polyhedron \mathcal{P} . In three steps. Being Figure 4.1a the original \mathcal{P} :

- the first order contribution is the extrusion of the facet l a distance ϵ along its normal, represented by the yellow volume in Figure 4.1b;
- the second order contributions are the volume differences created by extending the adjacent facets until they meet the plane of the newly extruded facet l , represented by the yellow (additions) and green (subtractions) volumes in Figure 4.1c;
- the third order contributions are the volume differences created by extending the facet vertices along their out-of-facet edges, represented in Figure 4.1d with the same color code.

This can be represented analytically by the following equation:

$$V' = V + A_l \epsilon + O(\epsilon^2), \quad (4.26)$$

whose first two components are the original volume and the 1st order contribution to the volume change. The latter is trivial to derive as the volume of the prism formed by the area extrusion associated to the first order contribution. The error term represents the 2nd order contribution and higher.

Applying the definition of the derivative as a limit yields

$$\frac{\partial V}{\partial h_l} = \lim_{\epsilon \rightarrow 0} \frac{V' - V}{\epsilon} = \lim_{\epsilon \rightarrow 0} A_l + O(\epsilon), \quad (4.27)$$

where $O(\epsilon)$ vanishes, so that this evaluates to Equation 4.25. \square

Next follows an analytical expression for the Hessian, original to this thesis in the context of light curve to shape inversion. As a preliminary step, the following corollary is necessary:

Corollary 1. *The double derivative of the volume of a polyhedron \mathcal{P} with the supports of two of its facets is equivalent to the derivative of the area of either facet with the support of the other:*

$$\frac{\partial^2 V}{\partial h_l \partial h_p} = \frac{\partial A_l}{\partial h_p} = \frac{\partial A_p}{\partial h_l}. \quad (4.28)$$

Proof. This is trivial to prove using the chain rule of the derivative on Equation 4.25. Furthermore, the fact that V can be expressed as a function of the polyhedron supports (Equation 4.19) guarantees that partial mixed derivatives can be interchanged (Aksoy et al., 2002). \square

With it, the equation for the Hessian can be deduced:

Theorem 3. *Given a polyhedron \mathcal{P} in \mathbb{R}^3 with N_f facets, the double derivative of its volume V with the support of its l^{th} and p^{th} facets is:*

$$\frac{\partial^2 V}{\partial h_l \partial h_p} = \begin{cases} - \left[\sum_{m=1}^M \|e_m\| \frac{n_0^T n_m}{\sqrt{1 - (n_0^T n_m)^2}} \right]_l & \text{if } l = p, \\ \left[\frac{\|e_m\|}{\sqrt{1 - (n_0^T n_m)^2}} \right]_p \mid m = \mathcal{I}_p^{-1}(l) & \text{if } p \text{ is adjacent to } l, \\ 0 & \text{else.} \end{cases} \quad (4.29)$$

Proof. The proof is split in three parts, one for each case in Equation 4.29.

Part 1: $l = p$

This proof starts with the $l = p$ case. It is necessary to expand Equation 4.26 to include second order contributions, too:

$$V' = V + A_l \epsilon - \frac{1}{2} \left[\sum_{m=1}^M \|e_m\| \frac{n_0^T n_m}{\sqrt{1 - (n_0^T n_m)^2}} \right]_l \epsilon^2 + O(\epsilon^3). \quad (4.30)$$

The example on Figure 4.2 is used to derive the 2nd order contribution along any edge e_m . The figure depicts a zoom on the leftmost edge of the polyhedron in Figure 4.1c, as seen

4. Shape Determination

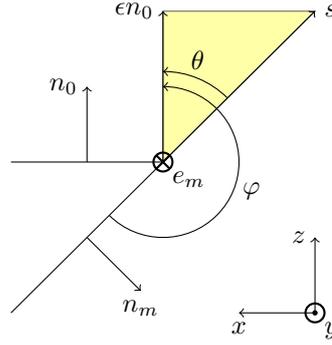


Figure 4.2.: Zoom on 2nd order contribution to the volume versus support change

from the tip of the y axis. Thus, this volume contribution is none other than the volume of a prism, obtained by multiplying its length (the facet edge) with the area of its base:

$$\Delta V_{0,m}^{2^{nd}} = \frac{1}{2} e_m^T (\epsilon n_0 \times s). \quad (4.31)$$

Note that, from Figure 4.2,

$$n_0^T n_m = \cos \varphi = -\sin \theta,$$

hence

$$|\cos \theta| = \sqrt{1 - (n_0^T n_m)^2}. \quad (4.32)$$

Furthermore, by realizing that

$$\|s\| = \frac{\epsilon}{|\cos \theta|},$$

together with

$$n_0 \times s = \|s\| \sin \theta \frac{e_m}{\|e_m\|},$$

one can easily rewrite Equation 4.31 to look like the third term in Equation 4.30.

Next, it is necessary to define a third polyhedron \mathcal{P}'' whose l^{th} support is increased by 2ϵ , so that its volume is

$$V'' = V + 2A_l \epsilon - 2 \left[\sum_{m=1}^M \|e_m\| \frac{n_0^T n_m}{\sqrt{1 - (n_0^T n_m)^2}} \right]_l \epsilon^2 + O(\epsilon^3). \quad (4.33)$$

This is easy to obtain, just substitute ϵ in Equation 4.30 with 2ϵ .

Now one can apply again the limit-definition of the derivative, but this time to the double derivative:

$$\begin{aligned} \frac{\partial^2 V}{\partial h_l^2} &= \lim_{\epsilon \rightarrow 0} \frac{\frac{V'' - V'}{\epsilon} - \frac{V' - V}{\epsilon}}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{V'' - 2V' + V}{\epsilon^2} = \\ &= \lim_{\epsilon \rightarrow 0} - \left[\sum_{m=1}^M \|e_m\| \frac{n_0^T n_m}{\sqrt{1 - (n_0^T n_m)^2}} \right]_l + O(\epsilon). \end{aligned} \quad (4.34)$$

After taking the limit, this is exactly the $l = p$ case of Equation 4.29.

The same result can be reached if working with the areas, using Corollary 1. From Figure 4.2, it is trivial to see that the first order area contribution to facet 0 due to the movement of edge m provoked by the support increase is

$$\Delta A_{0,m}^{1st} = n_0^T [(s - \epsilon n_0) \times e_m] = e_m^T [n_0 \times (s - \epsilon n_0)] = 2 \frac{\Delta V_{0,m}^{2nd}}{\epsilon}. \quad (4.35)$$

This is just the area of the polygon (perpendicular to the plane in Figure 4.2) whose sides are the edge e_m and the top of the yellow triangle. The contribution can be positive or negative depending whether n_0 is parallel or antiparallel to the cross product in the equation. As indicated, this corresponds to the same effect that generates the second order contribution to the volume.

Thus, $\partial^2 V / \partial h_l^2$ can be computed by taking the limit of $\epsilon \rightarrow 0$ of the sum of the area contribution of all adjacent edges:

$$\frac{\partial^2 V}{\partial h_l^2} = \frac{\partial A_l}{\partial h_l} = \lim_{\epsilon \rightarrow 0} \left[\sum_{m=1}^M \frac{\Delta A_{0,m}^{1st} + O(\epsilon^2)}{\epsilon} \right]_l. \quad (4.36)$$

This gives exactly the same result as Equation 4.34.

Part 2: l and p adjacent

Next comes the second case, where p is adjacent to l . One can prove it using Corollary 1 again. From Figure 4.2, it is evident that, when facet 0 increases its support by ϵ , the area of facet m increases as

$$A'_m = A_m + \|s \times e_m\| + O(\epsilon^2) = A_m + \frac{\|e_m\|}{|\cos \theta|} \epsilon + O(\epsilon^2) \quad (4.37)$$

With this, using again the limit definition of the derivative, as well as substituting with Equation 4.32, it follows that

$$\frac{\partial A_m}{\partial h_0} = \lim_{\epsilon \rightarrow 0} \frac{A'_m - A_m}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{\|e_m\|}{\sqrt{1 - (n_0^T n_m)^2}} + O(\epsilon). \quad (4.38)$$

Corollary 1 states that

$$\left. \frac{\partial^2 V}{\partial h_0 \partial h_m} \right|_{m \neq 0} = \frac{\partial A_m}{\partial h_0}. \quad (4.39)$$

Therefore, if one considers Figure 4.2 to be indexed in \mathfrak{M}_p , so that $l = \mathcal{I}_p(m)$ and $p = \mathcal{I}_p(0)$, Equation 4.38 becomes the $l \neq p$ case of Equation 4.29.

Part 3: l and p not adjacent

Finally, for the third case, as long as the adjacency information of the polyhedron does not change⁵, it is obvious that changing the support of a facet only affects the area of the facet

⁵If a change in a support produces e.g. a vertex to become an edge, the derivative of area of some facets becomes discontinuous. If e.g. a facet becomes an edge, the total number of facets N_f changes and then even the size of the Hessian of V becomes discontinuous. These are degenerate cases that are ignored in this work, as the use of the gradient and Hessian of V w.r.t to its supports is used to tune numerical optimization algorithms, which will likely skip or circumnavigate these singularities with ease, provided they are stable enough.

4. Shape Determination

itself, as well as the area of the adjacent facets. Therefore, from Corollary 1 it is obvious that $d^2V/d\mathbf{h}_l d\mathbf{h}_p$ must be 0 when facets l and p are not adjacent.

□

Having the vertex-facet adjacency information obtained by the convex-hull-dual process during the obtention of $\mathcal{P}(\mathbf{h})$, one can easily extract the list of edges, as well as which two facets touch each edge. One needs only to iterate through the vertices of each facet and log each edge as each adjacent vertex-pair; at the same time, one registers the facet each edge belongs to. At the end, each edge of the list shall have two facets associated to it. With this information, one can iterate all edges and compute the appropriate contributions to $\partial V/\partial \mathbf{h}_i \partial \mathbf{h}_j$ from Equation 4.29.

4.4. Observability w.r.t. the EGI Size

Friedman et al. (2019) studied the observability of this inversion problem based on the Gramian of $\hat{\mathcal{S}}$, *i.e.* $\mathcal{G} = \hat{\mathcal{S}}^T \hat{\mathcal{S}}$. The unconstrained (negative values allowed) solution to the least squares problem stated by Equation 4.4 has a closed form given by

$$\hat{\mathbf{a}} = \mathcal{G}^{-1} \hat{\mathcal{S}}^T \mathcal{L}. \quad (4.40)$$

Thus, they define the observability of the problem by requiring that \mathcal{G} is invertible—*i.e.* that it has full rank.

However, as explained above, this problem is rarely solved using Equation 4.40, because of the positivity restriction imposed on \mathbf{a} (see Section 4.1.1). Therefore, the Gramian \mathcal{G} is only a qualitative measure of observability. When either the exponential (Kaasalainen et al., 2001) or the constrained (NNLS) approach is used to force positivity, the solution to the problem has no closed form, and must be resolved using some iterative method—*i.e.* it is not guaranteed that the properties of \mathcal{G} define the observability of \mathbf{a} anymore.

In particular, the NNLS approach of Lawson et al. (1995) uses the so called active-set method. The active-set method divides the dimensions of the search space into an *active* subspace, composed by all the dimensions where the positivity constraint is active—*i.e.* the dependent variable is 0—, and the *passive* subspace, with the rest of dimensions where the constraint is not imposed—*i.e.* the dependent variable is positive. It does so by first assuming that all the dimensions are in the *active* set; then it iteratively selects the best candidate dimension to change from *active* to *passive* and solves the unconstrained least squares problem on the *passive* set; if positivity of the solution is maintained, the candidate *passive* dimension is accepted. Kim (2011) proved that the sub-matrix of $\hat{\mathcal{S}}$ projected onto the *passive* set remains full-rank until convergence is reached. In other words, a unique solution for \mathbf{a} exists despite \mathcal{G} being rank-deficient w.r.t. the number of directions in the EGI.

This implies that, in the noise-free ideal case, \mathcal{G} does not need to be full rank, but just have its rank high enough to capture independent information from each of the actual flat facets that compose the shape of the object being observed. These, of course, are unknown beforehand, hence the need for a shape-inversion method. However, especially in the case of objects with sharp edges and mostly flat facets, one shall expect that most directions in the EGI have zero

albedo-area, while only a few have significant values. Therefore, for most artificial objects, indeed a considerably reduced number of directions should suffice to describe a recognizable shape. In the noise-free scenario, then, it is interesting to tessellate the unit sphere with as many cells as computationally feasible, because this will ensure that non-zero albedo-area values are located where the normal is closest to that of the corresponding real facet of the observed object.

When noise is present, however, the rank-deficient \mathcal{G} case can still work, but increasing the dependent dimensions too much can lead to instability due to high sensitivity to noise (Calef et al., 2006). In these cases, the noise stabilization methods mentioned above (Calef et al., 2006; Fan et al., 2020; Fan et al., 2021) should be used.

4.5. Simulation Scenarios

Within the scope of this thesis, the shape inversion from light curves using the EGI plus Minkowski minimization method has been tested with different shapes, noise scenarios and ranks of the problem Gramian.

4.5.1. Test Shapes

The shapes being tested are:

- a cube with 4 m side—it represents one of the simplest shapes achievable;
- an icosahedron with 1 m edges—it represents a slightly more complex, approximately spherical bus;
- a cylinder with 10 m length and 1 m radius—analogue to a simplified rocket body; and
- a winged cuboid, whose base is a $4\text{ m} \times 6\text{ m} \times 12.5\text{ m}$ cuboid and whose wings on either side are $6\text{ m} \times 9\text{ m}$ rectangles without thickness, separated 1 m from the base—it emulates a non-convex object with solar panels.

In this work, the matrix $\hat{\mathcal{S}}$ used assumes Lambert’s diffuse law of reflection. This is, the CT BRDF model from Section 3.1.4 with $\rho = 1$ and $\mathfrak{d} = 1$.

Figure 4.3 illustrates these shapes in orthographic projection, conserving the aspect ratio of each shape. Figure 4.4 shows an example of the corresponding $d = 5$ EGI for each shape, both in orthographic and Hammer projections. The color code represents the area of each cell, as indicated by the color bars. Cell borders are shown as white lines. Additionally, the latitude/longitude grid is shown in red on the Hammer projections.

Note that for the specific case of the winged cuboid, which is a non-convex shape, the EGI is generated by assuming that each facet is not obscured along its normal by any other facet. Because the thickness of the wings is null in the simulation, this assumption does not introduce any error. However, the EGI to polyhedron algorithm presented above recovers a convex shape instead (see the results section).

4. Shape Determination

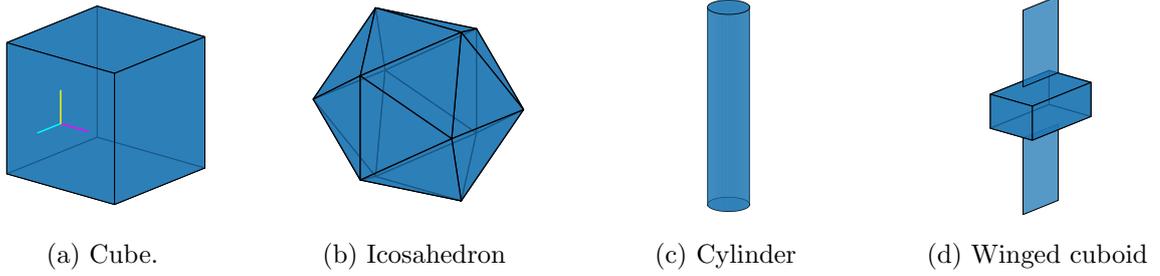


Figure 4.3.: Test shapes, body-frame axis represented in the Cube (cyan is x , magenta is y and yellow is z) (Vallverdú Cabrera et al., 2021b)

4.5.2. Scenario Geometry

The simulated light curve is measured by a telescope in the location of ART (see Appendix A.2) between 22:00 of 2019-01-01 and 3:45 of the following day, in Coordinated Universal Time (UTC). The simulated object always has the same circular Mid Earth Orbit (MEO) with 26 731 km semi-major axis, 40° inclination and 10.94° right ascension of the ascending node. At midnight UTC of the observation night, it has a true anomaly of 67.5° . All test cases exhibit a nutation, the motion typical of a body with axial symmetry⁶ whose angular momentum is constant and parallel to the inertial z -axis, with an overall magnitude of the angular rate of 3.44 min^{-1} , nutating with the body z -axis at 70.5° from the angular momentum, at a rate of 2.7 min^{-1} . Figure 4.5 shows the simulated light curve for each test shape: in orange, the light curve without noise; in blue, the light curve randomized with $\text{SNR} = 10$ (see subsection about measurement noise below). Circles are data points, and the lines are guides to the eye.

4.5.3. Light Curve Measurement

The light curves are modelled using the model from Chapter 3. For this particular experiment, source and sensor noise is added from a Gaussian distribution, based on the median signal level represented by a Signal to Noise Ratio (SNR). This is, the measured light curve $\mathcal{L} = [\mathcal{L}_1 \cdots \mathcal{L}_{N_t}]^T$ is obtained by polluting each simulated value with additive noise:

$$\mathcal{L}_i = \mathcal{L}_i^s + \mathcal{N} \left(0, \left(\frac{\bar{\mathcal{L}}}{\text{SNR}} \right)^2 \right), \quad (4.41)$$

where \mathcal{L}^s is a noise-less simulated value, and $\bar{\mathcal{L}}$ indicates the median of \mathcal{L}^s . Figure 4.5 shows an example noisy light curve with $\text{SNR} = 10$ (blue) along its no-noise counterpart (orange), for each test shape.

⁶Although some of the shapes introduced above do not present axial symmetry, during the simulations the same attitude profile has been forced to all of them, to remove the attitude as an added factor of comparison.

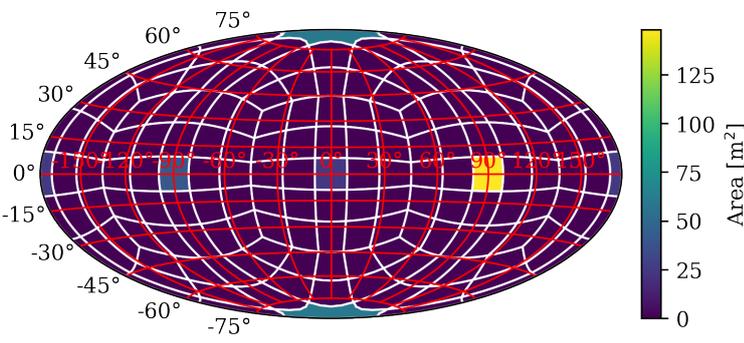
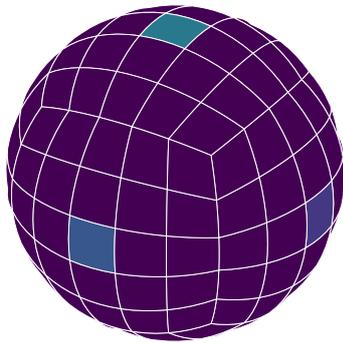
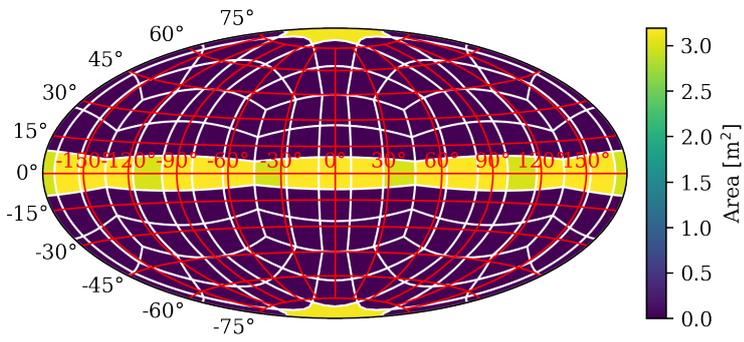
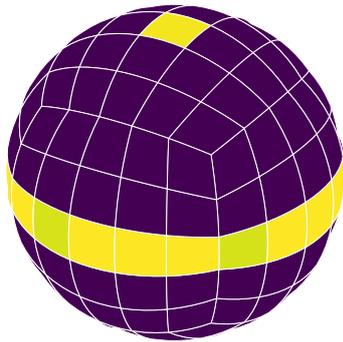
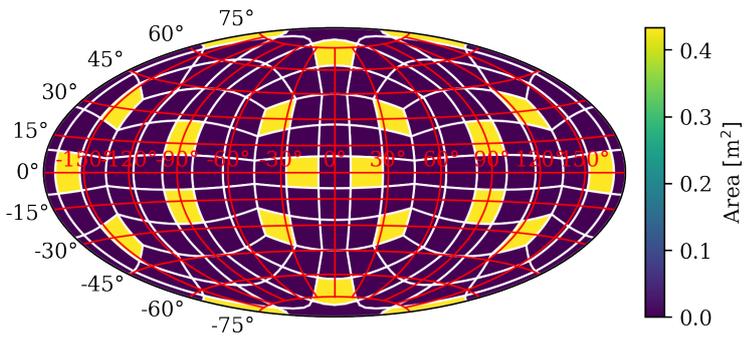
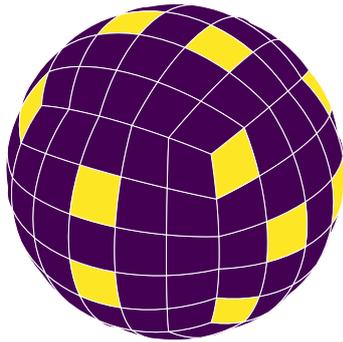
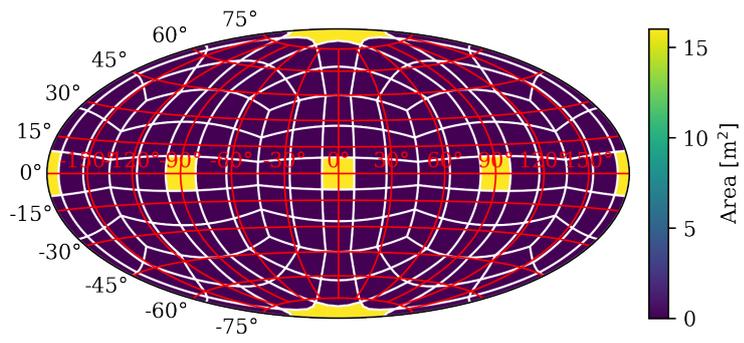
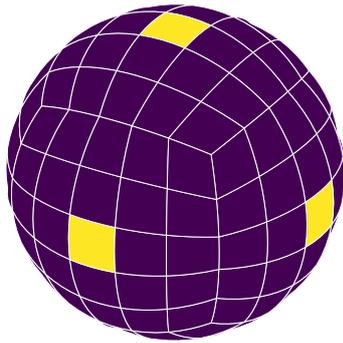
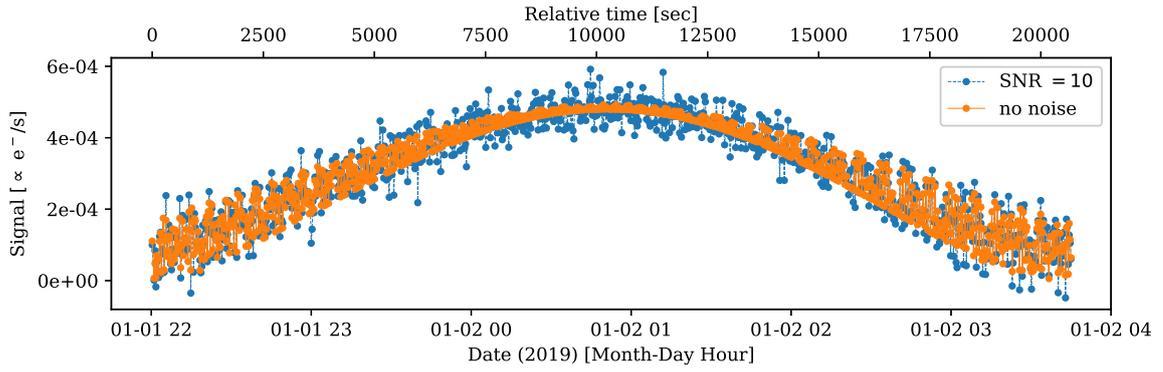
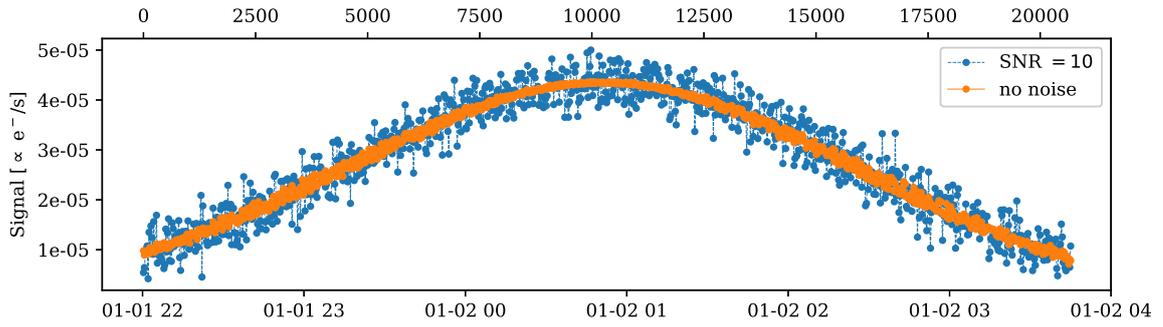


Figure 4.4.: EGIs of the test shapes, based on a $d = 5$ quad-cube tessellation (Vallverdú Cabrera et al., 2021b)

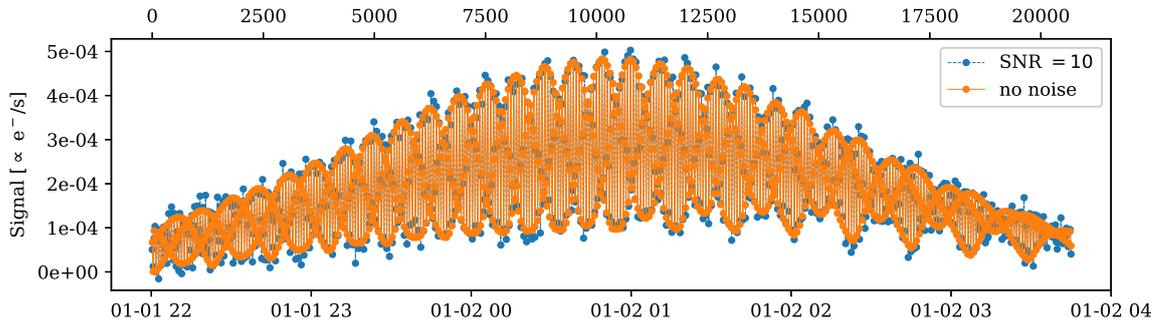
4. Shape Determination



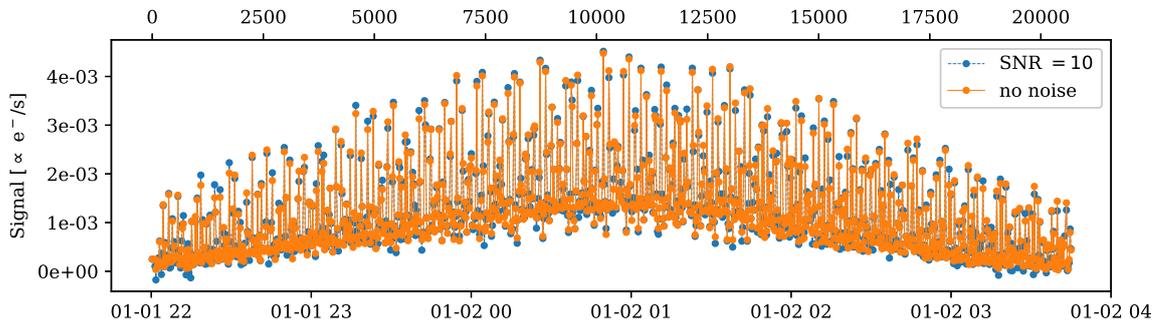
(a) Cube



(b) Icosahedron



(c) Cylinder



(d) Winged cuboid

Figure 4.5.: Simulated light curves (Vallverdú Cabrera et al., 2021b)

4.5.4. Observability Conditions

This thesis explores the shape recovery performance at different light curve sampling intervals (which affects the rank of the observability Gramian) and different EGI quad-cube tessellation divisions. This is quantified by the d parameter defined in Appendix E.2; the number of cells the resulting EGI is $N_a = 6d^2$. For the results in this work, the light curve has been sampled at $N_t \in \{150, 1000\}$ points, while the EGIs have been divided with $d \in \{3, 5, 7, 9\}$. The light curves from Figure 4.5 are sampled at $N_t = 1000$.

4.5.5. Noise Mitigation

Following the technique used in Fan et al. (2020) and companion papers, in this simulation the directions of the EGI with an area lower than $\epsilon \max(\hat{a})$ are ignored. This adds stability to the Minkowski minimization and filters out small area values that arise most likely from noise only. Then only those directions that survive this criterion are constrained for convexity with Equation 4.6. This study uses the value of $\epsilon = 5\%$.

Furthermore, the three regularization schemes (none, Tikhonov and Lasso) are compared in the scenarios of the cube and the icosahedron.

4.5.6. Hessian of the Volume

The results shown are obtained by using the (numerical) Broyden–Fletcher–Goldfarb–Shanno (BFGS) approach (Nocedal et al., 2006) to compute the Hessian of the volume constraint. In some cases, noted appropriately below, this has been compared against using the analytical expression from Equation 4.29.

4.6. Results and Discussion

This section presents and discusses the results of the shape recovery algorithm applied to the simulation scenarios described in the previous section. First, Table 4.1 shows the rank of the Gramian for each of the tested scenarios, which mainly depends on the number of cells of the EGI ($N_a = 6d^2$) and the number of samples in the input light curve (N_t). For all the cases tested, the numeric estimation of the rank always agrees with

$$\text{rank}(\hat{\mathcal{S}}^T \hat{\mathcal{S}}) = \min(N_a, N_t), \quad (4.42)$$

which means that the corresponding light curves have no redundant samples. In Table 4.1, rank-deficient scenarios—*i.e.* $N_a > \text{rank}(\hat{\mathcal{S}}^T \hat{\mathcal{S}})$ —are highlighted in **red**.

4.6.1. Cube

Starting with the cube-shaped object, Figure 4.6 shows the normalized homotheticity distance $\bar{H}(\mathcal{P}, \hat{\mathcal{P}})$ as in Equation 4.17 versus SNR, along different scenarios. \mathcal{P} is the original polyhedron (the cube, in this case), and $\hat{\mathcal{P}}$ is the estimated one. Each plot corresponds to a light curve sampling density (N_t). Each line corresponds to a specific EGI sampling d , color coded. Each (SNR, \bar{H}) point (\times) on the plots is the result of averaging

4. Shape Determination

Table 4.1.: Rank of the Gramian ($\hat{\mathcal{S}}^T \hat{\mathcal{S}}$) at different light curve and EGI sampling densities

d	N_a	rank($\hat{\mathcal{S}}^T \hat{\mathcal{S}}$)	
		$N_t = 150$	$N_t = 1000$
3	54	54	54
5	150	150	150
7	294	150	294
9	486	150	1000

the individual performance over $N_{\text{MC}} = 64$ runs within a Monte Carlo (MC) simulation, where the noise is randomly sampled according to Equation 4.41 at each run. This is

$$\hat{\mu} = \frac{1}{N_{\text{MC}}} \sum_{i=1}^{N_{\text{MC}}} \bar{H}(\mathcal{P}, \hat{\mathcal{P}}_i). \quad (4.43)$$

The vertical bars delimited by opposite triangles indicate the $\sim 95\%$ confidence interval of $\hat{\mu}$, computed from the unbiased estimation of the standard deviation:

$$s^2 = \frac{1}{N_{\text{MC}} - 1} \sum_{i=1}^{N_{\text{MC}}} \left(\bar{H}(\mathcal{P}, \hat{\mathcal{P}}_i) - \hat{\mu} \right)^2. \quad (4.44)$$

To estimate these error bars, it is assumed that $\hat{\mu}$ follows a Normal distribution. Thus, they span over the range

$$-1.96 \frac{s}{N_{\text{MC}}} < \bar{H} - \hat{\mu} < 1.96 \frac{s}{N_{\text{MC}}}. \quad (4.45)$$

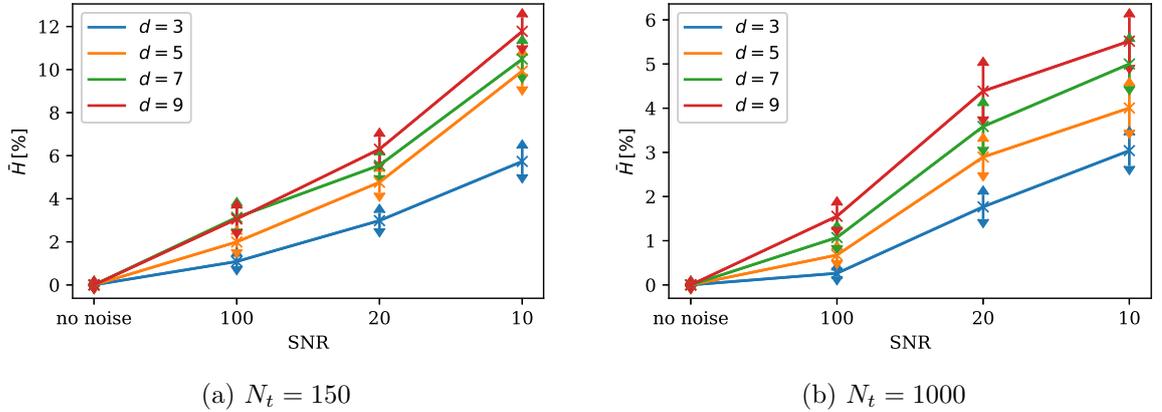


Figure 4.6.: Cube results, averaged over 64 Monte Carlo runs (Vallverdú Cabrera et al., 2021b)

For the *no noise* case, the shape can be recovered with maximum fidelity ($\bar{H} = 0$), even when the Gramian is rank-deficient. This stems from two main factors. First, both light curves ($N_t = \{150, 1000\}$) have enough samples, and come from a wide enough range of attitude states, so that the six facets of the cube are observable. Second, all the quad-cube tessellation sizes have one cell aligned to each of the six facets of the cube. Thus, all the

important information goes there, and the rest of the cells are virtually zero—they do not capture any significant amount of area-albedo that could *e.g.* affect the shape recovery differently for different values of N_t or d .

Moving over to the noisy scenarios shown on Figure 4.6, the first noticeable point is that fidelity decreases with higher SNR values, as expected. Unsurprisingly, the scenario with the denser light curve produces better fidelity—more samples implies more information on the object. Regarding the effect of d (or N_a), Figure 4.6 shows that, for the cube, a denser EGI always leads to higher error. Moreover, this trend does not seem to be highly influenced by whether the scenario is rank-deficient or not. This may again be explained by the affinity of the cube to the quad-cube tessellation: because the $d = 1$ EGI is enough to perfectly capture its six facets, more cells than these are just recipients to be populated by noise, each of which contributes further to fidelity error on the recovered shape.

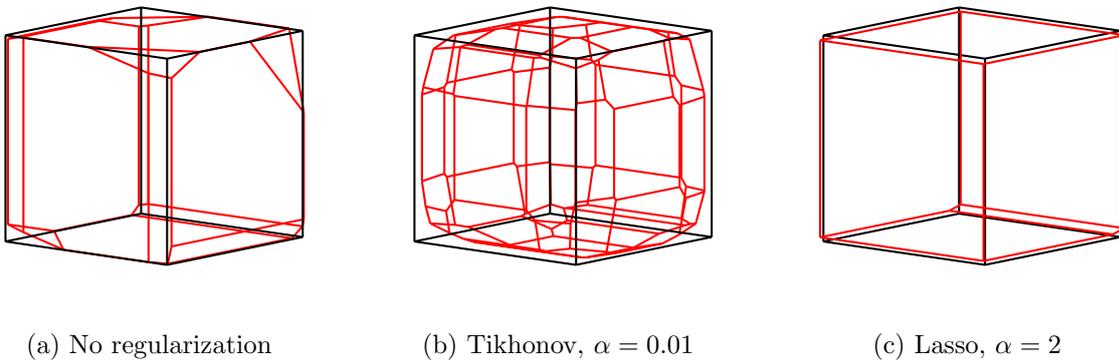


Figure 4.7.: Original (black) versus recovered (red) shape for the cube (Vallverdú Cabrera et al., 2021b)

The cube, being the simplest shape on the test bench, has been chosen to qualitatively assess the effects that regularization can have on the recovered shape. While Tikhonov (Equation 4.9) regularization favors many smaller values in the estimated parameter vector, Lasso (Equation 4.10) regularization encourages a sparser parameter vector, with a few big values and the rest almost null. Thus, when recovering an EGI from a light curve,

- shapes estimated using Tikhonov regularization tend to be more rounded, compared to not using regularization, because the area from big area cells *flows* into neighboring ones. In contrast,
- when using Lasso regularization, shapes have fewer facets and sharper edges, because neighboring cells with similar area values tend to merge into a single cell with bigger area.

This expected behaviour can be observed in Figure 4.7, where the original (black-edged frame) and recovered (red-edged frame) for three solutions with different regularization schemes are shown, in the scenario of $N_t = 1000$, $d = 7$ and $\text{SNR} = 10$, for one random Monte Carlo run. The values of α have been manually tuned to exemplify the features of each regularization scheme. However, as suggested in Calef et al. (2006), a more rigorous approach would be to use a cross-validation scheme to choose the ideal value of α .

4. Shape Determination

Qualitatively, Figure 4.7 shows that the non-regularized shape has some apparently random deviations from the cube; the Tikhonov solution smoothens them out by rounding the edges, while the Lasso one removes them by adding them to the biggest facets. In this case, because the cube is a clear example of a shape with a few big facets, Lasso clearly outperforms the other two approaches.

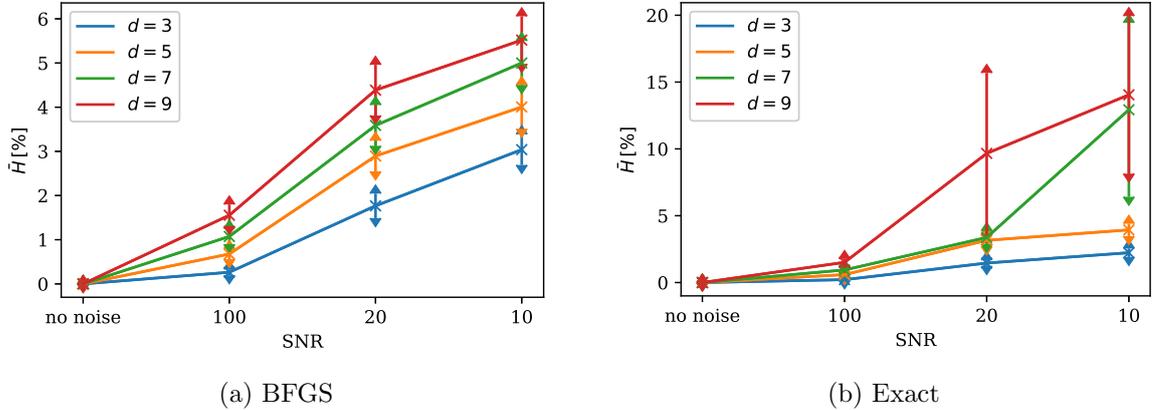


Figure 4.8.: Cube results, 64 Monte Carlo runs, $N_t = 1000$, different Hessian strategies

Finally, the analytical expression of the volume Hessian (Equation 4.29) has been tested only against the cube shape, too. Figure 4.8 shows a comparison between both strategies (BFGS and analytical/exact). Note that Figure 4.8a is exactly the same as Figure 4.6b above—it has been included again to make visual comparison easier. By comparing it against Figure 4.8b, it can be seen that both approaches deliver similar results, except for the low SNR, high tessellation density cases. In particular, SNR = 20 and $d = 9$, and SNR = 10 and $d \in \{7, 9\}$. For these cases, some of the Monte Carlo runs have failed to converge, which resulted in greater average homothetic distances of the recovered shape. This made the standard deviation of these tests widen significantly, too.

This indicates that the analytical expression from Equation 4.29 may not cover all the corner cases. In particular, there is the issue that Equation 4.29 assumes the small variations of the supports do not change the adjacency information of the polyhedron. However, the internal process to recover the polyhedron from the supports, via the dual transform, could mean that changing the value of one particular support changes the adjacency information of the polyhedron (*e.g.* a new facet appears, or two facets become adjacent, or cease to be so). Since the BFGS approach is numerical, and treats the whole volume computation from the polyhedron supports as a black box, it should capture these corner cases without problems, within numerical accuracy. Intuitively, it is likelier that these corner cases appear for higher tessellation densities—*i.e.* more supports to account for means more potential failure points. Therefore, the fact that the failure happens especially with the higher noise, higher density cases, constitutes circumstantial evidence to support this hypothesis.

In terms of computational speed, both methods took approximately the same time, meaning that the analytical and BFGS Hessian calculations cost approximately the same.

4.6.2. Icosahedron

The next shape on the test bed is the icosahedron, whose results are shown in Figure 4.9 (see Section 4.6.1 for a detailed explanation of the plots).

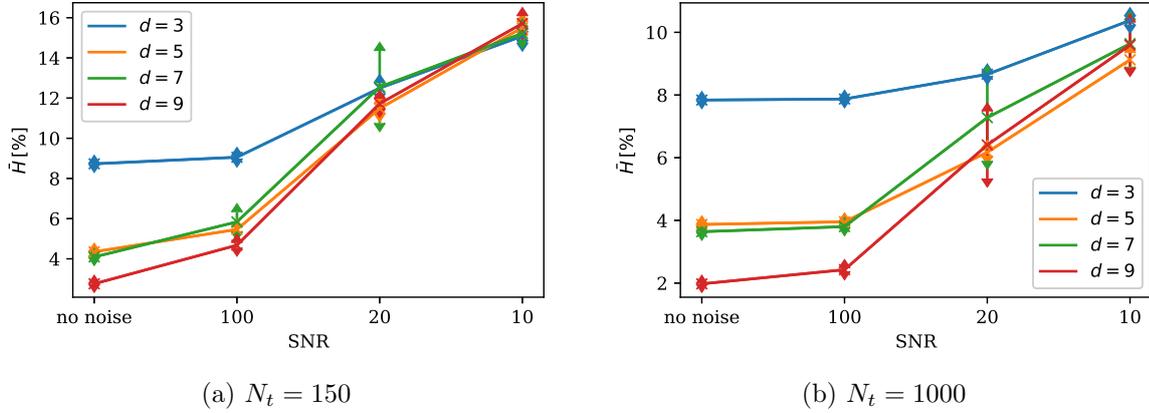


Figure 4.9.: Icosahedron results, averaged over 64 Monte Carlo runs (Vallverdú Cabrera et al., 2021b)

As happened with the cube, with the icosahedron fidelity error increases with higher noise levels (SNR) and sparser light curves (N_t). However, it differs from the cube in other aspects. First, with the icosahedron case one does not recover an exact replica of the original shape (*i.e.* $\bar{H} > 0$) even in the *no noise* scenarios. This is because the normals of the facets of the icosahedron are not aligned with any particular direction of the quad-cube tessellation. Thus, it can be seen that, for the *no noise* scenarios, the fidelity improves with a denser EGI, *even when the Gramian is rank-deficient*. This happens because, for a denser EGI, the closest direction to each normal of the icosahedron will be closer to it. Thus, the denser EGIs approximate the icosahedron facets better. For this same reason, at low noise levels (SNR = 100) denser EGIs still perform better. However, for the noisiest cases (SNR = 10), noise dominates and the EGI density has no significant impact anymore.

The Icosahedron plot presents some points where the error bars are unusually wide. These are $d = 7$, SNR = 20 on the $N_t = 150$ light curve, and $d = \{7, 9\}$, SNR = 20 for $N_t = 1000$. This may indicate some instability in the shape recovery process that produces higher variance homotheticity distance measurements, which should be investigated further with a wider range of scenarios—*e.g.* varying noise threshold (ϵ) values.

Regarding regularization, the icosahedron has also been tested qualitatively, for the same scenario as the cube (see previous section). Figure 4.10 shows the corresponding results. Opposite to the cube, the icosahedron has more facets with less area. As a consequence, the non-regularized and the Tikhonov schemes barely differ: the non-regularized solution is already quite rounded, so ridge regularization has little effect on it. On the contrary, the Lasso regularization with the same high value of α as in the cube case completely distorts the recovered shape away from the icosahedron—lower values of α for the Lasso scheme do not produce important changes w.r.t. the non-regularized case. Thus, the icosahedron is a good counter example to the regularization results of the cube.

4. Shape Determination

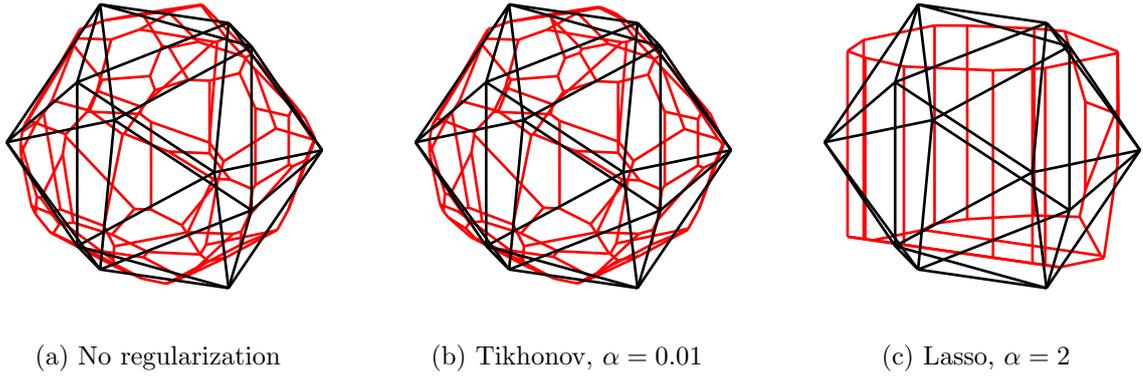


Figure 4.10.: Original (black) versus recovered (red) shape for the icosahedron (Vallverdú Cabrera et al., 2021b)

4.6.3. Cylinder

The last convex shape tested here, the cylinder, has its results in Figure 4.11 (see Section 4.6.1 for a detailed explanation of the plots).

In contrast to the two previous shapes, it has proven to be much more difficult to recover: $\bar{H}(\mathcal{P}, \hat{\mathcal{P}})$ is greater overall. This is probably due to the cylinder having a smooth curvature, which contradicts (or at least strains) the hypothesis of a shape represented by a polyhedron built from several flat facets—the side of the cylinder is the limit of infinite facets with infinitesimal area. Beyond this point, the effects of N_t and SNR share the same trends as with the cube and icosahedron cases: higher N_t and higher SNR improve fidelity, as is to be expected.

For the *no noise* cases, there is no clear trend whether denser EGIs perform better or not, although for $N_t = 150$ it seems that rank-deficient cases perform slightly worse. Instead, moving to the noisy scenarios, as noise level increases, denser EGIs clearly deteriorate fidelity, regardless of the rank of the Gramian—as with the cube. In fact, the two circles of the cylinder do match the polar facets of a quad-cube tessellation with an odd value of d , while its sides all map more or less well to the cells in the equator of the quad-cube. Thus, although the match is not as close as with the cube, it is plausible that the cylinder maps onto a low d quad-cube well enough so that the potential benefits of higher density are overrun by a wider spread of noise. This is, the cylinder is better approximated by a cleaner prism whose base is a polygon of less vertices, than one whose base has more vertices but is rougher.

Figure 4.13 shows the original (black-edged frame) and recovered (red-edged frame) shapes for the cylinder case, for $N_t = 1000$, $d = 7$ and $\text{SNR} = 10$, an arbitrary Monte Carlo run. A common feature with most of its iterations, is that the recovered shape is shorter along the cylinder axis, but wider orthogonality. This agrees with the conjecture that information is lost due to the real cylinder having infinite flat facets on its sides that are infinitesimally small, hence the axial shortness. Nonetheless, the recovered shape clearly resembles a cylinder.

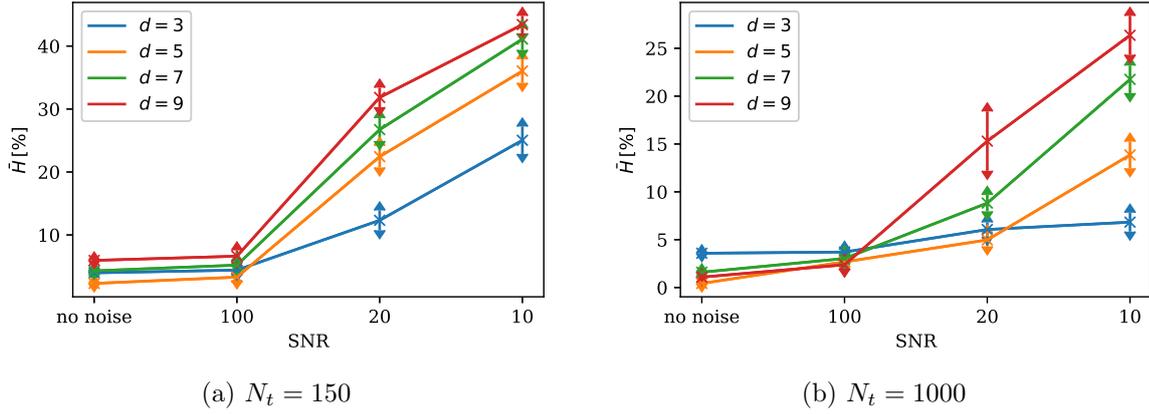


Figure 4.11.: Cylinder results, averaged over 64 Monte Carlo runs (Vallverdú Cabrera et al., 2021b)

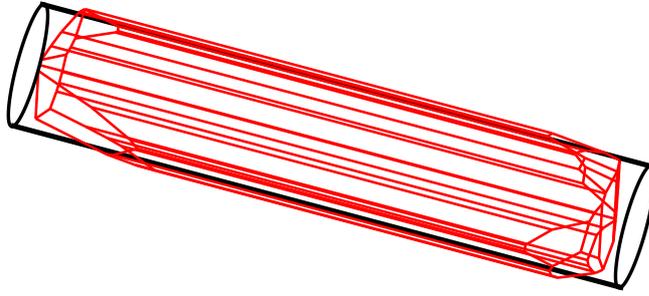


Figure 4.12.: Original (black) versus recovered (red) shape for the cylinder.

4.6.4. Winged Cuboid

The winged cuboid is a non-convex object. As such, it cannot be uniquely represented by an EGI, as per Minkowski's theorem (Minkowski, 1897). Nonetheless, it is still possible to recover an EGI from its light curve. In fact, in the particular case of the winged cuboid, for the $N_t = 1000$, $d = 7$ and $\text{SNR} \rightarrow \infty$ scenario, the area residual is $4.29 \cdot 10^{-3}$, low enough to pass as a convex object. In other words, Minkowski minimization can be applied to recover an *equivalent* convex shape—*i.e.* the shape that would produce a light curve as similar as possible to the measured one (Kaasalainen et al., 2001). This recovered shape may be significantly different than the original one, however. Figure 4.13 shows the original (black-edged frame) and recovered (red-edged frame) shapes for the winged case, for $N_t = 1000$, $d = 7$ and $\text{SNR} = 10$.

Minkowski minimization only renders convex shapes, hence the recovered shape approximates a cuboid that significantly differs from the original winged cuboid. Its only similitude is that the ratio of cross section areas in any two different observation directions is approximately equal when calculated for the original or for the recovered shape—this is a necessary condition if the recovered shape has to reproduce the measured light curve as faithfully as possible. Despite this, the two shapes are clearly far from homothetic. For this

4. Shape Determination

reason, a wide range of tests to assess the impact N_t , d and SNR on the homotheticity distance of original and recovered polyhedron has not been performed within the scope of this thesis.

Nonetheless, this experiment shows that *equivalent* convex shapes can be recovered even for non-convex objects. These could be used to help identify, classify or simulate the light curves of these objects, even if the *equivalent* convex shape is far from the original shape of the object.

Another important conclusion from this experiment is that the recovered shape is highly dissimilar from the convex hull of the original shape, even if scaled, which differs from the results presented in Kaasalainen et al. (2001), where recovered shapes of asteroid-like non-convex bodies were found to be close to the corresponding convex hulls. This is an extreme example of the conclusion therein, which states that the greater the concavities in the original object, the farther the *equivalent* convex shape is from the convex hull.

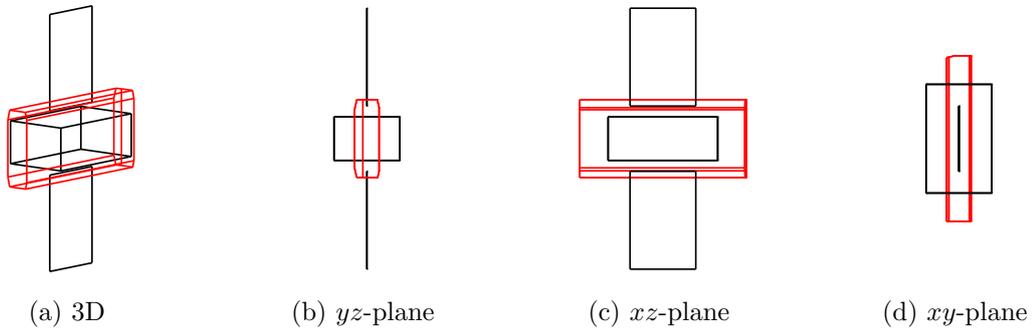


Figure 4.13.: Original (black) versus recovered (red) shape for the winged cuboid (Vallverdú Cabrera et al., 2021b)

4.6.5. General Discussion

This subsection closes the results part of the chapter by summarizing and merging the conclusions extracted from all the tested scenarios together.

Overall, regarding convex shapes, via Monte Carlo simulations of 64 runs each and evaluating the homotheticity distance, it has been shown that

- the sampling density of a light curve,
- the density of the EGI and
- the noise level present in the light curve

have a strong impact on the shape inversion results. As expected, more samples and less noise means better results. However, depending on the studied shape, denser or sparser EGIs can perform better. It has been conjectured that the main reason is the affinity/dissimilitude of the real shape and the quad-cube tessellation scheme. Further studies should investigate whether different tessellation schemes, such as the geodesic dome proposed in Horn (1984) and Calef et al. (2006), lead to significantly different performance curves—*e.g.* the icosahedron would have more affinity with an icosahedron-based geodesic

dome. The rank-deficiency/fullness of the Gramian of the problem has been found to have a marginal impact, significantly lower than the other aforementioned parameters: similar trends are observed with $N_t = 150$, where half of the cases are rank-deficient, and with $N_t = 1000$, where all cases are full-rank.

Still on convex shapes, regarding regularization, it has been shown qualitatively that Lasso stabilizes noisy scenarios for shapes with few, big facets, while Thikhonov is better suited for naturally rounder shapes. However, care must be used when applying regularization. Both Lasso and Thikhonov smoothen the parameter vector by adding some bias (sharper/rounder features). For objects with *a priori* unknown shapes, it is not possible to know if it is better to sharpen or to round the noisy recovered shape. Thus, it is not known if the added bias corrects or further deforms the result. All in all, regularization should only be used if some *a priori* information on the general shape of the object is known beforehand—*e.g.* it is likely that asteroids are rounder, hence Thikhonov may improve smoothen sharp features that originate from noise.

Finally, the analytical formula to compute the volume Hessian from the polyhedron supports, which has been derived in this chapter, has been shown to fail for the most demanding cases. It has been hypothesised that changes in the adjacency information triggered by support variations could be the cause.

4.7. Potential Extensions

This section discusses possible extensions to recover additional information, or to improve the results of the shape inversion method based on EGI estimation plus Minkowski minimization.

4.7.1. Albedo Variegation

For the case of asteroids, Kaasalainen et al. (2001) propose a regularization scheme in which they minimize the albedo and area as separate terms so that, when multiplied element-wise, they are equal to the obtained $\hat{\mathbf{a}}$. To constrain the infinite-solutions problem, they add two regularization components, one for the convexity condition, and one that describes how albedo should vary across the surface of the object.

In the case of asteroids, although albedo distributions are typically unknown, they in general can be assumed to be smooth enough across different regions of the asteroid. Therefore, it is plausible to use some constraint that penalizes fast albedo variations in any direction along the surface (Kaasalainen et al., 2001). Man-made objects, however, cannot be expected to have constant or smooth albedo distributions; on the contrary, one should expect artificial satellites to have external components with radically different optical properties (*e.g.* solar panel vs. antenna dish vs. multi-layer insulation). This means that, unless *a priori* information on the albedo distribution of the particular object is known, using single-channel⁷ light curves the albedo-area separation for man-made objects is not

⁷By single channel, it is meant light curves measured with a single sensor, at a specific single (color) band. In contrast, multiple channel light curves are simultaneous light curves of the same object, from the same location, observed through different bands.

4. Shape Determination

possible.

Research on this line should follow the work of Dianetti et al. (2019) and Dianetti (2020), which explores using multi-channel and/or polarized light curves to estimate the albedo of the surface of RSOs.

4.7.2. Non-Convexity

As shown above, the estimated *equivalent* convex shape of a non-convex object may share some characteristics with the original one, but has a great homotheticity distance and is, of course, convex. Kaasalainen et al. (2001) managed to qualitatively characterize concavities in asteroid-like non-convex objects, not being able to fully recover the depth of the valleys, but only to infer their presence. They did so by modifying the vertices of the *equivalent* convex shape by a set of coefficients, which then they used to optimize, in the least squares sense, the difference between the measured light curve and the one obtained by simulating it with the candidate non-convex shape.

However, Viikinkoski et al. (2017) argue that photometric data alone is not enough to determine a non-convex shape uniquely: one big concavity can be replaced by a set of smaller concavities that produce the same light curve signature. Thus, non-convex shape recovery must involve other types of measurements, which constrain the possible shape enough to overcome the non-uniqueness of the problem: in Kaasalainen (2016), it is shown that low resolution images that provide a rough contour of cross-sections of the object can help toward this goal.

4.8. Conclusion

This chapter has presented the EGI plus Minkowski minimization algorithm to recover convex shapes from light curves, when the attitude is known. The method has been tested against several simulated scenarios.

Focusing on its scientific contribution, first a new metric has been introduced, which can be used to assess how close two polyhedrons are to each other in the homothetic sense. This metric has been used to assess the accuracy of the EGI plus Minkowski minimization shape determination.

Second, the expression of the Hessian of the volume of a polyhedron relative to its supports has been derived. It has been compared against a common numerical alternative: it has been found to fail for complex cases, probably because it does not account for topological changes triggered by support variations.

Third, it has been shown, through analysis of the active-set method used in NNLS, that EGIs with a higher number of tessellation cells than the number of samples in the light curve can be used, without hindering the observability conditions. This conclusion has been ratified with simulation results, and put to the test against noisy cases, too. In the particular case of the icosahedron, whose facet normals do not align with the quad-cube tessellation scheme naturally, the scenarios with more number of EGI cells outperform the lower-density ones for high SNR, even when the albedo-areas are not observable according to the rank of the

Gramian. Meanwhile, performance remained similar across all densities at low SNR, where noise dominates.

The next step would be to obtain good quality light curves of operated satellites whose telemetry is known, so that attitude is an input, and test the shape inversion on them. This would stress the method further, although problems such as albedo variegation and non-convexity would still have to be addressed.

5. Attitude Determination

This chapter researches a new avenue for sequential Attitude Determination (AD) from light curves. From Section 2.2.2 of the Literature Review Chapter, it follows that both the dynamic and measurement models involved in estimating attitude from light are highly non-linear functions of the state, to varying degrees. Moreover, the attitude state may not be fully observable in some scenarios (Dianetti, 2020). This leads to the attitude state often following highly non-Gaussian and potentially multimodal Probability Density Functions (PDFs). When additional states are estimated simultaneously, this condition aggravates further (Linares et al., 2018; Yun et al., 2020). The proliferation of advanced non-linear filtering techniques to approach the light-curve-based AD, summarized in Section 2.2.2, stems from this additional complexity. The goal of these methods is to 1) outperform the classical Extended Kalman Filter (EKF)/Unscented Kalman Filter (UKF)¹, while 2) reducing the computational burden of the Particle Filter (PF).

This chapter proposes to use an alternative non-linear filter that has never been applied before to the problem of AD from light curves: the Adaptive Gaussian Mixtures Unscented Kalman Filter (AGMUKF). Like some approaches presented in Section 2.2.2, the AGMUKF approximates the PDF of the state as a Gaussian Mixture (GM). Each kernel is then treated as a UKF, like in Yun et al. (2020). However, the AGMUKF differs in that it incorporates a non-linearity assessment metric, based on the Unscented Transform (UT). This metric is used to detect the rise of non-linearities and split the kernels accordingly. These may eventually be merged back together as soon as they become too close.

On the one hand, a fine enough GM can represent any arbitrarily complex PDF in L1 sense, something unachievable with classical EKF/UKF. On the other hand, the adaptive character of this approach strives to reduce the computational burden typical of the particle or GM filters—it naturally avoids the degeneration problem of the PF, as well as the need for regularization or other such techniques.

The Gaussian splitting libraries developed by Vittaldev et al. (2016) enable the design of this filter, which Schiemenz et al. (2020) specialized to perform OD extended with air density realism. This dissertation adapts these two works to the AD case. Furthermore, refer to Wu et al. (2005) for a generic definition of the UKF. Analogously, Elfring et al. (2021) is a hands-on tutorial on the PFs.

Within the scope of this work, the AGMUKF for light-curve-based AD is implemented assuming that all other factors are known: shape, size, reflective properties, observation geometry (orbit) and telescope performance parameters.

The rest of this chapter is divided as follows. First, the following Section 5.1 outlines the non-linear sequential estimation problem, and presents the AD-specific dynamic and

¹typical on *e.g.* the OD problem (Vallado et al., 2013)

measurement models. Then, Section 5.2 analyses the non-linearity of these models, which motivates and justifies the AGMUKF proposed in Section 5.3. The AGMUKF is validated under simulated scenarios in Section 5.4, after which the chapter closes with the chapter conclusions in Section 5.5.

This chapter is adapted from and expands on Vallverdú Cabrera et al. (2023), the third and last paper published as part of the doctoral research of this thesis. As such, parts of Vallverdú Cabrera et al. (2023) are included herein.

5.1. Problem Definition

This section poses the problem of RSO attitude estimation using light curves. Thus, it defines the conventions, hypotheses, variables, parameters and models involved. In particular, Section 5.1.1 describes the generic non-linear sequential estimation problem. It is followed by the detailed description of the attitude dynamic model (Section 5.1.2) and the measurement model (Section 5.1.3).

5.1.1. Non-Linear Sequential Estimation

This subsection defines the non-linear sequential estimation problem based on Bayes theorem, as described in *e.g.* Elfring et al. (2021).

The generic non-linear sequential estimation problem has the goal to estimate the Probability Density Function (PDF) of some hidden random variable $X(t) \in \mathbb{R}^n$ that varies with time and follows a continuous Markov chain, conditional to a sequence of observables $Z_k \sim \text{pdf}(z_k)$, which have been measured as z_k at times $t_k, \forall k \in 1..N_t$ such that $k > l \Rightarrow t_k > t_l$. This is

$$X(t_{N_t}) \sim \text{pdf}(x(t_{N_t}) | z_1, \dots, z_{N_t}). \quad (5.1)$$

Note that this work assumes the convention that random variables are defined in upper case, while their realizations are written in lower case. X is often referred to as the state, while Z_k are the measurements. The operator pdf represents some arbitrary Probability Density Function.

Conditionality on a measurement implies conditionality on the whole sequence of measurements up to that point. Thus, from here onward,

$$\text{pdf}(\cdot | z_k) = \text{pdf}(\cdot | z_1, \dots, z_k), \forall k \in 1..N_t. \quad (5.2)$$

Furthermore, since the PDF of X is only evaluated at measurement times t_k , it is shortened as $X_k = X(t_k)$. This is, $\{X_k\}$ is a discrete Markov chain.

The initial PDF

$$X_0 \sim \text{pdf}(x_0 | \emptyset) \quad (5.3)$$

is known, which is not conditional to any of the measurements, hence the explicit use of the empty set. The dynamic model that dictates the PDF of the state at some time t_k , conditional to the state at some other time t_l , is known, too,

$$X_{k|l} \sim \text{pdf}(x_k | x_l), \quad (5.4)$$

as well as the measurement model that returns the PDF of Z_k at time t_k , given a known state x_k ,

$$Z_k \sim \text{pdf}(z_k | x_k) . \quad (5.5)$$

The non-linear filter cycles over two steps, starting from t_1 , until all measurements are processed. After the filter is initialized with

$$X_{k-1|k-1} = X_0 , \quad (5.6)$$

these steps are:

- **Propagation.** During the propagation step, the posterior PDF of the state from the previous cycle, $\text{pdf}(x_{k-1} | z_{k-1})$, is brought to the new measurement time t_k by marginalization, obtaining the prior

$$\text{pdf}(x_k | z_{k-1}) = \int \text{pdf}(x_k | x_{k-1}) \text{pdf}(x_{k-1} | z_{k-1}) dx_{k-1} . \quad (5.7)$$

- **Update.** In the update step, the new posterior PDF is calculated using the Bayes rule, from the new information brought by the measurement z_k :

$$\text{pdf}(x_k | z_k) = \frac{\text{pdf}(z_k | x_k) \text{pdf}(x_k | z_{k-1})}{\text{pdf}(z_k | z_{k-1})} , \quad (5.8)$$

where $\text{pdf}(z_k | z_{k-1}) = \int \text{pdf}(z_k | x'_k) \text{pdf}(x'_k | z_{k-1}) dx'_k$ is the PDF of the measurement z_k , conditional to all previous measurements.

After having processed all the measurements, the output of the filter is the PDF of the state at the last time step, conditioned to the entire sequence of measurements (Equation 5.1).

Different implementations approximate the propagation and the update steps (Equations 5.7 and 5.8) under different assumptions, such as Gaussianity and local linearisation (*i.e.* EKF), or by modelling the PDF as a weighted sum of Dirac's deltas (*i.e.* PF).

Apart from an initial state, any implementation of a non-linear estimation filter needs to define the dynamic and the measurement models. Such models for the AD problem are described in the following subsections.

5.1.2. Attitude Dynamic Model

As indicated in Section 3.1.3, the attitude model represents attitude as the quaternion that transforms a given vector in body frame B —the frame attached to the object—to the inertial frame I —the GCRF as defined in Dianetti et al. (2018). This is q_B^I . Its kinematics and dynamics follow Equations 3.8 and 3.9, respectively.

However, quaternion representation of attitude conflicts with the non-linear filter scheme presented above: such filter needs to deal with the fact that, although $q_B^I \subset \mathbb{R}^4$, because $\|q_B^I\| = 1$, there are only three DoFs involved. To guarantee uniqueness of the state, the state itself should have as many dimensions as DoFs. This problem is often overcome using Generalized Rodrigues Parameters (GRP), as in *e.g.* Linares et al. (2014a). This attitude

5. Attitude Determination

representation is described in Appendix D. GRP are defined by the map from Equation D.7, which transforms from attitude quaternion to GRP, denoted as $p_{a,f}$, and its inverse (Equation D.8). These equations are reproduced here for the sake of readability:

$$p_{a,f}(\mathbf{q}) = f \frac{\Im(\mathbf{q})}{a + \Re(\mathbf{q})}, \quad (5.9)$$

and

$$\Re(\mathbf{q}) = \frac{-a \|p\|^2 + f \sqrt{f^2 + (1 - a^2) \|p\|^2}}{f^2 + \|p\|^2} \text{ and } \Im(\mathbf{q}) = \frac{a + \Re(\mathbf{q})}{f} p, \quad (5.10)$$

with $a \in \mathbb{R}_{\geq 0}$ and $f \in \mathbb{R}_{> 0}$.

The classical UKF (Julier et al., 2004) can then be implemented by using a reference quaternion plus an error-state based on the GRP (Crassidis et al., 2003; Wetterer et al., 2009). This is,

$$x = \begin{bmatrix} \delta p \\ \omega_{B/I}^B \end{bmatrix}, \quad (5.11)$$

where $\delta p = p(\delta \mathbf{q})$. The error-quaternion $\delta \mathbf{q}$ transforms from some reference body frame B_{ref} to the body frame B represented by state x :

$$\mathbf{q}_{B_{ref}}^B = \delta \mathbf{q}. \quad (5.12)$$

Thus, the quaternion represented by the state is

$$\mathbf{q}_B^I = (\delta \mathbf{q})^* \odot \mathbf{q}_{B_{ref}}^I. \quad (5.13)$$

The dynamic model of the state is represented by

$$x(t + \Delta t) = f(x(t), t, \Delta t, \dots) + w(t, \Delta t), \quad (5.14)$$

where w is white process noise (typically Gaussian) with zero mean and covariance Q , and f is the solution to the system of differential equations composed by Equations 3.8 and 3.9, substituting by Equations 5.9, 5.10 and 5.13 where necessary. All other known parameters, such as the net torque in body frame T^B , are represented by the three dots.

5.1.3. Light Curve Measurement Model

Similar to the shape determination case, this chapter uses the light curve model presented in Chapter 3. The measurement used in this case is the logarithmic version of the brightness measurement, however. This is instrumental magnitude from Equation 3.37. As explained in the aforementioned chapter, it can be expressed as a function of attitude, among other parameters. Thus, in this case, Equation 3.37 can be expressed as

$$m_I = -2.5 \log_{10} F_a = h(\mathbf{q}_B^I, \dots), \quad (5.15)$$

where F_a is the measured photon flux in e^-s^{-1} . Recall from Chapter 3 that this flux does not depend only on attitude, though, hence the dots. It changes with the sensor parameters, which encompass the sensor transmissivity and quantum efficiency; the object shape (a

closed region in space bounded by flat facets) and reflective properties, which define how the light is reflected on its surfaces; the relative geometry between object, observer and Sun. All these factors are considered parameters here, hence their exclusion from the above equation. Thus, for the case at hand, the attitude of the object \mathbf{q}_B^I , which determines the orientation of each of the facets of the object, is the independent variable of interest in this model.

As seen in Chapter 3, one of the cornerstones of the measurement model is the choice of the BRDF that models how light reflects off the surfaces of the object. The AS and CT equations, like any other physically realistic BRDF, are non-linear with respect to the incident/reflected directions of each facet.

Another driving factor of the measurement model is the shape of the object and the effect of self-shadowing, which can be costly to compute for non-convex objects—recall the shadow projection algorithm (Algorithm 1). For small changes in attitude or the observation geometry, sharp edges and other shape features can easily result in abrupt changes in the reflected intensity.

Therefore both the BRDF model and the sensitivity to complex shapes, as well as the logarithm in Equation 5.15, make the function $h(\mathbf{q}_B^I, \dots)$ highly non-linear against the attitude of the object.

To conclude, the measurement space for non-linear filtering is

$$z = [m_I] , \quad (5.16)$$

with m_I being the instrumental magnitude of the object measured at the telescope. The hypothesis taken in this work is that the overall noise effects in the measurement can be modelled as Gaussian noise. Thus, the measurement estimation model is ruled by

$$z(x) = h\left(\mathbf{q}\left(\delta p, \mathbf{q}_{B_{ref}}^I\right), \dots\right) + v(t). \quad (5.17)$$

The function $\mathbf{q}(\delta p, \mathbf{q}_{B_{ref}}^I)$ symbolizes the process of obtaining the attitude quaternion \mathbf{q}_B^I represented by δp (the first three components of x), by means of Equations 5.10 and 5.13. The sensor model noise is represented by $V \sim \mathcal{N}(v | 0, R)$, where $R_k := R$ is equal for all measurements.²

5.2. Motivation for a New Attitude Determination Method

The purpose of this section is to analyse the non-linearities introduced by the dynamic and measurement models of the AD problem defined above, and to pave the way for an efficient non-linear filter implementation that is aware of the degree of non-linearity at all times.

For the sake of clarity, the following shorter notations shall be used from here onward: $\mathbf{q} \equiv \mathbf{q}_B^I$, $\tilde{\mathbf{q}} \equiv \mathbf{q}_{B_{ref}}^I$ and $\omega \equiv \omega_{B/I}^B$.

²From here onward, in this document, $X \sim \mathcal{N}(x | \mu, P)$ reads as *the probability density of a sample x from the random variable X is that of a normal distribution centred around mean μ , with covariance P .*

5.2.1. The Non-Linearity Index

Vittaldev et al. (2016) propose a Non-Linearity Index (NLI) to assess the degree in which an arbitrary transformation deviates from linearity along a given direction, by means of a numerical approximation of the second order derivative of the transformed space against the original one. Schiemenz et al. (2020) adapt the concept to exploit the sigma points of an Unscented Transform (UT) (Julier et al., 2004)—see Appendix F for a detailed description of the UT.

Let $X \in \mathbb{R}^n$ be a random variable that follows a multivariate normal distribution, $X \sim \mathcal{N}(x \mid \mu_x, P_x)$, whose mean vector μ_x and covariance matrix P_x are encoded in a sigma cloud \mathcal{X} , as defined by Equation F.2. Let $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a map that applies a non-linear transformation to x —e.g. let it model a physical process whose associated uncertainty is $V \sim \mathcal{N}(v \mid 0, P_v)$. Being \mathcal{Y} the sigma cloud transformed under g ,

$$\mathcal{Y}(i) = g(\mathcal{X}(i)), \forall i \in 0 \dots 2n, \quad (5.18)$$

which represents the PDF of another random variable Y (under UT assumptions), one can assess the non-linearity introduced by g along each independent direction of the original sigma cloud with the NLI $\phi \in \mathbb{R}^n$ (Schiemenz et al., 2020):

$$\phi_i = \left\| \frac{\mathcal{Y}(i) + \mathcal{Y}(i+n) - \mathcal{Y}(0)}{2(n+\lambda)} \right\|, \forall i \in 1 \dots n, \quad (5.19)$$

where λ is the parameter of the UT that controls the size of the sigma cloud (see Equation F.3). Each dimension i of ϕ , ϕ_i , corresponds to the NLI associated to the direction

$$u_i = \frac{S_i}{\|S_i\|}, \quad (5.20)$$

in the original x space, where S_i is the i -th column of the Cholesky decomposition of $P_x = SS^T$. This is, each direction u_i is the direction along which the pair of sigma points $\mathcal{X}(i)$ and $\mathcal{X}(i+n)$ have been spanned, on opposite sides of $\mathcal{X}(0)$.

One problem of this approach is that, when the characteristic dimensions of the state differ between independent directions, it is difficult to interpret ϕ unless there is some canonical rescaling that can be applied to \mathcal{Y} . The work of Schiemenz et al. (2020) uses a scaling *ad hoc* to the OD problem. In the present work, however the NLI is further modified to circumvent this problem with an elegant and general approach by applying Equation 5.19 to the standardized sigma cloud

$$\bar{\mathcal{Y}}(i) = T^{-1}\mathcal{Y}(i), \forall i \in 0 \dots 2n, \quad (5.21)$$

where $P_y = TT^T$ is the covariance of Y as estimated by the UT, and T its Cholesky decomposition. This is, if Y were normally distributed as $Y \sim \mathcal{N}(y \mid \mu_y, P_y)$, then $\bar{Y} \sim \mathcal{N}(\bar{y} \mid T^{-1}\mu_y, I_{m \times m})$.

Using Equation 5.21, one rescales the transformed sigma cloud so that each dimension of \bar{Y} is associated to each principal direction of P_y , and each of these directions is given the same importance within the NLI. This re-interpretation of the NLI still assesses non-linearity by to what extent straight lines in x space become curves in y space, but it considers two

independent directions equally non-linear if their transformed curves have the same shape relative to the uncertainty associated to each of them. In practice, this could mean that even for transformations where P_x is very small, high NLI values could rise if g is strongly non-linear. However, the fact that the transformed P_y includes the model noise P_v , even when P_x is small, any curvatures in y of a scale smaller than the error volume of P_v will yield a relatively low NLI.

With this new NLI, one can compare the non-linearity, not only between two directions with different characteristic sizes for one given transformation, but also between different models applied to the same random variable. In the case at hand, for instance, it is possible to compare how much non-linearity is introduced by the dynamic (f) and the measurement (h) models relative to each other.

Simple, yet non-linear test functions g (e.g. high order polynomials, exponentials, sines and logarithms) have been explored within the scope of this thesis. Visually, the density histogram of a transformed Monte Carlo (MC) sample starts to differ from Gaussianity for values of $\phi \sim 0.1$, while maximum observed values were around $\phi = 1$. These thresholds, however, still depend on the particularities of each transformation g , and possibly on the size of its input and output spaces.

5.2.2. Non-Linearity Analysis of the Attitude Determination Problem

This subsection focuses on the analysis of the non-linearities in the problem of AD, using the NLI (or ϕ) defined in the previous subsection.

Within the AD problem exposed in the previous section, simulation results suggest that the highest non-linearity contribution to the filtering cycle arises through the measurement model, not through the dynamic model. This is to be expected, given the highly non-linear characteristics of the measurement model (Bernard et al., 2018; Yun et al., 2020).

Figure 5.1 shows the evaluation of the NLI along one UKF run on the scenario used by Wetterer et al. (2009), the details of which are presented below in Section 5.4.1. Each point is the NLI in the maximum-NLI direction—i.e. $\max(\{\phi_i\})$. While blue + marks are the non-linearity computed for the dynamic model (propagation step), orange × marks are those that result from the measurement model (update step). It is clear from the plot that the measurement update has a maximum NLI of at least two orders of magnitude higher than the prediction step, for this case.

Focusing on the measurement update step, Figure 5.2 shows the estimation of the probability distribution of the measurement for the cases of $k = 1, \dots, 5$ of the Wetterer et al. (2009) scenario. In these plots, the grey histogram is a Monte Carlo approximation of the PDF of the estimated measurement obtained by drawing 100 random samples from $\mathcal{N}(x_k | \hat{x}_{k|k-1}, P_{k|k-1})$, where $\hat{x}_{k|k-1}$ and $P_{k|k-1}$ are the prior mean vector and covariance matrix of the state—the histogram contains the magnitude estimated from each random sample. Then, the dashed green curve is the PDF $\mathcal{N}(z_k | \tilde{z}_k, R_k)$, a Gaussian centred around the measured (sample) \tilde{z}_k with uncertainty covariance R_k . Finally, the solid orange one is $\mathcal{N}(z_k | \hat{z}_k, P_k^{zz} - R_k)$, where \hat{z}_k and P_k^{zz} are the mean and covariance obtained from applying the UT with the measurement model to $\hat{x}_{k|k-1}$ and $P_{k|k-1}$. In these figures, it is

5. Attitude Determination

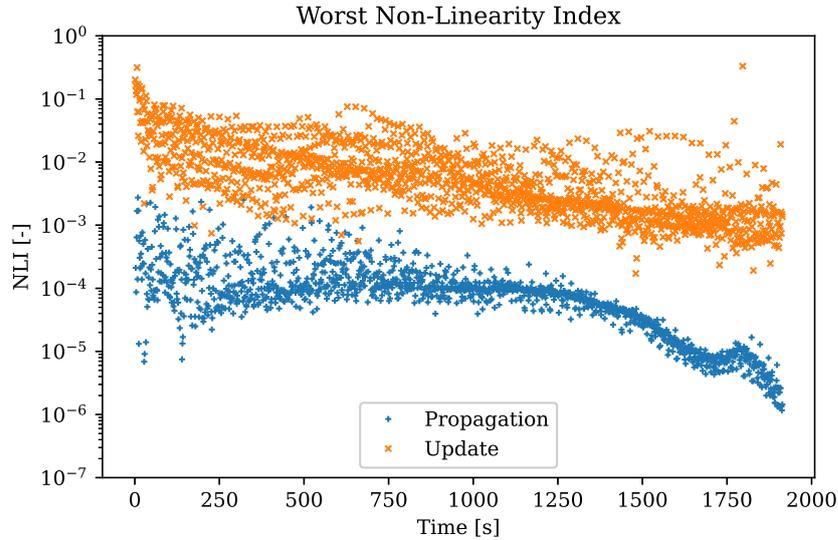


Figure 5.1.: Non-linearity index of one UKF run, based on the scenario from Section 5.4.1 (Vallverdú Cabrera et al., 2023)

obvious that the UT-estimated PDF (orange solid curve) and the more realistic MC-estimated one (grey histogram) are in significant disagreement.

To sum up, the goal of this chapter is to define a non-linear adaptive GM filter architecture that focuses on dealing with the measurement model non-linearity.

5.3. The Adaptive Gaussian Mixtures Unscented Kalman Filter

This section defines the Adaptive Gaussian Mixtures Unscented Kalman Filter for Attitude Determination. The algorithm proposed below adapts the idea introduced by Vittaldev et al. (2016), which Schiemenz et al. (2020) apply to the classical Orbit Determination use case with density uncertainty realism. The AGMUKF version developed in the present work focuses on capturing the high non-linearity introduced by the measurement model, as justified by the non-linearity analysis above.

The working principle of the AGMUKF is to estimate the state PDF as a mixture of Gaussian distributions, and then to apply normal UKF steps to each member of the mixture. In contrast to the classical EKF/UKF, which represent the state PDF as a single Gaussian, the mixture approach can represent arbitrarily complex PDFs. Furthermore, the AGMUKF is labelled as *adaptive* because it exploits 1) the NLI defined in Equation 5.19 to refine the mixture, together with 2) a distance metric to coarsen it. Its aim is for the mixture to always approximate the state PDF as faithfully as possible, yet with the minimum number of underlying Gaussian kernels. Thus, careful analysis of the problem is necessary to decide where in the filter steps refinement and coarsening should happen. If this is done right, the AGMUKF may be more efficient than other complex-PDF-modelling filters, such as the PF and its variants (Du et al., 2018; Coder et al., 2017).

In the OD case of Schiemenz et al. (2020), the major non-linearity contribution happens

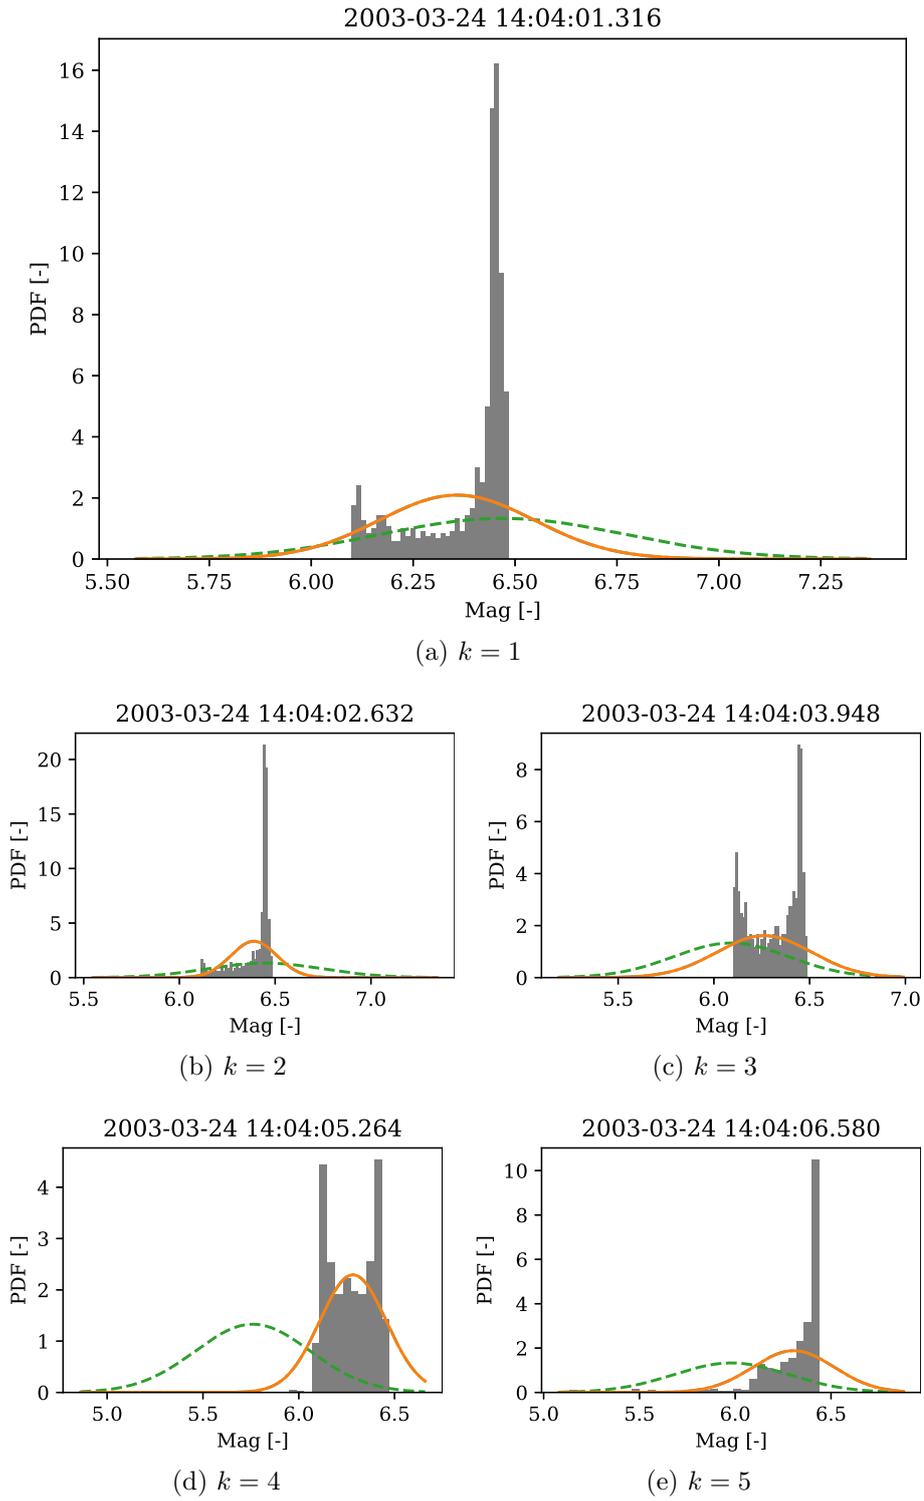


Figure 5.2.: Probability density function of \hat{z}_k for the first five measurements in the scenario from Section 5.4.1 when using the UKF. Legend: PDF of the observation (dashed green line); PDF of the measurement estimated from the state using UT (solid orange line) or Monte Carlo (gray histogram). Dates in the titles expressed in ISO format: YYYY-MM-DD hh:mm.ss.fff (f stands for millisecond) (Vallverdú Cabrera et al., 2023)

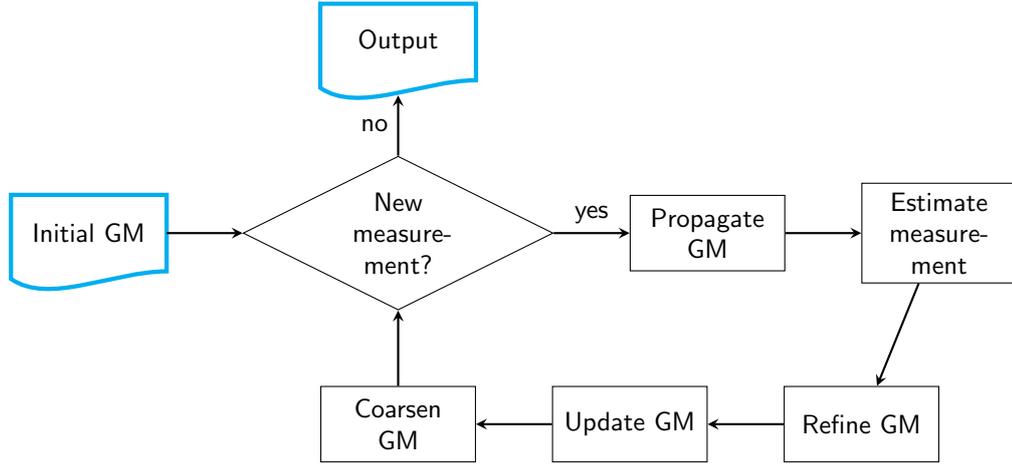


Figure 5.3.: Diagram of the AGMUKF algorithm for AD

within the dynamic model, where orbit propagation over long spans of time is coupled with realistic modelling of density uncertainty. This motivates the refinement and coarsening steps to alternate with consecutive propagation steps (without updates in-between). In this case, the yardstick used to decide whether and how to split is the NLI of dynamic model.

In the AD problem, however, non-linearity comes mainly from the measurement model, and propagation time spans are relatively short for a dense light curve. Therefore, the implementation of the AGMUKF for AD introduced here refines the mixture according to the measurement model NLI, instead, and does so just before the measurement update step. This way, the measurement PDF is modelled with finer detail. Right after the update, the filter merges the mixture back, so that the state PDF does not remain oversampled. Figure 5.3 summarizes this workflow.

Following, each AGMUKF step is described in detail. Of these, the UKF steps of the AGMUKF are mainly adapted from and inspired by Crassidis et al. (2003), Wetterer et al. (2009), and Du et al. (2018).

5.3.1. Probability Distribution

Following the non-linear filtering notions introduced in Section 5.1.1, the AGMUKF estimates the PDF of the state X as a Gaussian Mixture (GM) with M kernels:

$$X \sim \text{pdf}(x) = \sum_{j=1}^M w^j \mathcal{N}(x | \mu^j, P^j), \quad (5.22)$$

where w^j is the weight of kernel j , which follows a normal distribution with mean vector μ^j and covariance matrix P^j . Note that $w^j = \text{pdf}(\eta = j)$ is the value of the Probability Mass Function (PMF) of a discrete random variable $H \sim \text{pdf}(\eta)$, with $\eta \in 1 \dots M$. Thus,

$$\sum_{j=1}^M w^j = 1. \quad (5.23)$$

Furthermore, $\text{pdf}(x | \eta = j) = \mathcal{N}(x | \mu^j, P^j)$, so the expression in Equation 5.22 is just $\text{pdf}(x)$ obtained through marginalization.

The first and second statistical moments of a GM are, then

$$\mu = \sum_{j=1}^M w^j \mu^j \quad (5.24)$$

and

$$P = \sum_{j=1}^M w^j [P^j + (\mu^j - \mu)(\mu^j - \mu)^T] \quad (5.25)$$

respectively (Schiemenz et al., 2020).

5.3.2. Filter Initialization

The filter is initialized with the known PDF of the state at some time t_0 , represented by the reference quaternion $\tilde{\mathbf{q}}_0$, and initial GM with M_0 kernels represented by

$$X_0 \sim \sum_{j=1}^{M_0} w_0^j \mathcal{N}(x_0 | \hat{x}_0^j, P_0^j), \quad (5.26)$$

where \hat{x}_0^j and P_0^j are the a priori estimated means and covariances of the state. Each kernel of the filter is initialized as in the conventional UKF algorithm. This is, the initial sigma cloud of each kernel, \mathcal{X}_0^j , composed by $2n + 1$ points in state space (where $n = 6$) is obtained using Equation F.2.

The rest of the steps are executed in a loop until all measurements have been processed.

5.3.3. Propagation Step

For each new measurement sample \tilde{z}_k , measured at t_k , the state PDF first needs to be propagated from the previous measurement update, at t_{k-1} , to t_k . This is, this subsection deals with obtaining the prior PDF $\text{pdf}(x_k | z_{k-1})$ from the last known posterior $\text{pdf}(x_{k-1} | z_{k-1})$.

First, the sigma cloud of each kernel is propagated using the dynamic model defined by the map $f : \mathbb{R}^n \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^n$ (Equation 5.14). To do so, the δp component of each sigma point needs to be converted back to quaternion state using Equations 5.10 and 5.13, obtaining $\mathbf{q}_{k-1|k-1}^j(i)$. With this, each sigma point is propagated from t_{k-1} to t_k using Equations 3.8 and 3.9 and a numerical integration routine (Shampine et al., 1975), which yields $\mathbf{q}_{k|k-1}^j(i)$.

The prior reference quaternion needs to be redefined, so that the sigma cloud of each kernel, whose points are in quaternion form, can be converted back to state space—*i.e.* the GRP form. To do so, one propagates the first statistical moment of the previous step, a.k.a. the merged mean $\hat{x}_{k-1|k-1}$ (obtained from the set $\{w^j, \hat{x}_{k-1|k-1}^j\}$ with Equation 5.24) to t_k , using the same method of GRP-to-quaternion as with the sigma points. This yields the

5. Attitude Determination

merged, propagated attitude quaternion estimate $\mathbf{q}_{k|k-1}$, which becomes the new reference quaternion:

$$\tilde{\mathbf{q}}_{k|k-1} = \mathbf{q}_{k|k-1}. \quad (5.27)$$

With the new reference, the sigma cloud \mathcal{X}_k^j of each kernel is recovered by using the inverse of Equation 5.13 followed by Equation 5.9 to obtain the δp representation of each $\mathbf{q}_{k|k-1}^j(i)$.

At this point, the estimated prior mean $\hat{x}_{k|k-1}^j$ and covariance $P_{k|k-1}^j$ of each kernel are obtained with the conventional UT: Equations F.5 and F.6 are used on the propagated cloud \mathcal{X}_k^j , using the process noise covariance matrix of step k , Q_k , as model noise. Note that for the clouds to include the additional process noise covariance Q_k , they need to be redrawn using Equation F.2 (Wu et al., 2005). Under the assumption that all kernels are affected by the same process noise, Q_k can be added to each one of them with the overall effect of increasing the merged covariance by the same amount (Schiemenz et al., 2020).

Finally, assuming the propagation is linear enough, the kernel weights can remain constant (Schiemenz et al., 2020):

$$w_{k|k-1}^j = w_{k-1|k-1}^j. \quad (5.28)$$

Likewise, the number of kernels does not change either, so $M_{k|k-1} = M_{k-1|k-1}$. With this, the propagated prior state PDF is

$$X_{k|k-1} \sim \sum_{j=1}^{M_{k|k-1}} w_{k|k-1}^j \mathcal{N}\left(x_k \mid \hat{x}_{k|k-1}^j, P_{k|k-1}^j\right). \quad (5.29)$$

5.3.4. Measurement Estimation

The traditional update step of the UKF, where the measurement PDF is used to update the prior state PDF $\text{pdf}(x_k \mid z_{k-1})$ to the posterior $\text{pdf}(x_k \mid z_k)$, has been split into multiple steps, to accommodate the Gaussian splitting and merging strategies that characterize the AGMUKF for AD.

This section describes the first part of the UKF update step, which consists in applying the measurement model on the sigma cloud of each kernel, using the map $h : \mathbb{R}^3 \times \mathbb{A} \rightarrow \mathbb{R}$ from Equation 5.17. This yields the measurement sigma clouds \mathcal{Z}_k^j . Next, the means \hat{z}_k^j and innovation covariance matrices $P_k^{zz,j}$ are calculated using again the UT (Equations F.5 and F.6), this time applied on \mathcal{Z}_k^j and using the measurement noise covariance matrix R_k to represent model noise.

This step finishes with an estimation of the innovation PDF,

$$\text{pdf}(z_k \mid x_k) = \sum_{j=1}^{M_{k|k-1}} w_{k|k-1}^j \mathcal{N}\left(z_k \mid \hat{z}_k^j, P_k^{zz,j}\right). \quad (5.30)$$

5.3.5. Refinement Step

Because of the high non-linearity of h , before updating the state with the measurement information, it is necessary to check whether the innovation PDF has become too non linear, in which case the prior PDF should be split into a finer GM, to minimize the rise of non-linearities in the measurement estimation step from Section 5.3.4.

To do so, the NLIs of the measurement model in each kernel are computed from each sigma cloud \mathcal{Z}_k^j with Equation 5.19, using $P_k^{zz,j}$ as the scale term in Equation 5.21. These NLIs are used to decide which kernels should be split, into how many parts, and along which directions.

The split library used is the one developed by Vittaldev et al. (2016), which is homoscedastic—meaning that the resulting new kernels have the same common covariance. This library allows splitting one kernel among up to 39 new ones, only in odd numbers.

Following the implementation in Schiemenz et al. (2020), the split recommendation for kernel j in each direction i in state space takes the form

$$n_i^j = \text{odd} \left(a \left(\log_{10}(\phi_i^j) + b \right) \right), \quad (5.31)$$

where $\text{odd}(x) = 2\lceil x/2 \rceil - 1$.

The splitting library works in the standard normal univariate distribution. For a split of L new kernels, it provides the set of weights and means $\{\bar{w}_{l|L}, \bar{\mu}_{l|L} \mid l \in 1..L\}$, as well as the covariance $\bar{\sigma}_L^2$, that result in the symmetric univariate GM with merged mean 0 and covariance 1—*i.e.* the moments of the standard normal distribution. Therefore, to split along the direction $u \equiv u_i^j$ (Equation 5.20) from kernel j with weight $w = w_{k|k-1}^j$, mean $\mu \equiv \hat{x}_{k|k-1}^j$ and covariance $P \equiv P_{k|k-1}^j$, first the original multivariate normal of the kernel needs to be projected onto the standard normal along direction u . This is done by linearly transforming the state using *e.g.* the Cholesky decomposition S of $P = SS^T$ (Vittaldev et al., 2016), so that

$$\bar{u} = \frac{S^{-1}u}{\|S^{-1}u\|}. \quad (5.32)$$

Then, the $L \equiv n_i^j$ new kernels are obtained as the sequence of new weights

$$w_l = w \cdot \bar{w}_{l|L}, \quad (5.33)$$

means

$$\mu_l = \mu + \bar{\mu}_{l|L} S \bar{u} \quad (5.34)$$

and covariances

$$P_l = S \left(I_n + (\bar{\sigma}_L^2 - 1) \bar{u} \bar{u}^T \right) S^T \quad (5.35)$$

for all $l \in 1..L$, where I_n is the $n \times n$ identity matrix.

After the split, all the newly created kernels need to have their sigma cloud \mathcal{X}_k^j redrawn (Equation F.2), so that the measurement estimation step from Section 5.3.4 can be applied on them, too.

If several directions in the same kernel are recommended for refinement, the split operation can be applied first in one direction, and then in the other one to the new kernels. As long as

5. Attitude Determination

the directions are orthogonal, the order in which these are applied is irrelevant (Schiemenz et al., 2020).

Now, the prior and innovation PDF equations are equal to Equations 5.29 and 5.30, with the difference that there are $M'_{k|k-1} \geq M_{k|k-1}$ kernels. This results in a denser mixture that captures the PDF in measurement space with significantly higher accuracy, even if the PDF in state space remains virtually the same as before the split.

5.3.6. Update step

Having split the state PDF to mitigate the appearance of non-linearities in each UKF kernel, one can proceed with the normal UKF steps to obtain the posterior PDF of the state, $\text{pdf}(x_k | z_k)$.

On a kernel basis, the UKF cross covariance $P_k^{xz,j}$ is computed by applying Equation F.8 on the state \mathcal{X}_k^j and measurement \mathcal{Z}_k^j sigma clouds, in this order. With this, one obtains the Kalman gains

$$K_k^j = P_k^{xz,j} \left(P_k^{zz,j} \right)^{-1} \quad (5.36)$$

and updates the kernel means and covariances with

$$\hat{x}_{k|k}^j = \hat{x}_{k|k-1}^j + K_k^j \left(\tilde{z}_k - \hat{z}_k^j \right) \quad (5.37)$$

and

$$P_{k|k}^j = P_{k|k-1}^j - K_k^j P_k^{zz,j} \left(K_k^j \right)^T . \quad (5.38)$$

In the update step, actual new information has been introduced in the state PDF. Thus, the kernel weights must be updated. This is done by rescaling each kernel weight with the innovation likelihood of the prior residual (Schiemenz et al., 2020),

$$w_{k|k}^j = w_{k|k-1}^j \mathcal{N} \left(\tilde{z}_k | \hat{z}_k^j, P_k^{zz,j} \right) . \quad (5.39)$$

Now, the posterior PDF looks like

$$X_{k|k} \sim \sum_{j=1}^{M'_{k|k-1}} w_{k|k}^j \mathcal{N} \left(x_k | \hat{x}_{k|k}^j, P_{k|k}^j \right) . \quad (5.40)$$

5.3.7. Coarsening step

After the measurement update, it is often the case that the GM that represents the posterior state PDF is oversampled, due to all the extra kernels added to the mixture to properly model the non-linearity in the measurement estimation. Thus, a merge step is performed, where kernels that are too close to each other in state space are joined within a moment-preserving merge.

The moment-preserving merge of any two kernels i and j in the mixture is the weight-normalized version of Equations 5.24 and 5.25 applied to the two kernels only—the subscript

$_{k|k}$ of weights, means and covariances has been omitted for brevity since it applies everywhere in this subsection:

$$w^{ij} = w^i + w^j, \quad (5.41)$$

$$\mu^{ij} = \frac{w^i \mu^i + w^j \mu^j}{w^{ij}} \quad (5.42)$$

and

$$P^{ij} = \frac{w^i P^i + w^j P^j}{w^{ij}} + \frac{w^i w^j}{(w^{ij})^2} (\mu^i - \mu^j)(\mu^i - \mu^j)^T. \quad (5.43)$$

To coarsen the mixture, the same strategy used by Schiemenz et al. (2020) is used here. This is, the first step is to compute the Kullback-Leibler (KL) distance (Runnalls, 2007) between all kernel pairs (i, j) ,

$$B(i, j) = \frac{1}{2} [w^{ij} \log(|P^{ij}|) - w^i \log(|P^i|) - w^j \log(|P^j|)]. \quad (5.44)$$

Then, if the closest pair is at a distance lower than a given threshold, the two affected kernels are substituted by the new merged one. The distance of the new kernel relative to all the other kernels is computed anew, and the cycle is repeated until the closest pair is at a distance greater than the threshold. Refer to Schiemenz et al. (2020) for the reasons to choose the KL from all the distance metrics available in the literature.

At the end of the merge step, the posterior PDF looks similar to Equation 5.40, but with $M_{k|k} \leq M'_{k|k-1}$ kernels due to the merge.

5.3.8. Filter Reset

To close the filtering cycle, the posterior reference quaternion needs to be updated again. To do so, the merged mean of the posterior state PDF, $\hat{x}_{k|k}$, is computed as in Equation 5.24, from which the posterior merged error GRP, $\delta p_{k|k}$, is extracted. The merged error quaternion $\delta \mathbf{q}_{k|k}$ is obtained using Equation 5.10 on $\delta p_{k|k}$. Similar to the propagation step, the posterior reference is computed from the prior one as

$$\tilde{\mathbf{q}}_{k|k} = \delta \hat{\mathbf{q}}_{k|k}^* \odot \tilde{\mathbf{q}}_{k|k-1}. \quad (5.45)$$

With the new reference, the kernel means need to be shifted, too. In UKF implementations for attitude states with reference quaternion plus error state (Crassidis et al., 2003), this is commonly achieved by resetting them with the merged error GRP:

$$\delta \hat{p}_{k|k}^j \leftarrow \delta \hat{p}_{k|k}^j - \delta p_{k|k}. \quad (5.46)$$

This process makes the merged mean of the mixture yield exactly $\delta p_{k|k} = 0$, meaning that the reference quaternion is exactly the one represented by the merged mean of the state. Equation 5.46, however, is only a first order approximation (see Appendix D.4) that is only valid if both $\delta \hat{p}_{k|k}^j$ and $\delta p_{k|k}$ are very small. This is appropriate for filters like the plain UKF, which assume small uncertainties relative to the total surface of the unit quaternion hypersphere. However, the PDF of the state within the AGMUKF may be composed, at any given step, by a disperse mixture spanning across wide regions of the attitude space. This

5. Attitude Determination

means that the errors introduced by Equation 5.46 are likely to be significant if used within the AGMUKF, especially when the second statistical moment of the mixture is relatively high.

In this thesis, the proposed solution is to exploit the UKF kernel of the AGMUKF to reset the state of each kernel relative to the new reference quaternion by means of an UT, so that at least the second order statistical moment of each kernel is preserved. To do so, first, the sigma cloud of each kernel is redrawn using Equation F.2 from its mean $\hat{x}_{k|k}^j$ and covariance $\hat{P}_{k|k}^j$. Then, each sigma point i of each kernel j is modified as

$$\delta p_{k|k}^j(i) \leftarrow p \left((p^{-1}(\delta p_{k|k}))^* \odot p^{-1}(\delta p_{k|k}^j(i)) \right), \quad (5.47)$$

which is the actual quaternion composition that ensures that each sigma point still represents the same quaternion when composing it with the new reference quaternion. Then, the mean and covariance of each kernel are recovered using the normal UT from Equations F.5 and F.6, as was done during the propagation step. Due to the non-linearity of this operation, the new merged mean will not yield exactly $\delta p_{k|k} = 0$; but because it will still approach Equation 5.46 for small error states, it will in general imply a reduction of $\|\delta p_{k|k}\|$ proportional to second order error—*i.e.* the reference quaternion will still be a closer representation to the updated merged state than before the reset.

5.4. Simulation Scenarios and Results

In this section, different scenarios of the AD problem are presented and discussed. It is worth mentioning that all the results shown in this section have been obtained using the Square Root Unscented Kalman Filter (SRUKF), which is mathematically equivalent to the plain UKF steps presented in the previous section, but guarantees the positive-definiteness of the covariance matrix. It synergises well with the AGMUKF, since the Cholesky decompositions of the covariance matrices are needed anyway to establish the split directions. Refer to Van der Merwe et al. (2001) for a detailed development of the core SRUKF.

Furthermore, all cases presented here use $a = 1$ and $f = 2(a + 1) = 4$ for the GRP implementation, as *e.g.* in Du et al. (2018). The choice of $a = 1$ prevents the δp part of the state from tending to infinity at the extremes, in contrast to $a = 0$ (Crassidis et al., 2007), while $f = 2(a + 1) = 4$ makes it approach Euler angle errors in first-order Taylor expansion (see Appendix D.4).

5.4.1. AGMUKF vs. UKF

The first validation step of the AGMUKF has been to compare it against a plain UKF, which is in this case equivalent to the AGMUKF algorithm presented above, with only one kernel and skipping the splitting and merging steps, plus using Equation 5.46 to reset the error state.

The simulated scenario reproduces the exact same case presented in Wetterer et al. (2009). This is, the Atlas-Centaur II upper-stage with NORAD number 12069 is simulated based on its TLE, available from Space-Track.org (U.S. Government, 2019). The initial attitude,

angular velocity, measurement acquisition rate and simulation time span are unchanged w.r.t. Wetterer et al. (2009). The shape model conserves the same characteristics: a cylinder with 9 m length, 1.5 m radius, and two caps on either side. In the current work, the caps are modelled as cones of 1.5 m height, instead of flattened hemispherical ones (see Figure 5.4). Both cylinder and cones are discretized using 100 flat facets each. As in Wetterer et al. (2009), the sides of the cylinder (blue) follow a Cook Torrance BRDF (Section 3.1.4) with $\rho = F_0 = 0.2$ and $\mathfrak{d} = 0.3$, with $\mathfrak{m} = 0.17$. The cones, also Cook Torrance, have the same \mathfrak{m} and $\mathfrak{d} = 0.8$; one cone has $\rho = F_0 = 0.6$ (orange), while the other one has $\rho = F_0 = 0.3$ (green). The inertia tensor is computed assuming the object is solid, with constant density. Regarding the filter, the same measurement model uncertainty, initial state uncertainty and process noise are used. Finally, the simulated observer is the AEOS telescope.

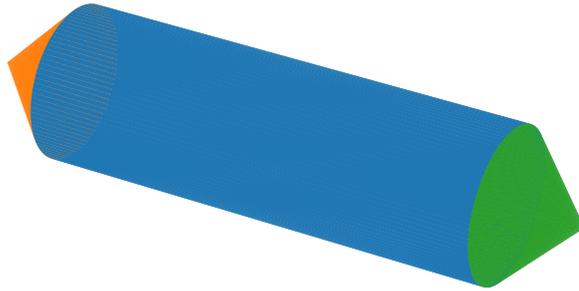


Figure 5.4.: Simulated object shape (Vallverdú Cabrera et al., 2023)

The goal of this experiment is to reproduce the Wetterer et al. (2009) scenario as faithfully as possible. The main reason is to assess the AGMUKF against a known-working case of the UKF. However, the UKF is the work-horse underneath the AGMUKF. Therefore, the second reason is to validate the UKF implementation itself used in this thesis.

The UKF (as well as the AGMUKF) has been tuned with $\alpha = 0.5$, $\beta = 2$ and $\kappa = 0$ (refer to Appendix F). Regarding the parameters only relevant to the AGMUKF, the split criteria from Equation 5.31 have been set to $a = 15$ and $b = 1$. The KL-distance threshold has been set to 10^{-4} —*i.e.* kernel pairs at a relative distance of less than that value are merged in the merge step. These values have been obtained empirically by trying to balance the number of kernels and filter performance. Additionally, the maximum number of kernels allowed after any refinement step is set to $M_{\max} = 999$. Whenever a split step would generate more than that, the following measures are taken:

1. The kernels are split by order of decreasing weight. This ensures that the splitting effort is focused on the kernels with the highest likelihoods.
2. While splitting any particular kernel j , if it would exceed the maximum number of total kernels, the split recommendation vector is reduced by 2 (*i.e.* $n_i^j \leftarrow \max(n_i^j - 2, 0)$) iteratively until the split operation does not generate more kernels than the threshold, or until all $n_i^j = 0$, in which case the kernel is skipped from splitting. The jumps of 2 by 2 arise from the fact that the splitting libraries of Vittaldev et al. (2016) only allow odd-numbered splits.
3. The kernel splitting step is stopped as soon as the mixture has $M_{\max} - 1$ kernels, since the minimum possible split would add 2 new kernels as per the libraries of Vittaldev

5. Attitude Determination

et al. (2016).

From the three available splitting libraries (Vittaldev et al., 2016), the one that yields the smallest individual kernel covariance is used ($\sigma^2 = M^{-1}$), because it is the one that can split a kernel in more pieces (39), and because the purpose of splitting is to reduce the uncertainty region of each kernel, so that the subsequent UTs are closer to local linearity.

Having set the scenario, a Monte Carlo study of 100 iterations is performed. In each Monte Carlo iteration, both a UKF and an AGMUKF are initialized with the same random realization \hat{x}_0 from $\hat{X}_0 \sim \mathcal{N}(\hat{x}_0 | \mu_0, P_0)$, where P_0 is the one used in Wetterer et al. (2009),

$$P_0 = \text{diag} \left([4 \cdot 10^{-2} \quad 4 \cdot 10^{-2} \quad 4 \cdot 10^{-2} \quad 10^{-24} \quad 10^{-24} \quad 10^{-6}] \right), \quad (5.48)$$

while

$$\mu_0 = [0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0.252 \frac{\text{rad}}{\text{s}}]^T. \quad (5.49)$$

The quaternion represented by the $\delta\hat{p}_0$ portion of \hat{x}_0 is set as the initial reference quaternion,

$$\tilde{\mathbf{q}}_0 = (q(\delta\hat{p}_0))^* \odot \mathbf{q}_0, \quad (5.50)$$

after which $\delta\hat{p}_0$ is reset to 0. The quaternion

$$\mathbf{q}_0 = 0.743 - 0.171\mathbf{i} + 0.050\mathbf{j} - 0.645\mathbf{k} \quad (5.51)$$

is the true initial attitude quaternion of the object. The Wetterer et al. (2009) covariance P_0 is used to initialize the filter—the AGMUKF starts with a single kernel. Measurements \tilde{z}_k are likewise realizations of $Z_k \sim \mathcal{N}(z_k | z(t_k), R_k)$, where $z(t_k)$ is the true measurement in the absence of sensor noise at time t_k , obtained by propagating the initial true attitude and angular velocities to t_k with Equations 3.8 and 3.9, and then applying the measurement model from Equation 5.17 with $v = 0$. Likewise, the filter uses R_k as the measurement covariance, too.

The first part of this subsection focuses on the analysis of one arbitrary filter run whose initial reference quaternion and angular velocity are

$$\tilde{\mathbf{q}}_0 = 0.756 - 0.070\mathbf{i} + 0.237\mathbf{j} - 0.605\mathbf{k} \quad (5.52)$$

and

$$\hat{\omega}_0 = [-1.10 \cdot 10^{-12} \quad -7.79 \cdot 10^{-13} \quad 0.253]^T \frac{\text{rad}}{\text{s}}. \quad (5.53)$$

Both simulation and filter assume zero net torque, and the filter assumes a process noise of (Wetterer et al., 2009)

$$Q_k = \text{diag} \left([4 \cdot 10^{-8} \quad 4 \cdot 10^{-8} \quad 4 \cdot 10^{-8} \quad 10^{-24} \quad 10^{-24} \quad 10^{-12}] \right). \quad (5.54)$$

The measured (*i.e.* noisy) light curve is shown in Figure 5.5. In the three plots, the solid blue line connects the measurements, which are marked with a \times symbol. While Figure 5.5a shows the whole simulation time span, Figures 5.5b and 5.5c contain the first and last 3 minutes, respectively, so that finer detail can be appreciated. The dashed orange curves symbolize the upper and lower bounds of the regions under which the cumulative probability of the measurement is about 99.73%. For the measurement, which follows a Gaussian distribution

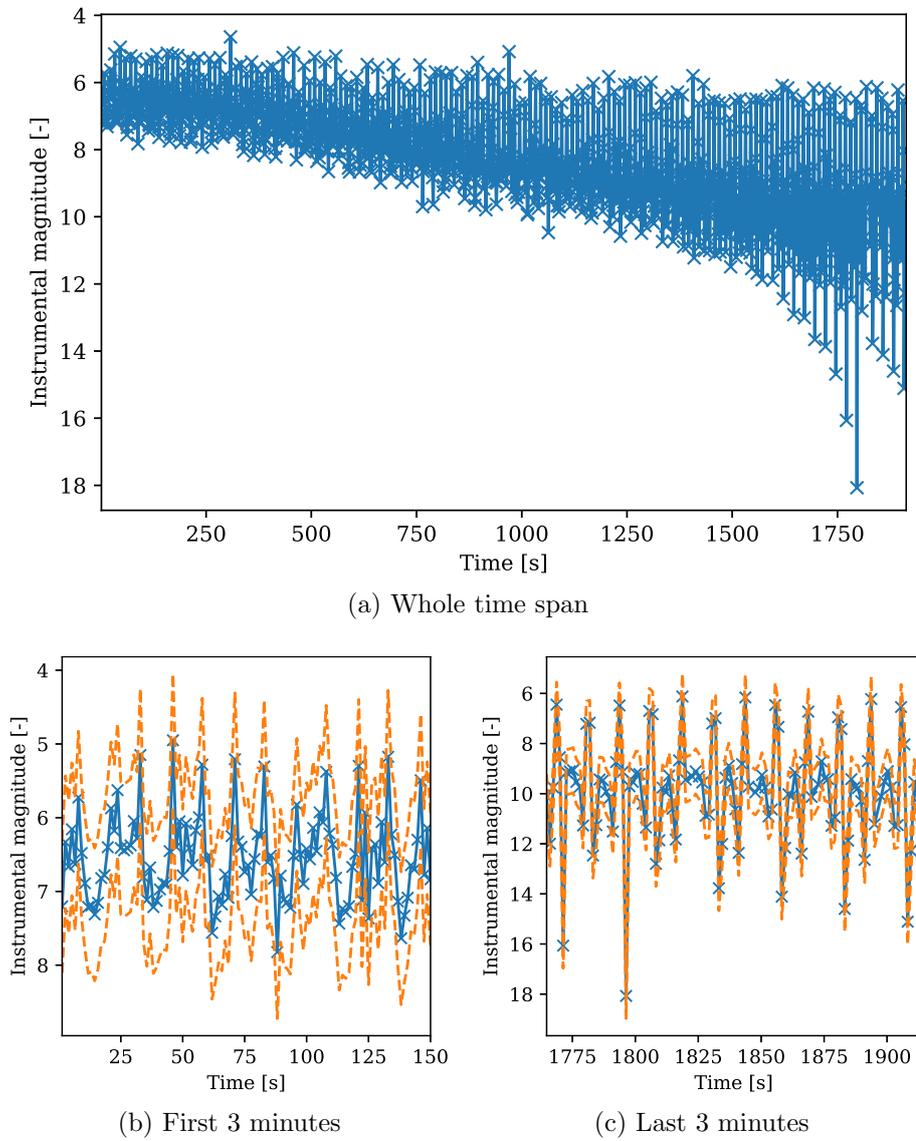


Figure 5.5.: One noisy light curve of the Wetterer et al. (2009) scenario. Blue crosses are brightness measurements, while the blue line is a visual guide. The orange dashed line is the $3\text{-}\sigma$ envelope of uncertainty around the light curve (Vallverdú Cabrera et al., 2023)

5. Attitude Determination

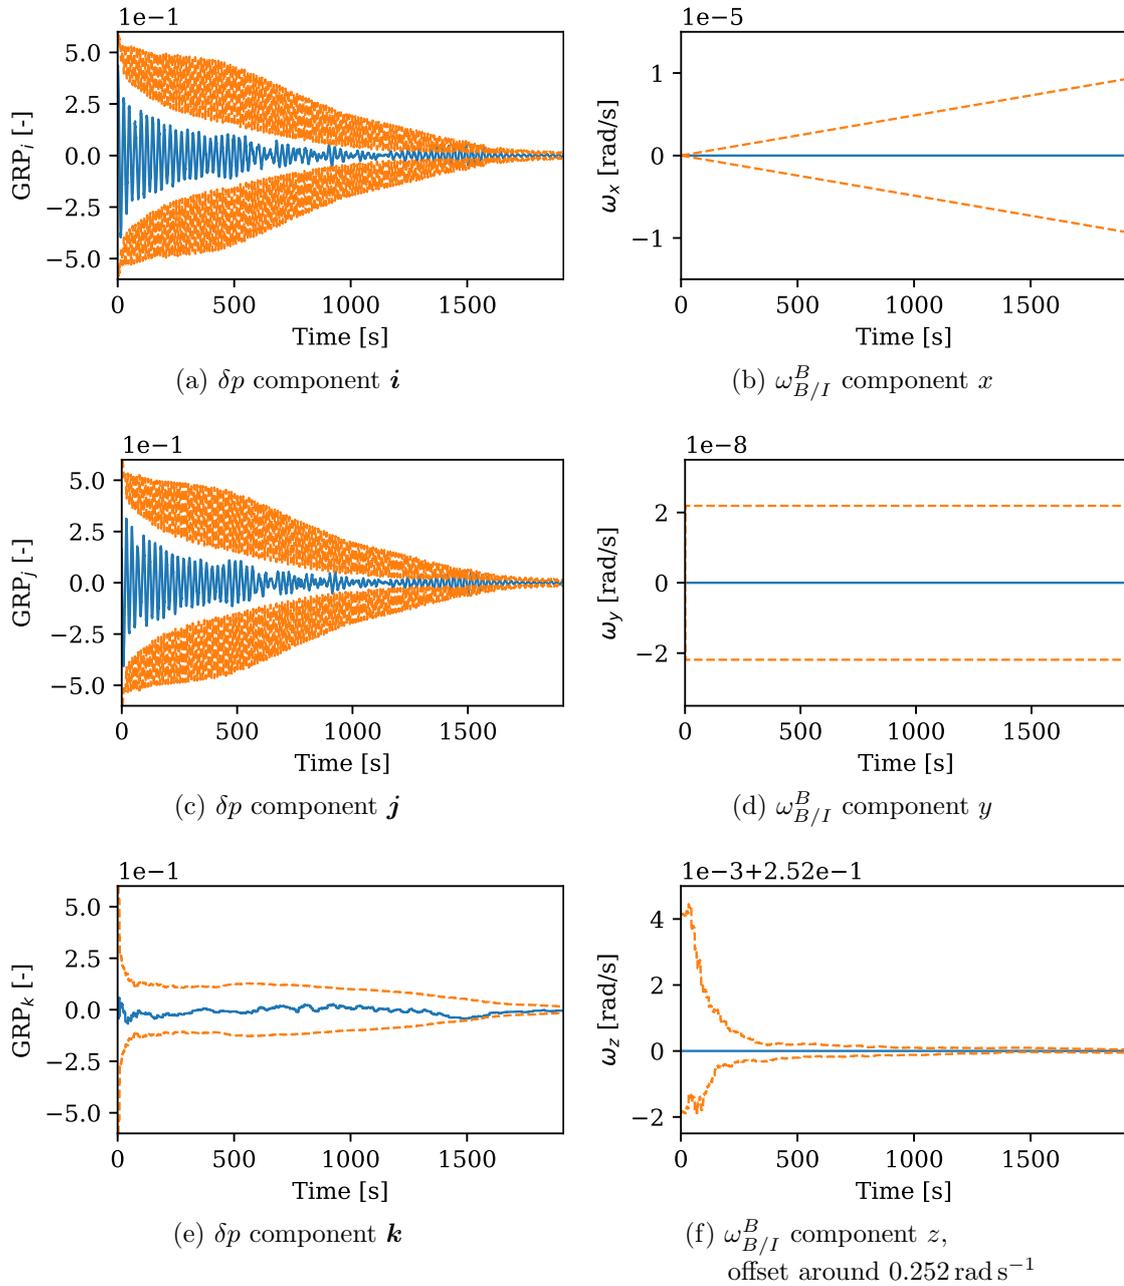


Figure 5.6.: One UKF run on the Wetterer et al. (2009) scenario. Legend: truth (blue line); 3- σ envelope of the estimated PDF (dashed orange line) (Vallverdú Cabrera et al., 2023)

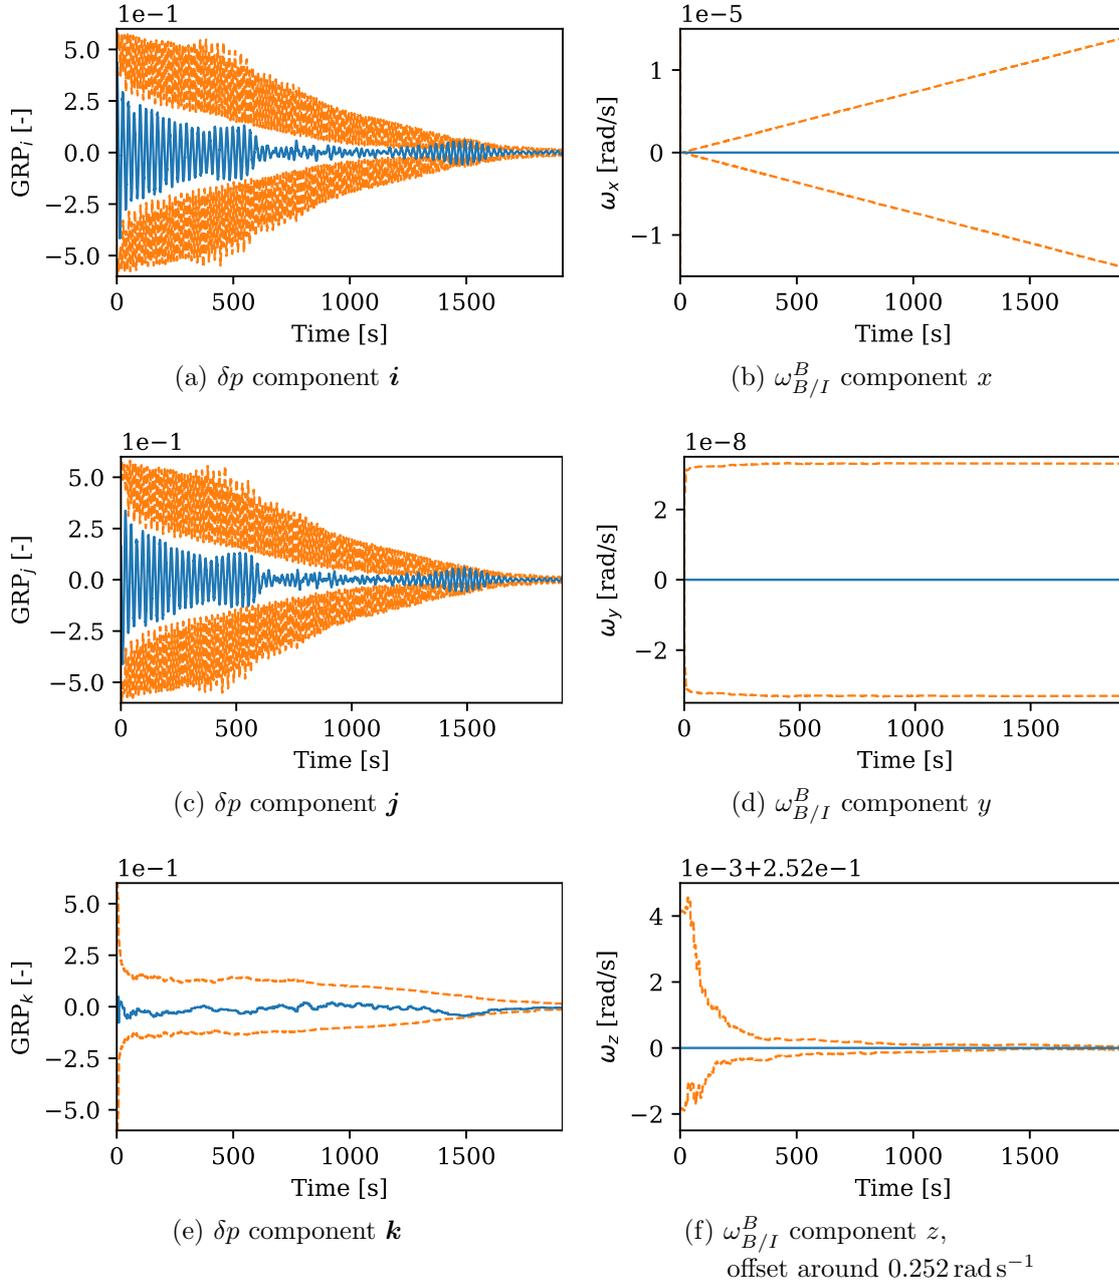


Figure 5.7.: One AGMUKF run on the Wetterer et al. (2009) scenario. Legend: truth (blue line); 99.73% probability envelope of the estimated PDF (dashed orange line) (Vallverdú Cabrera et al., 2023)

5. Attitude Determination

with covariance $R = 0.3^2$ this is the $\pm 3\sqrt{R}$ region around the measurement realization (solid blue curve), which acts as an estimated mean in this case.

Figure 5.6 shows the evolution of the estimated state PDF as per the UKF, for this particular Monte Carlo run, while Figure 5.7 has the same data, but for the AGMUKF. The solid blue curve in each plot is the true state in the specified dimension—*i.e.* either error GRP w.r.t. the reference quaternion of the filter run, or plain angular velocity. The orange dashed curves enclose the 99.73% cumulative probability of the marginal PDF in that dimension.

For the UKF case, since the PDF is modelled as a Gaussian, for each state $i \in 1 \dots n$, this still is the $\pm 3\sigma_i$ offset from the filter estimated mean \hat{x}_i , where σ_i is the square root of the i -th diagonal element of P . Note that this applies to all prior and posterior estimations in all time steps—*i.e.* here sub-indices do not indicate the time step, as before, but the state dimension.

For the AGMUKF, however, at each time step of the filter (both prior and posterior), this region is defined as the sub-set Ω_i of the domain of each state such that, for probability $p_r = 99.73\%$,

$$p_r = \int_{\Omega_i} \text{pdf}_i(x_i) dx_i \quad (5.55)$$

and

$$\text{pdf}_i(x_i) \geq \text{pdf}_i(x'_i), \forall x_i \in \Omega_i \wedge x'_i \notin \Omega_i, \quad (5.56)$$

where

$$\text{pdf}_i(x_i) = \sum_{j=1}^M w^j \mathcal{N}\left(x_i \mid \hat{x}_i^j, (\sigma_i^j)^2\right) \quad (5.57)$$

is the univariate marginal GM of state dimension i . In words, this is the region with cumulative probability 99.73% such that the PDF at any point within this region evaluates higher than at any point without. It can be proven that Ω_i is unique given p_r and pdf_i (Kesemen et al., 2018).

The marginal PDFs estimated by each filter, overall, do not seem to be extremely different when comparing them within this particular Monte Carlo run. The AGMUKF estimates the PDF of the angular velocity as a bit more disperse, for the non rotating directions. Nonetheless, in both filters the truth stays within the 99.73% region for all states, even if it goes close to the border around the 1500s for the AGMUKF. Overall, both converge quite well by the end. Since this is a case tuned to work with the UKF version, it is to be expected that the AGMUKF reduces to almost a normal UKF.

Despite all this, the fact that each filter converges well once, for a specific initial realization of the state and a specific realization of the light curve, does not guarantee that the filter will converge with all combinations. To assess this, the results of the 100 Monte Carlo runs need to be analysed as a whole, and not individually. This is done with the Log-Likelihood (LL) of all the Monte Carlo runs (Schiemenz et al., 2020; Vittaldev et al., 2016)

$$LL = \sum_{i=1}^{MC} \left[\log \left(\sum_{j=1}^M w^j \mathcal{N}(x \mid \hat{x}^j, P^j) \right) \right]_i, \quad (5.58)$$

where the square brackets indicate each distinct Monte Carlo iteration i . In the UKF case, this reduces naturally to $M = 1$. In any circumstance, the higher the LL, the likelier it is that a filter estimates a consistent PDF of the state.

Figure 5.8 shows the LL for the UKF; while Figure 5.9, that of the AGMUKF. In both cases, the solid blue curve is the LL of the whole state, the orange dotted one is the LL applied only to the three attitude states, and the green dashed one is that of the angular velocity states.

These two plots do show a significant performance difference between the two filters. On the one hand, the LL of the UKF plummets as time goes by, which implies that the filter diverges in several of the runs. On the other hand, the LL of the AGMUKF steadily increases over time, indicating that it converges over all runs, consistently closing in around the true state as simulation time increases. This shows that the AGMUKF is clearly superior to the plain UKF in terms of stability, and that it is significantly likelier to estimate a consistent PDF.

The root cause for the superiority of the AGMUKF is likely the fact that it can model the PDF of the estimated measurement from the state with significantly higher accuracy. This is demonstrated by the results shown in Figure 5.10, which is the equivalent of Figure 5.2 but for the AGMUKF case. An additional curve is shown here—dash-dotted cyan—which represents the measurement PDF before the split. The solid orange one, then, is the measurement PDF that results after the split. When they coincide, it means that there is no need for splitting in that particular time step. By comparing Figure 5.2 against Figure 5.10, it is clear that the AGMUKF is able to model the measurement PDF, which is highly non Gaussian and frequently multimodal, with a higher degree of accuracy than the UKF.

Figure 5.11 shows the history of the number of kernels for the AGMUKF run in question. The solid blue curve indicates mixture size over time, while the orange \times marks appear whenever the mixture is refined, right after the split takes place. The first four fifths of the abscissa, up to the dash-dotted vertical line, is represented in a zoomed scale, to highlight the first 100s of simulation, when most of the splits happen. For the rest of the simulation time, except at the very end, the kernel population gradually decreases as the kernels approach each other in KL-distance sense, which triggers merges once in a while.

Note, then, that even if at peak moments there are up to 100 kernels active, for the most part of the run the mixture is composed by less than 20. This demonstrates the efficiency of the AGMUKF architecture, which refines the mixture only when necessary, coarsening it as soon as possible to reduce the computational burden. Without parallelization, under a constant measurement frame rate, the AGMUKF uses approximately c times the computational resources of the UKF, where

$$c = \frac{1}{t_{N_t}} \int_0^{t_{N_t}} M(t) dt, \quad (5.59)$$

plus the additional time it takes to split and merge kernels. For the particular run presented here, this is $c \sim 11.7$. In particular, the computational needs of the merge step grow quadratically with M , since it needs to compute the relative distance between all kernels. Fortunately, both the merge step and all the kernel-by-kernel steps (*e.g.* propagation) can be parallelized, since the operations done on one kernel do not depend on the other kernels for

5. Attitude Determination

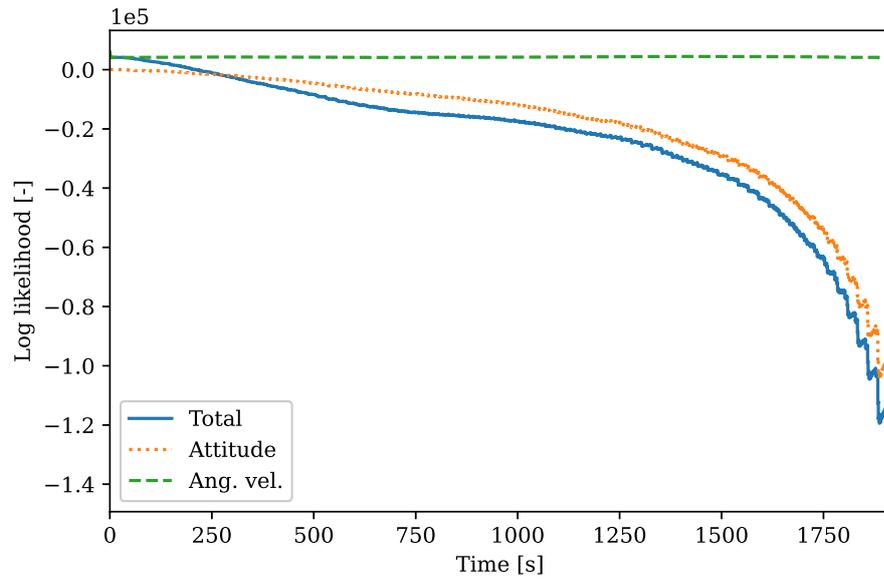


Figure 5.8.: Log-likelihood over 100 MC UKF runs for the Wetterer et al. (2009) scenario (Vallverdú Cabrera et al., 2023)

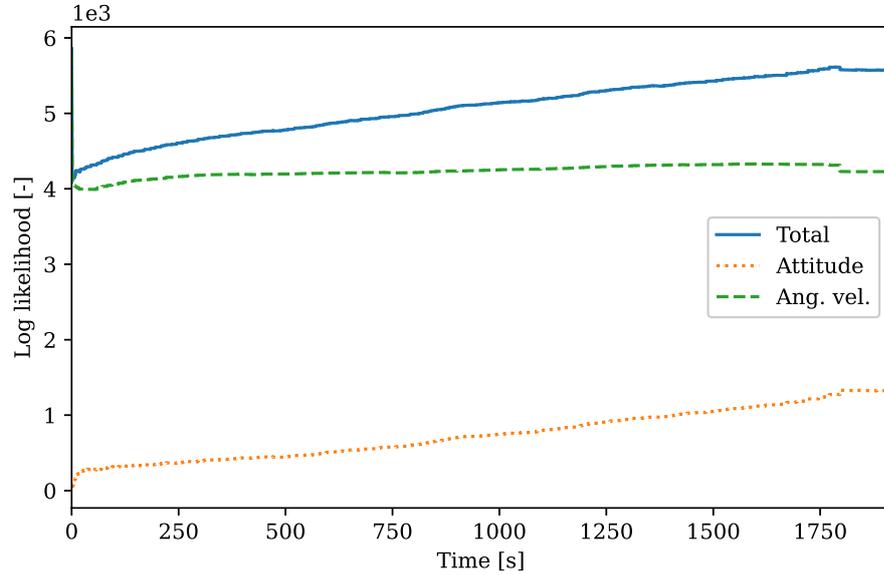


Figure 5.9.: Log-likelihood over 100 MC AGMUKF runs for the Wetterer et al. (2009) scenario (Vallverdú Cabrera et al., 2023)

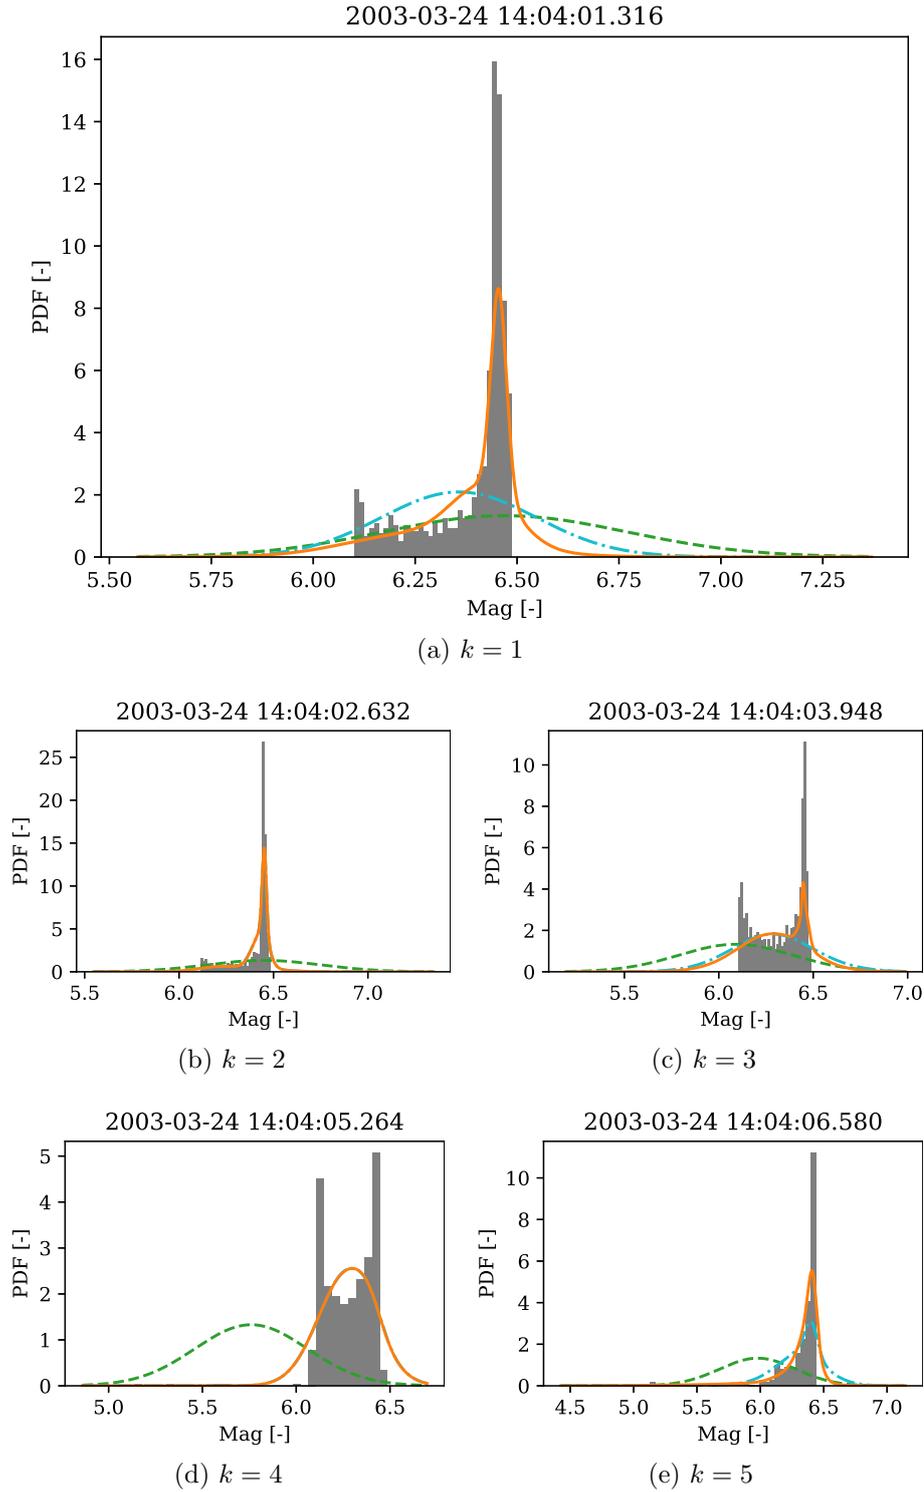


Figure 5.10.: Probability density function of \hat{z}_k for the first five measurements in the Wetterer et al. (2009) scenario when using the AGMUKF Legend: PDF of the observation (dashed green line); PDF of the measurement estimated from the state using UT before the split (dash-dotted cyan line), UT after the split (solid orange line) or Monte Carlo (gray histogram). Dates in the titles expressed in ISO format: YYYY-MM-DD hh:mm:ss.fff (f stands for millisecond) (Vallverdú Cabrera et al., 2023)

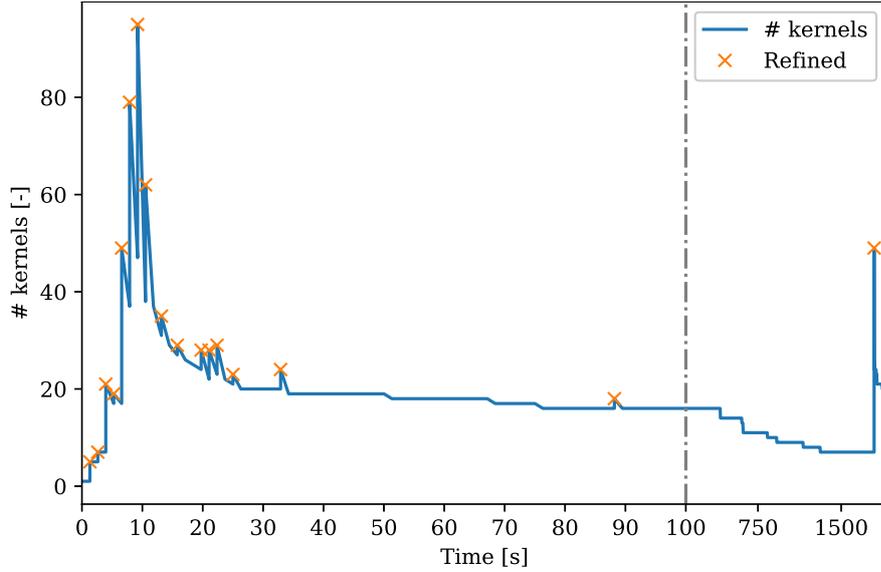


Figure 5.11.: Kernel number history for one AGMUKF run of the Wetterer et al. (2009) scenario (Vallverdú Cabrera et al., 2023)

most of the algorithm. Additionally, for each kernel, the UT propagation can be parallelized, too, since each sigma point can be processed independently from the others. Thus, overall, the execution time of both the UKF and the AGMUKF can be approximately divided by the number of available threads.

5.4.2. Highly non-linear case

The other test result studied in this chapter is that of a highly non-linear case, where the initial uncertainty of the state is greater than the previous case. Keeping on with the comparison against existing literature of AD using different techniques, this subsection studies a scenario similar to the one presented in Linares et al. (2014a), Example 2. This scenario has been chosen for the similarity of the models used therein w.r.t. to the dynamic and measurement models used here, and because Linares et al. (2014a) use the PF for AD, the base of the current state-of-the-art, and main competitor to the AGMUKF.

The main modification of the scenario used here w.r.t. Linares et al. (2014a) is that a different shape is used, since they do not provide information on the reflective properties of the surfaces of their test shape. To follow with the example from the previous section, the shape used is a 100-faceted cylinder approximation with 9 m length and 3 m diameter, with conic caps 1.5 m long on either end, like the one in Figure 5.4. The body x axis is parallel to the axial direction of the cylinder, while y and z lie on the orthogonal plane, forming a right-handed frame. The same reflective properties as in the previous section have been used. Regarding the filter, the same UKF- and AGMUKF-specific tunable parameters are kept from the previous section, too.

In the Monte Carlo run presented here, the filter is initialized with an initial diagonal covariance, with the δp components being 0.3932 and the ω ones, $2.350 \cdot 10^{-4} \text{rad}^2/\text{s}^2$, as in

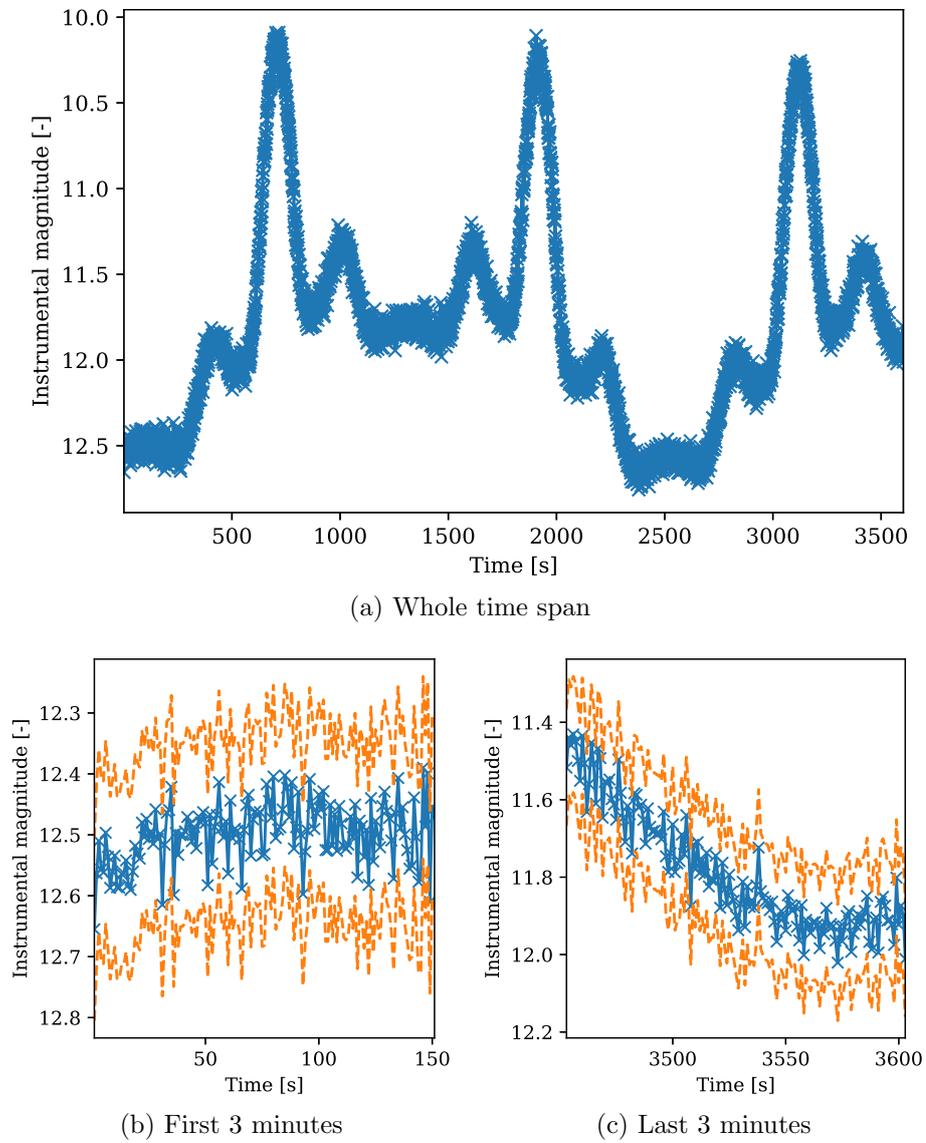


Figure 5.12.: One noisy light curve of the Linares et al. (2014a) scenario. Blue crosses are brightness measurements, while the blue line is a visual guide. The orange dashed line is the $3\text{-}\sigma$ envelope of uncertainty around the light curve (Vallverdú Cabrera et al., 2023)

5. Attitude Determination

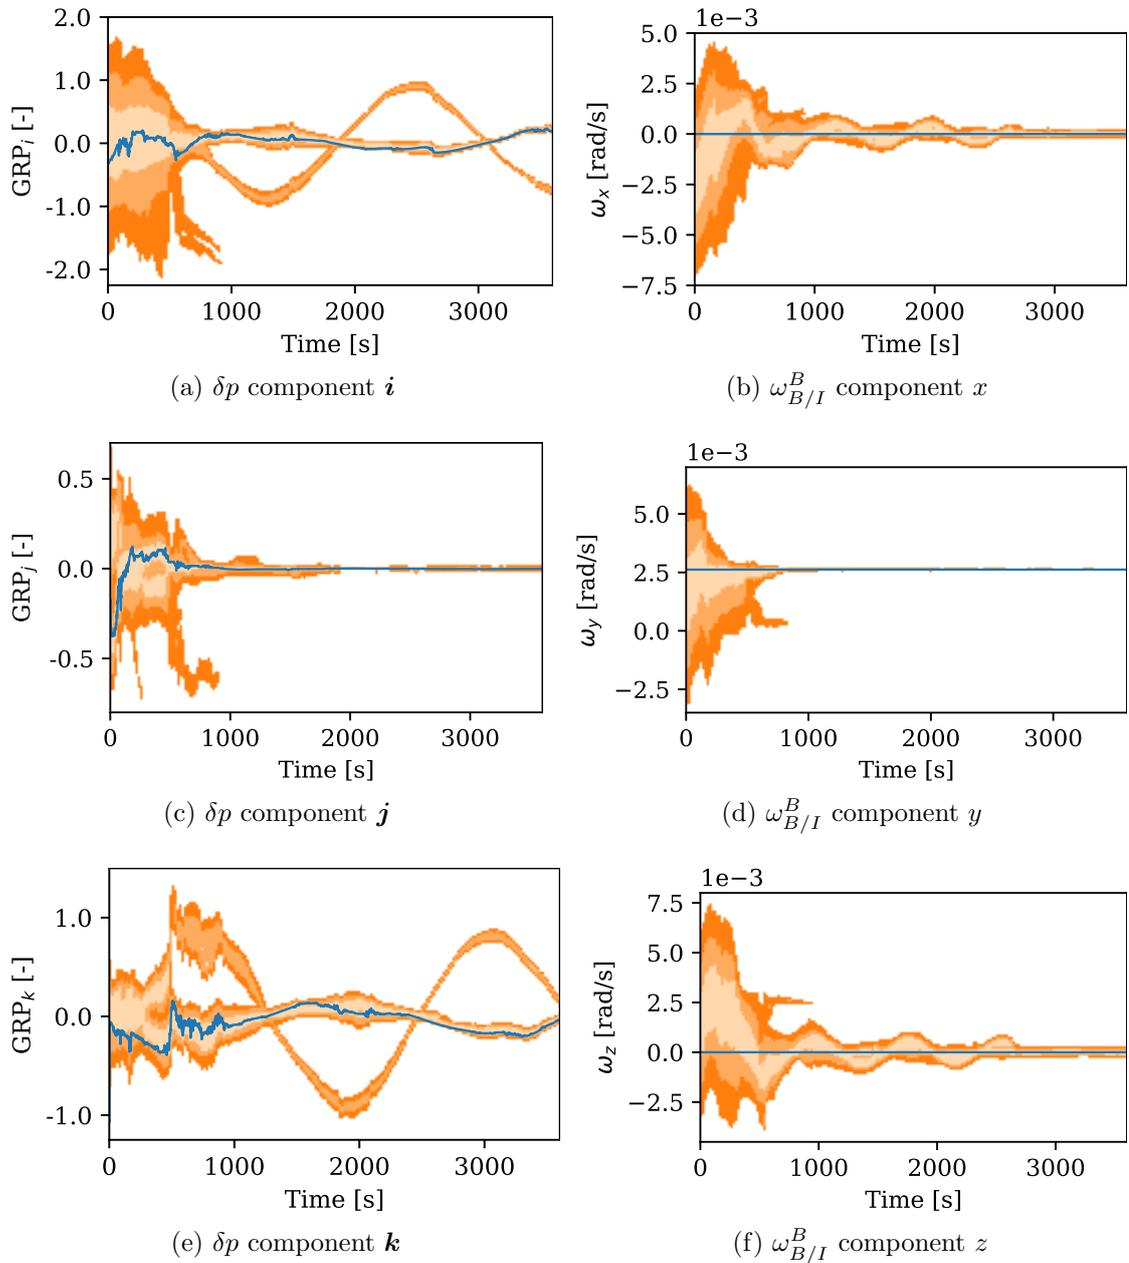


Figure 5.13.: One AGMUKF run on the Linares et al. (2014a) scenario. Legend: truth (blue line); 99.73% envelope of the estimated PDF (dark orange region), 95.45% envelope (middle orange region), 68.27% envelope (pale orange region) (Vallverdú Cabrera et al., 2023)

Linares et al. (2014a). The initial state is

$$\tilde{\mathbf{q}}_0 = 0.495 - 0.710\mathbf{i} + 0.293\mathbf{j} - 0.407\mathbf{k} \quad (5.60)$$

and

$$\hat{\omega}_0 = [-2.50 \cdot 10^{-3} \quad 1.50 \cdot 10^{-3} \quad 1.31 \cdot 10^{-3}]^T \frac{\text{rad}}{\text{s}}. \quad (5.61)$$

This results in the initial $\delta\hat{p}_0 = 0$ being at an Euclidean distance of ~ 1.13 from the true δp_0 , in GRP space. Figure 5.12 shows the measured light curve, where the zoomed plots include the bounds of the 3σ , or 99.73% probability region, with $R = 0.05^2$. The process noise used is

$$Q_k = \text{diag}([10^{-12} \quad 10^{-12} \quad 10^{-12} \quad 10^{-14} \quad 10^{-14} \quad 10^{-14}]). \quad (5.62)$$

Again, net torque is 0.

Figure 5.13 shows the evolution of the estimated pdf of this particular filter run. In contrast to Figures 5.6 and 5.7, here the 1σ ($P_r = 68.27\%$), 2σ ($P_r = 95.45\%$) and 3σ ($P_r = 99.73\%$) probability regions are all represented in solid orange, from lighter to darker tone, instead of using bounds.

Within this run, the AGMUKF is able to constrain the state pdf well around the truth, as shown in Figure 5.13. However, due to the high non-linearity of the problem, coupled with the wide initial uncertainty in attitude, which almost gives no constrain to the initial state, the filter cannot fully decide among two different modes. These two modes have the same angular rate and \mathbf{j} GRP component, but show a constant offset on the \mathbf{i} and \mathbf{k} GRP states (Figures 5.13a and 5.13e). This means that there is a non-negligible probability that the measured light curve could be the result of either alternative attitude states, as far as the filter is concerned.

This does not imply that the filter has failed to converge. In fact, the true state stays within the 1σ (palest orange) region through most of the simulation time. Thus, the mode that differs more from the truth is the one with a significantly lower share of the total probability. However, most importantly, the log likelihood of this single run increases throughout the simulation (Figure 5.14), meaning that the filter tends to converge around the truth as new measurements come in.

This particular Monte Carlo run has been chosen in this work precisely as a clear example of convergence around a multi-modal solution, a feature that is enabled by the use of Gaussian Mixtures.

This stability, however, is not preserved in a 100-samples Monte Carlo study, as in the previous scenario: while some runs have converged around a single attitude mode, or with multi-modal PDF, such as the one presented here, some other runs diverge, making the total LL plummet. The specific ratio of failed runs for this simulation is 14%.

Since the initial greater covariance results in higher NLIs, both in the dynamic and the measurement models, but especially in the latter (see Figure 5.15), the splitting step recommends significantly more splits, compared to the Wetterer et al. (2009) scenario. This is clear in the plot from Figure 5.16, where the first 700s are zoomed in to show in more detail the region where most splits happen. The total number of kernels in a given

5. Attitude Determination

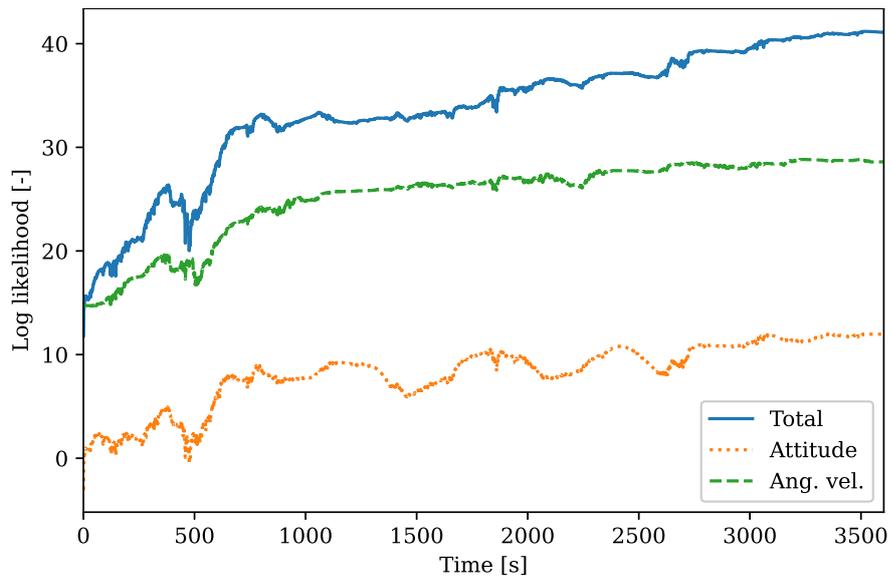


Figure 5.14.: Log-likelihood over a single MC AGMUKF run for the Linares et al. (2014a) scenario (Vallverdú Cabrera et al., 2023)

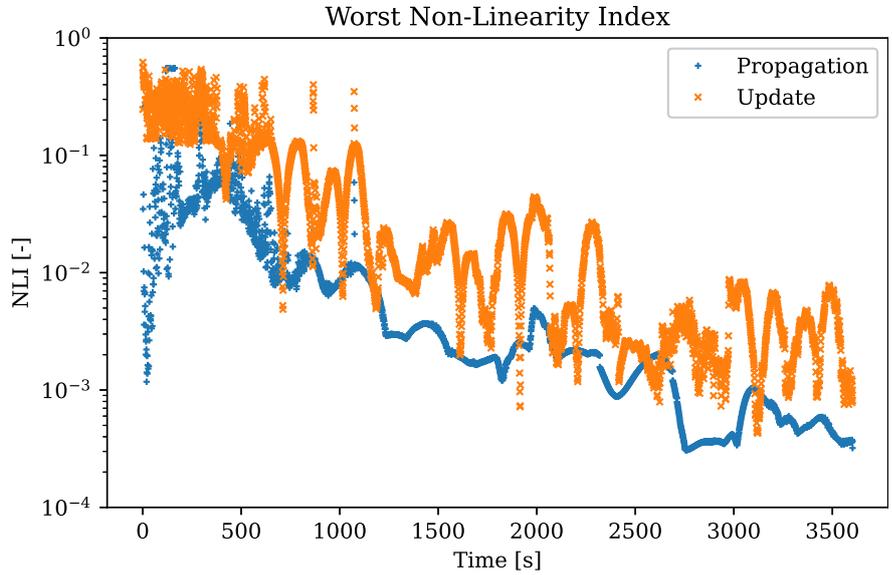


Figure 5.15.: Non-linearity index of one AGMUKF run for the Linares et al. (2014a) scenario (Vallverdú Cabrera et al., 2023)

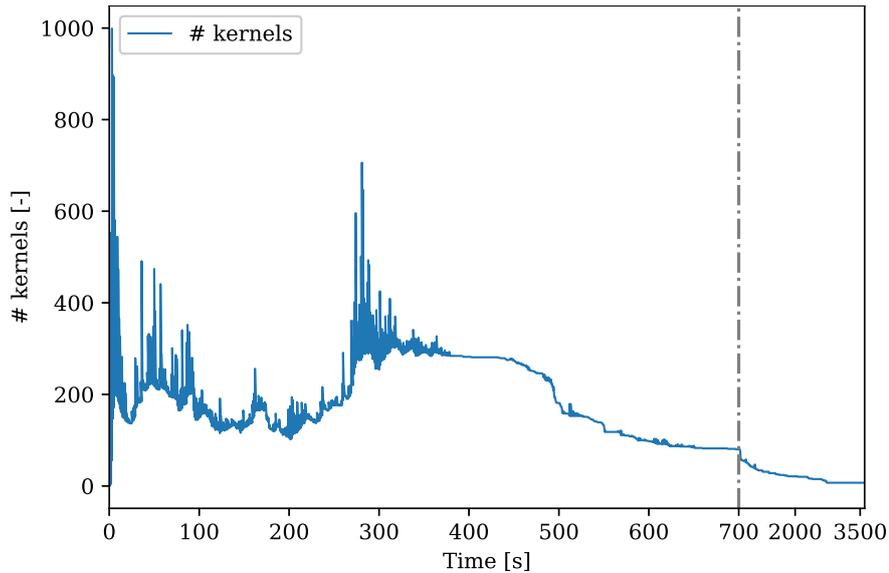


Figure 5.16.: Kernel number history for one AGMUKF run of the Linares et al. (2014a) scenario (Vallverdú Cabrera et al., 2023)

time reaches the hard limit of 999 at the very beginning, suggesting that a higher bound might improve the results for this scenario in a general Monte Carlo sense.

Given that the 100-runs Monte Carlo study does not result in LL convergence, it is clear that several improvements could be attempted to stabilize the AGMUKF for this kind of higher initial uncertainty scenario. Some viable options would be to:

- increase the maximum allowed size of the population, as suggested above;
- make the splitting algorithm (Equation 5.31) more aggressive by making it start splitting earlier (higher b) and at a faster rate (higher a)—this would result in a finer modelling of the measurement PDF;
- since the NLI is significantly high for the propagation step (Figure 5.15), incorporate the split/merge strategy around it, too, as in Schiemenz et al. (2020) for the OD case; in this case, doing multiple prediction steps between measurements could help conserve a more faithful estimate of the state PDF;
- include better process noise tuning/modelling, as in Crassidis et al. (2003) and Coder et al. (2017);
- make the filter over-estimate the initial covariance, to ensure outlier initial conditions are still covered;
- initialize the filter with a GM of more than one kernel, with the goal to off-load the splitting algorithm with the (currently) inherent limit of 39 splits per direction per update of the Vittaldev et al. (2016) library; this could be combined with the Expectation-Maximization algorithm, as in Yun et al. (2020), to fit the initial mixture to some arbitrary initial PDF, such as the uniform distribution (Shuster, 2003);
- the 39-splits limit might also be palliated by recursively splitting plus consecutive

5. Attitude Determination

merging—although this is not the optimal split solution, it might still benefit the filter by reducing the non-linearity of each kernel.

Computationally speaking, applying Equation 5.59 to Figure 5.16 yields an average number of kernels per iteration of roughly 70. Each kernel consists of an UKF composed by 13 particles (twice the number of error-states plus one). Thus, the results here have been obtained with an average of $13 \times 70 = 910$ particles. In contrast, the results from Linares et al. (2014a, Fig. 3) use roughly 20 000 particles, one order of magnitude higher. Further tests are needed to assess how the AGMUKF and the PF perform in exactly identical scenarios, and how much of the computational overhead is taken by the AGMUKF splitting/merging part, versus the resampling and regularization tasks required by the PF. Nonetheless, these results show significant promise for the AGMUKF, and its potential efficiency benefits w.r.t. the PF.

5.5. Conclusion

This chapter has introduced the process of non-linear sequential estimation aimed at the Attitude Determination (AD) of Resident Space Objects (RSOs) using light curves, when the shape and surface reflectance model is known. The dynamic and measurement models of the problem have been presented, and their significantly non-linear character has been emphasized.

An existing Non-Linearity Index (NLI) based on the Unscented Transform (UT) has been modified to allow comparison between different transformation models. This new NLI has been used within a Monte Carlo study of the measurement model to assess the inherent non-linearity of the problem. Results have shown that the measurement model is relatively less linear than the dynamic one. This has motivated the design of the Adaptive Gaussian Mixtures Unscented Kalman Filter (AGMUKF) for Attitude Determination (AD), based on a similar implementation of the AGMUKF for Orbit Determination (OD).

First, the convergence performance and stability of this new filter has been assessed against a scenario that had previously been proven to converge under the classical Unscented Kalman Filter (UKF) for AD. The results of a Monte Carlo study have shown that, although the convergence rate can be similar for particular runs, the AGMUKF is significantly likelier to always converge than the UKF.

Second, the AGMUKF has been tested against a more demanding scenario, for which the Particle Filter (PF), its main rival, had been previously proven to work. The AGMUKF has been shown to converge for an example run with significant initial uncertainty, successfully capturing the multimodal character of the state probability density function. Several suggestions have been given to potentially improve its stability against varied initial and noise conditions.

The results of this chapter have introduced the novel AGMUKF for AD as potential alternative to PF-derived techniques, since it has been shown to converge in scenarios where the traditional PF does, too. In contrast to the PF, the AGMUKF manages the number of kernels (particles in PF lingo) more efficiently, since it can concurrently assess the non-linearity present in the model. Furthermore, the AGMUKF avoids the PF

degeneracy problem, as well as its dependence on regularization. Finally, the AGMUKF naturally avoids the need to marginalize the PF when a part of the state evolves linearly through the involved models, as in Coder et al. (2017), because the NLI will naturally stay 0 in these directions.

The one caveat of the AGMUKF w.r.t. the PF is that the computational complexity of the merging step increases quadratically with the number of kernels. Significant improvement would be achieved by devising more efficient merging criteria.

Further work should include the application of the AGMUKF to other dynamic and measurement models, which *e.g.* target agile objects (Holzinger et al., 2012), to include shape- and reflectance-model uncertainties (Coder et al., 2017) or to fuse attitude and orbit estimation (Yun et al., 2020), or attitude and shape (Linares et al., 2018). Finally, the AGMUKF for AD should be tested against real-world data, as in Coder et al. (2018).

6. Use Cases

So far, this thesis has researched light curve simulation, plus two state-of-the-art Object Characterization (OC) methods, one for shape and one for attitude characterization of RSOs. This chapter explores some use cases for these contributions. Thus, it serves as a wrap-up of this dissertation, and lays out the ground for further development, too.

In other words, this chapter tackles the last research question: *Which SST use cases does direct Object Characterization enable?*

This chapter is divided in three sections. First, Section 6.1 summarizes the different use-case-enabling OC pipelines researched within this thesis. Then, Section 6.2 demonstrates a simple use case of using frequency analysis of light curves to deduce properties of the status of a RSO. Next, Section 6.3 conceptualizes a satellite diagnosis tool that exploits the main shape and attitude determination contributions of this thesis. Finally, Section 6.4 reviews other use cases that could benefit from exploiting OC techniques.

6.1. Object Characterization Pipeline

First, this section summarizes the OC pipelines studied within this thesis as a whole, so that the rest of the chapter can focus on which OC-derived use cases are enabled as a consequence.

Figure 6.1 shows a block diagram of said OC pipelines. The undulated blocks to the left of the vertical dashed line represent the input data used by each pipeline¹. To the right of the line, the different OC pipelines are represented by the interconnected rectangular blocks. On this side, undulated blocks represent the outputs of each pipeline.

The blocks *Frequency analysis* and *Epoch method* that form the middle row represent two shape-independent Attitude Determination (AD) methods presented in Vallverdú Cabrera et al. (2021a):

- The *Frequency analysis* block represents all the Fourier-related methods used to extract dominant frequencies from light curves and generate phase plots. These capabilities are demonstrated in the following section.
- The *Epoch method* block is the shape-independent AD algorithm presented in Hall et al. (2006) and Hall et al. (2014) that bears this name. As explained in the Literature Review, it exploits the fact that the apparent period of a light curve changes over time due to changes of the Sun-object-observer triangle relative to the spin axis of a

¹Note that common inputs such as the orbit of the RSO or the properties of the observer are not shown, since they apply everywhere.

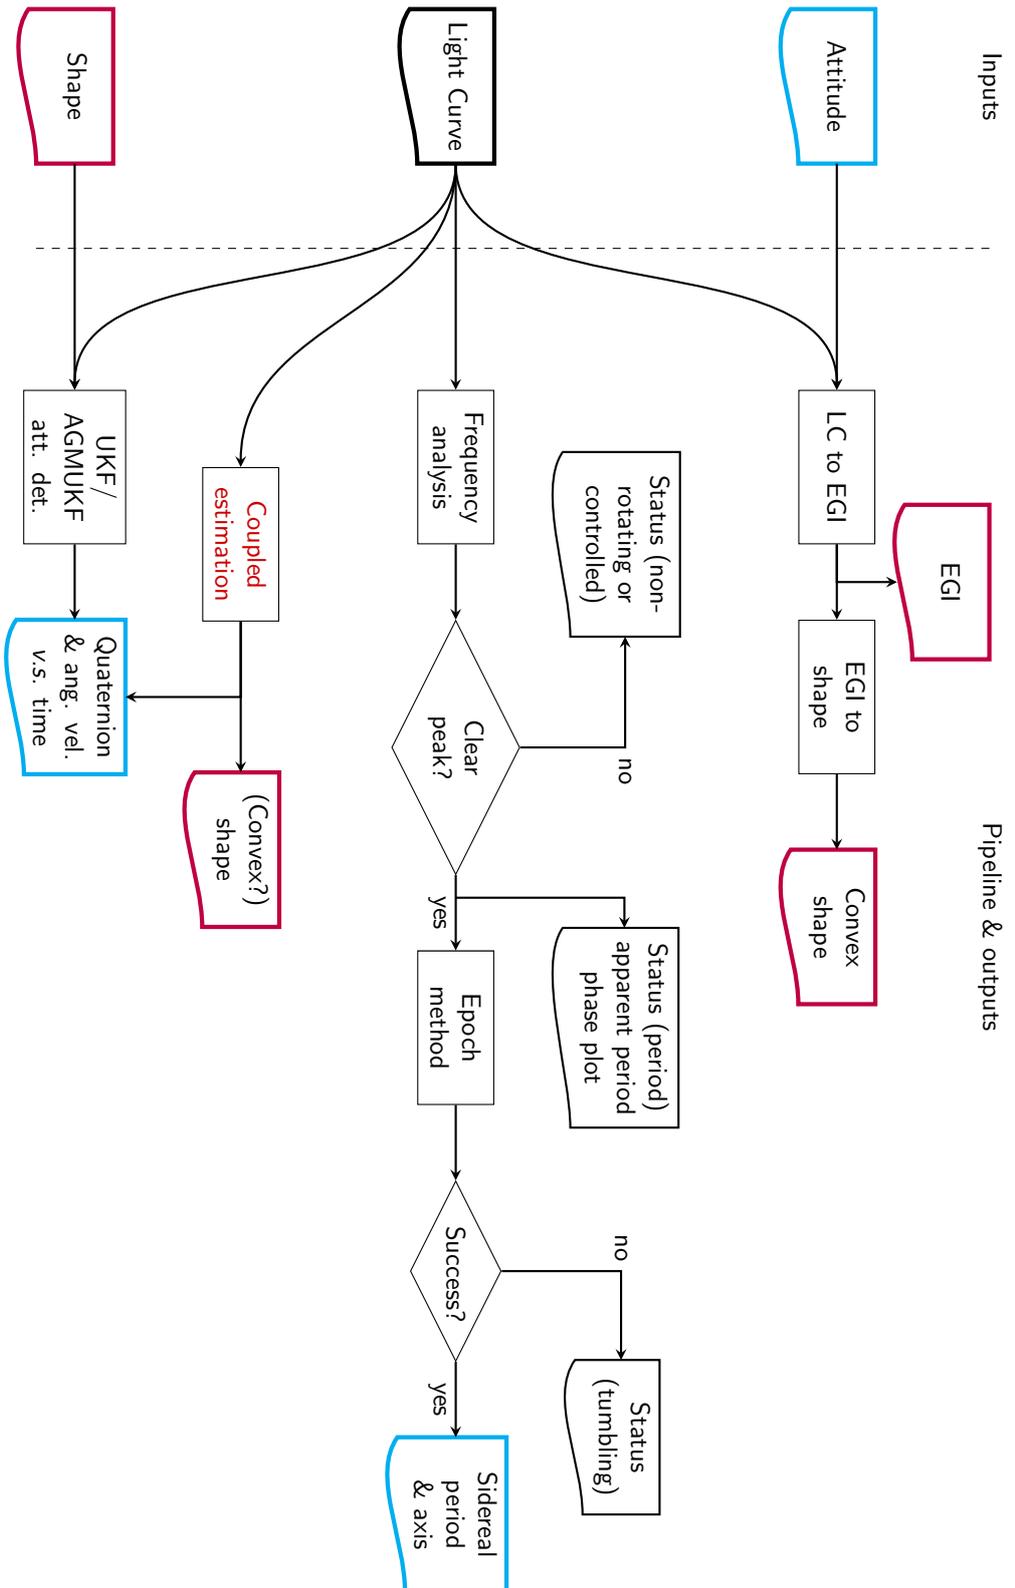


Figure 6.1.: Flow-chart of the OC capabilities analysed in this thesis

spin-stabilized object. This allows the method, under strongly restricting observability conditions, to recover the inertial period and direction of the spin axis.

Refer to Vallverdú Cabrera et al. (2021a) for a detailed development of this particular pipeline.

Next, the blocks connected on the top row, *LC to EGI* and *EGI to shape*, represent the light curve to EGI plus Minkowski minimization (*EGI to shape*) from Chapter 4. Meanwhile, the bottom block *UKF/AGMUKF att. det.* is the algorithm from Chapter 5. Finally, the *Coupled estimation* block symbolizes the simultaneous characterization of attitude and shape from a light curve. This is not part of the scope of this thesis, hence the red font, but it has been included in the diagram in the role of further work.

6.2. Frequency Analysis

This first use case is adapted from Vallverdú Cabrera et al. (2021a, Section 7.1), the first paper published within the scope of this doctoral research. As such, parts of Vallverdú Cabrera et al. (2021a) are included in this section. It intends to show one simple use case in which real observations are used to assess the rotational status of the ISO satellite (Cospar 1995-062A). It is included here to illustrate a simple, yet effective use case for OC, validated against true measurements.

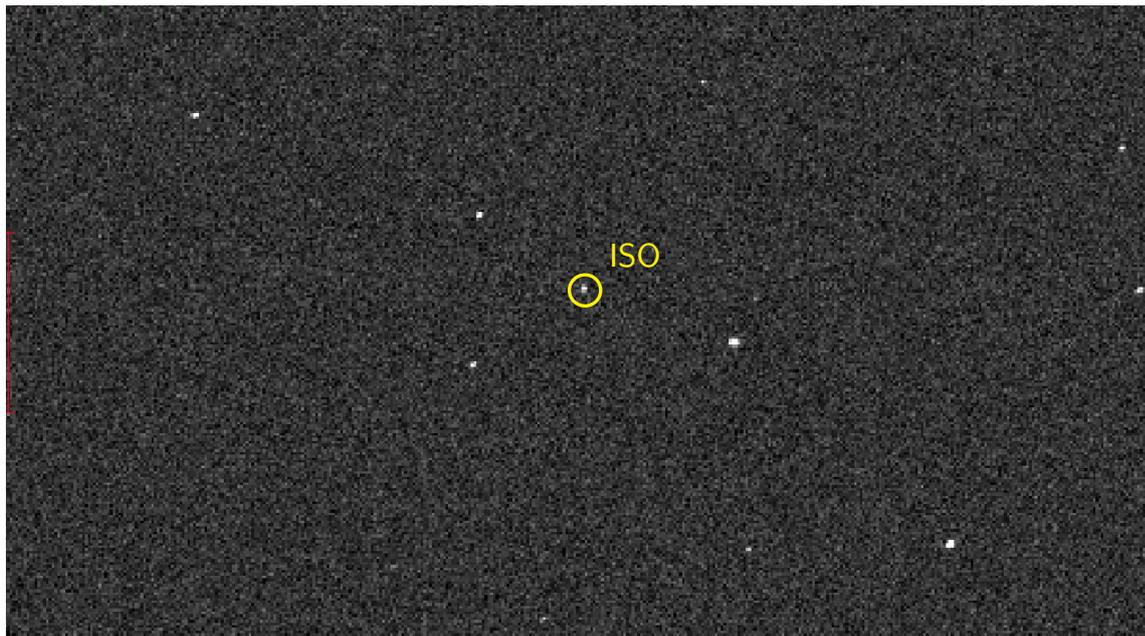


Figure 6.2.: Image of the ISO satellite taken with ART.

First, Figure 6.2 is a zoom on one of the images used to build the light curve used for this analysis. It is part of a series of images obtained by ART (see Appendix A.2) on the night of 2019-10-06. Each image was exposed for 0.2s, and the pixels were binned on two by two squares. In the example image, the ISO satellite is the highlighted dot.

6. Use Cases

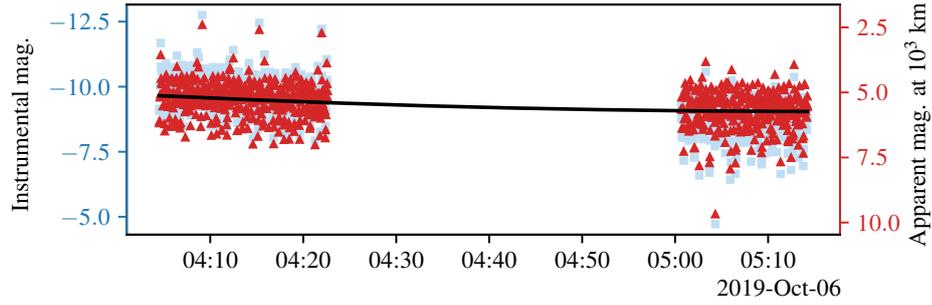


Figure 6.3.: Light curve of the ISO satellite measured on the morning of 2019-10-06 with ART (Vallverdú Cabrera et al., 2021a)

As explained in Vallverdú Cabrera et al. (2021a), all the images are reduced with AstroImageJ (AIJ) (Collins et al., 2017) to obtain the light curve shown in Figure 6.3. The faded blue squares represent the instrumental magnitude (Equation 3.37), computed from the electron counts reported by AIJ. These shall be read on the left-side ordinates. The red triangles, read on the right-side ordinates, are the apparent magnitude as computed by AIJ—it uses differential photometry to extrapolate the apparent magnitude of the object of interest, based on the known apparent magnitude of surrounding stars Collins et al. (2017). The faint blue squares, read on the left-side ordinates, are the same apparent magnitude samples, but further de-trended with a constant range (observer to object distance) of 1000 km. The black line is a polynomial fit on the apparent light curve.

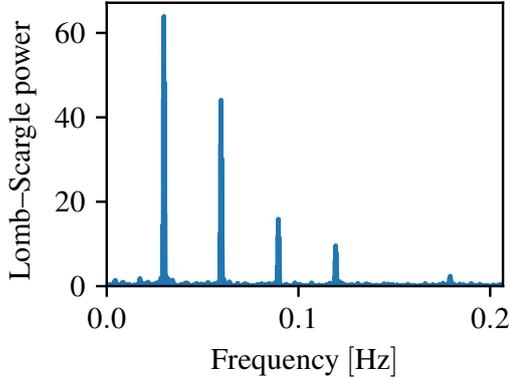


Figure 6.4.: Lomb–Scargle periodogram of the light curve from Figure 6.3, after detrending (Vallverdú Cabrera et al., 2021a)

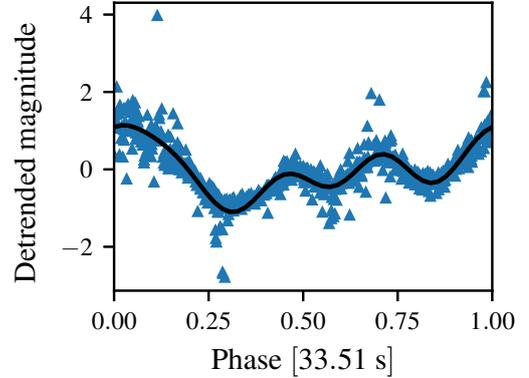


Figure 6.5.: Phase plot of the detrended light curve from Figure 6.3 (Vallverdú Cabrera et al., 2021a)

Figure 6.4 shows the Lomb–Scargle periodogram (VanderPlas, 2018) of this light curve. A clear series of harmonics can be observed, with a fundamental period of $T \simeq 33.51$ s. Given the non-symmetric shape of the ISO satellite, this is probably the actual apparent period at which it was rotating around itself on the night of observation—symmetries could mean that

the observed period is a divisor of the actual full rotation of the object shape from the point of view of the PAB (Linder et al., 2015). Furthermore, no other peaks that are not multiple of the fundamental frequency appear, which suggests that its rotation is mostly around a principal axis.

This allows the creation of a phase plot, shown in Figure 6.5. This is a plot of the light curve against the modulus of the relative time since the beginning of the light curve, divided by the apparent period, and normalized by the apparent period itself. Blue triangles are individual light curve points, while the black line is a fit of a Fourier series into the data.

This brief example shows one simple, but effective use case enabled by the most basic and well established OC applications (Linder et al., 2015). It has proven that frequency-domain analysis of actual light curves, without quantitative a priori knowledge of the object observed, can be used to extract qualitative assessments of the status of space object. In this case, its rotational status. Furthermore, curated data such as the power spectrum (or one of its generalizations such as the Lomb-Scargle periodogram) or a phase plot can be used as pre-processed input to ML-based classification algorithms, as suggested in Jia et al. (2018).

6.3. Diagnosis of an Operational Satellite

The use case presented in this section strives to verify the statement from Enomoto et al. (2016), which suggests that satellite health monitoring can benefit from the outputs of OC.

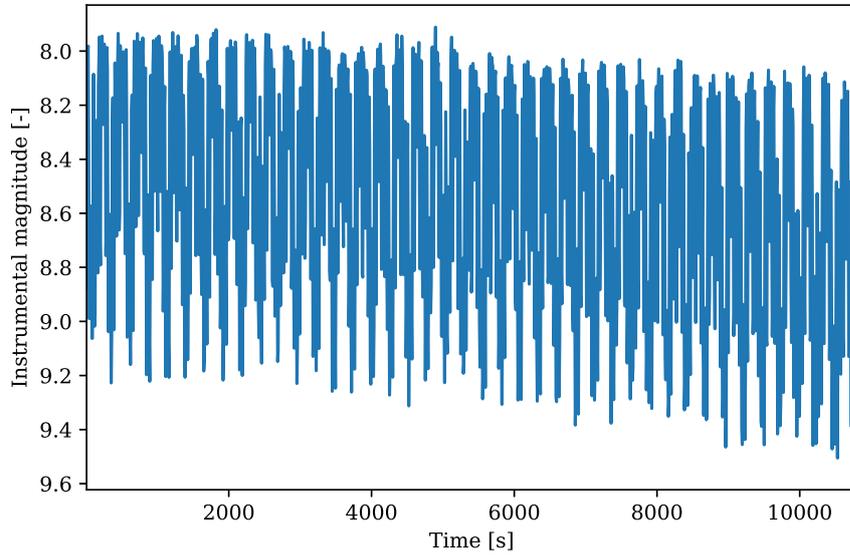
On the one hand, the EGI-based shape-determination approach, which this thesis has revisited and verified in several scenarios, requires knowledge of the RSO attitude. On the other hand, the AGMUKF for AD, which has been introduced and validated in this dissertation, needs the shape (and reflective properties) model that represents the RSO as an input. This makes these two OC methods complementary. Motivated by this, this section combines these two methods into a demonstrator for a diagnosis tool for operational satellites.

The use case assumes the scenario in which the attitude of the satellite is known via telemetry provided to the SST service provider by the operator entity. The design shape model, including the reflective properties of the satellite surface materials, are provided, too. The question to answer is: how can one detect whether the surface integrity of the satellite is intact?

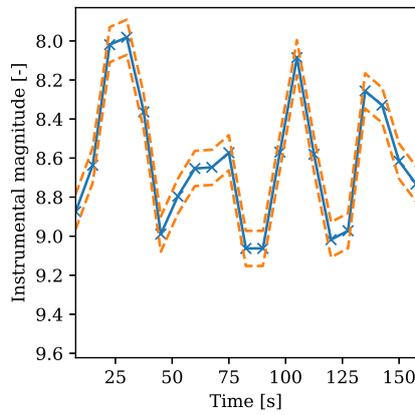
One solution is to exploit the Attitude Determination method developed in Chapter 5. The hypothesis is that, if the shape/reflective model of the satellite remains valid, the result of AD should return an attitude state PDF for which the real attitude has a significant likelihood. Otherwise, the likelihood of the real attitude state for the obtained PDF should be close to zero.

This section tries this approach via the following example. Observations from the Airbus Robotic Telescope (ART) are simulated, of a GEO object with the shape of the winged cube used in Chapter 4 (Figure 4.3d). The design reflectances correspond to Lambertian diffuse BRDF (Equation 3.16) with $\rho = 0.5$ on the body, and a specular Cook-Torrance BRDF (Equations 3.16 and 3.17) with $\mathfrak{d} = 0.25$, $\rho = 0.45$, $F_0 = 0.75$ and $\mathfrak{m} = 1/\sqrt{2}$ on the

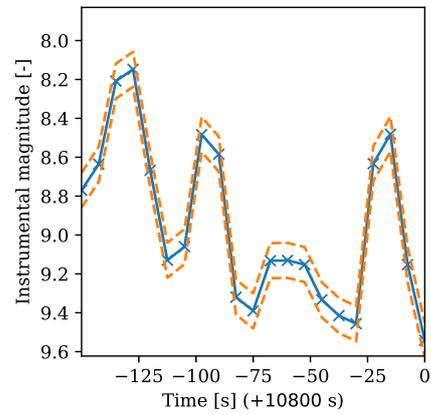
6. Use Cases



(a) Whole time span



(b) First 3 minutes.



(c) Last 3 minutes.

Figure 6.6.: One noisy light curve of the diagnosis scenario. Blue crosses are brightness measurements, while the blue line is a visual guide. The orange dashed line is the 3- σ envelope of uncertainty around the light curve

wings. The simulation starts on 2022-07-25 at 23:45:00 UTC (t_0 from here onward). At the time when the first measurement is taken, $t_0 + 7.5$ s, the object position and velocity GCRF vectors are

$$\mathbf{r}^I = [23871.3 \quad -34766.6 \quad -67.5]^T \text{ km}, \quad (6.1a)$$

$$\mathbf{v}^I = [2.534 \quad 1.740 \quad -0.005]^T \text{ km s}^{-1}. \quad (6.1b)$$

At that time, the attitude state is

$$\mathbf{q}_B^I = 0.7015 - 0.4425\mathbf{i} - 0.2896\mathbf{j} + 0.4777\mathbf{k}, \quad (6.2a)$$

$$\boldsymbol{\omega}_{B/I}^I = [-3.947 \quad 0.711 \quad 2.865]^T \text{ }^\circ \text{ s}^{-1}. \quad (6.2b)$$

The exposure time is 5 s, and the time between consecutive measurements is 7.5 s. The attitude during the 3 h long simulation follows torque-free conservation of angular momentum. It is assumed that the object has a symmetric² inertia tensor that is diagonal in body frame, $\text{diag} = [I_{xx} \quad I_{yy} \quad I_{zz}]$, with $I_{zz}/I_{xx} = 5.852 \cdot 10^{-2}$ and $I_{xx} = I_{yy}$. This gives rise to a precession movement in which the angular momentum of the object lies at 87.6° from the z -body axis, and rotates around it $2.698^\circ \text{ s}^{-1}$. In this use case, this is the attitude reported by the satellite telemetry after *e.g.* a suspected collision with a piece of space debris.

Let us imagine that the operators of the satellite want to determine whether the integrity of the surface of the satellite is intact. One could do so via the AGMUKF for AD algorithm presented in Chapter 5. To demonstrate, let us assume that the operator can provide the attitude state of the satellite at a Gaussian $1\text{-}\sigma$ accuracy of approximately 1° in any of the Euler angles, while the attitude rate error is $0.006^\circ \text{ s}^{-1}$. It is assumed that the total sensor noise amounts to a Gaussian covariance of $R = 0.03^2$ magnitudes squared. Figure 6.6 shows the simulated light curve, with the usual zoom at the beginning and the end, to show the $3\text{-}\sigma$ envelope.

The experiment goes as follows. If the result of AD using the known shape and reflective model matches the known attitude, it is probable that the satellite integrity is fine. If this is not the case, this could be an indicator of degraded integrity, because the shape/reflective model from design is not applicable: some properties have changed. The likelihood metric can be used for this purpose (Equation 5.58). The filter is tuned as in Section 5.4, but with diagonal process noise of $\text{diag}(Q) = [4 \cdot 10^{-8} \quad 4 \cdot 10^{-8} \quad 4 \cdot 10^{-8} \quad 0 \quad 10^{-12} \quad 0]$, in the corresponding Interational System (SI) units.

On the one hand, Figure 6.7a shows the likelihood of one run of the AGMUKF, for the case when the design shape and reflectance model are true. The likelihood steadily increases. This is expected of a filter with a correct physical model, where each new piece of information (each measurement) contributes to keep the hidden state knowledge (the PDF) concentrated narrowly around the truth. On the other hand, Figure 6.7b shows the same metric, but for the case when one of the two solar panels has been degraded, in such a way that its specular reflectance F_0 has changed from the design value of 0.75 to the degraded value of 0.15. The filter still thinks it has the old value, though. This is the reason why this filter run, where

²This assumption may not be exactly realistic with the shape from Figure 4.3d, but it is merely used here as an artefact to justify a tumbling motion of the object. Without loss of generality, a more detailed inertia tensor could be used.

6. Use Cases

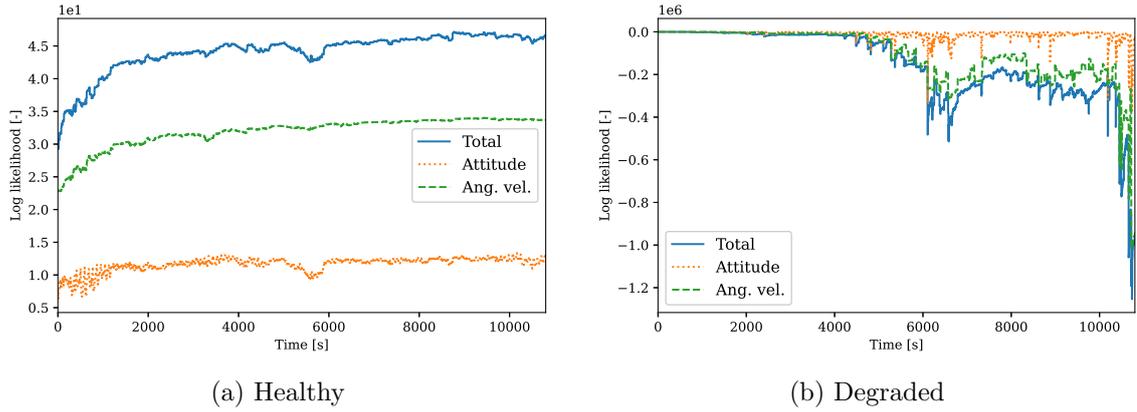


Figure 6.7.: Log-likelihood over a single AGMUKF run for the diagnosis scenario

any other parameters are equal to the case of Figure 6.7a, fails to estimate the true attitude: the likelihood stays at almost null values for most of the duration of the light curve.

The AGMUKF for AD algorithm used here is rather new. As commented in Chapter 5, there is ample room for improvement, so that it becomes faster and more robust. This means that in its current state, a convergence failure of the filter in a real scenario (where it is unknown whether the actual shape model matches the design one or not), might not be caused by a degradation of the object surface, but by the instability of the method.

Despite this, this small example still accomplishes the goal to demonstrate, in a simplified and controlled scenario, how OC algorithms could potentially be used for health diagnosis of operational satellites. This pipeline could be further automatized by state-of-the art ML-based analysis, applied either on the results of OC, such as attitude, or on the light curves themselves. This is *e.g.* what Kucharski et al. (2021) do to detect attitude anomalies in mega constellations of satellites.

Finally, this exercise could have been done using the shape determination algorithm, instead. One would then compare whether the convex-equivalent shape obtained from real measurements is equal to the one obtained from simulated measurements (where the design shape is used to simulate them). However, the AGMUKF was chosen to demonstrate the diagnosis use cases because it can handle non-convex object without having to resort to convex-equivalent shapes.

6.4. Other Use Cases

Other use cases of Object Characterization (OC) have been suggested in the Literature Review. In particular, OC for ADR and OC for Orbit Determination (OD).

For the ADR case, the most interesting usage of OC is Attitude Determination (AD)—it has been shown in Chapter 2 that knowing the attitude state of an ADR target a priori is key to the mission success. This is, in fact, a direct application of the AGMUKF algorithm developed in Chapter 5. The algorithm needs to undergo a more in-depth and systematic tuning process, and can diverge in high-initial-uncertainty cases. Nonetheless, for

demonstration purposes, Section 5.4 gives a perfectly valid example of the ADR use case enabled by OC.

Another use case would be to fuse OD with the rest of OC. This would enable enhanced accuracy Collision Avoidance (CA) and Re-entry Analysis (RE), two of the typical SST use cases, as pointed out in Chapter 2. Partial solutions to this use case have been implemented in the literature (Section 2.2.5). However, observability conditions are weak at best. In fact, most often, the full set of characteristics that would enable enhanced orbital models (attitude, shape and reflective properties to account for precise perturbation forces and torques) are unobservable (Section 2.3). This observability weakness might be improved using multiple simultaneous observers from different locations (different observation geometries contribute different coverage of the hidden state), mixing sensor types (*e.g.* Radar and optical) and/or enhanced light curves (polarized or spectral ones). Furthermore, OD enhanced by OC requires the merging and coupling of the dynamic and measurement models of the orbit, attitude, shape and reflective properties estimation problems.

Providing tangible application examples of this use case falls outside the scope of this thesis. Nonetheless, this work has advanced the status of some of the individual components that enable it.

7. Conclusion and Outlook

This is the final chapter of this thesis. It synthesizes the conclusions of this research work (Section 7.1) and lays out recommendations for future work (Section 7.2). This thesis closes with a final word in Section 7.3.

7.1. Summary

Motivated by the need to mitigate the ever-growing threat that overpopulation of Earth orbits poses to humanity, the aim of this research was to advance the technologies used to characterize the Resident Space Objects (RSOs) that inhabit these orbits—the field known as Object Characterization (OC). In particular, it explored which characteristics can be estimated from light curves, how, and under which observability conditions. Additionally, it sought to determine which Space Surveillance and Tracking (SST) use cases are enabled by OC.

A thorough literature research on OC has been performed to fetch a preliminary answer to the research questions of this thesis. The outcome of this effort has led to the identification of the particular and specific goals that have driven the core research effort of this work.

First, within the tangential, but necessary field of light curve simulation, a new rendering algorithm optimized for objects with a few, big flat facets has been developed. This algorithm, dubbed Shadow Projection, has been proven to outperform the state-of-the-art ray tracing alternative when rendering simple polyhedrons.

Second, the shape determination algorithm that exploits the Extended Gaussian Image (EGI) parametrization of arbitrarily complex convex polyhedrons and Minkowski minimization has been explored. This work has computed the analytical expressions for the Hessian of the volume of a polyhedron with its supports. It strived to improve upon previous implementations of this shape recovery algorithm, which used finite differences to compute the Hessian. However, the accuracy of the algorithm has, at best, stayed the same. It has been hypothesized that the proposed analytical solution of the Hessian dismisses the corner cases when facets may appear/disappear upon small changes of support values—these rare cases would make high-density EGI tessellation cases more likely to fail.

Still on the shape determination topic, it has been shown that EGI tessellations with more cells than points in the light curve can be used even when the Gramian of the problem is rank-deficient, making the hidden state non-observable from the point of view of linear control theory. This is thanks to the non-negativity constraint imposed on the area values of the EGI. It has been observed that this property holds even in simulation scenarios with noisy light curves, depending on the shape of the object.

7. Conclusion and Outlook

The study of the EGI plus Minkowski minimization shape determination algorithm performed in this thesis has resulted in one final contribution: a technique to quantitatively assess how close two polyhedrons are from homotheticity. This method improves upon previous evaluation techniques by assessing the final result of the shape recovery pipeline (the polyhedron), instead of the intermediate one (the EGI).

Third, this research has explored the sequential estimation approach for Attitude Determination (AD) from light curves. Sequential estimation being the main choice for state-of-the-art AD, this dissertation has introduced the new Adaptive Gaussian Mixtures Unscented Kalman Filter (AGMUKF) for AD. This filter has been adapted from the Orbit Determination (OD) field from previous research. It models the Probability Density Function (PDF) of the attitude and angular velocity state of the object with a Gaussian Mixture (GM), and exploits recent Gaussian splitting libraries that optimize the management of Gaussian kernels. To adapt the AGMUKF concept to AD, a modified Non-Linearity Index (NLI) has been introduced, which has been used to assess and compare the non-linearity present in the attitude dynamic model and the light curve measurement models. The results obtained have motivated the design of an architecture that focuses on mitigating the measurement non-linearity.

Simulation results have shown that the AGMUKF for AD is superior to the Unscented Kalman Filter (UKF) classical alternative for sequential estimation of attitude from light curves. Furthermore, it has been shown that the algorithm has the potential to handle highly non-linear models proven to work under the state-of-the-art Particle Filter (PF), but at a lower computational cost. Although more work is needed to properly tune it for it to be truly comparable to the PF, the AGMUKF is shown to be a promising competitor in the field of sequential AD: it would fill the gap between the low-accuracy, low-cost Extended Kalman Filter (EKF)/UKF families, and the high-accuracy, high-cost PF ones.

Finally, it has been shown how the different OC algorithms introduced throughout this work, including some mentioned in the Literature Review, have the potential to become key pieces of the SST service pool. In particular, examples have been included on how 1) Fourier analysis can be used to obtain rotational properties of objects, and 2) to propose and demonstrate a simplified architecture that exploits OC algorithms to assess the health status of an operational satellite.

7.2. Outlook

Object Characterization (OC) as a field has a considerable history, especially rich during the last three decades. Nonetheless, new methods and technologies keep appearing continuously, and new insights are being published at SST-related conferences, as well as new issues of relevant journals. This indicates that there is still work to be done, and that there is still room for OC to become more mature as a field. This dissertation is no exception. Thus, this section outlines what short- and mid-term research efforts should focus on, to further advance the work carried out here. These suggestions are mainly motivated by the results obtained above.

With respect to the topic of light curve simulation, and rendering in particular, it is necessary to properly parallelize the Shadow Projection algorithm (Algorithm 1). An

implementation based on Graphics Processing Units (GPUs) could increase its computational performance.

Concerning shape determination and the EGI plus Minkowski minimization algorithm, it would be interesting to see how well the regularization schemes studied in this thesis, as well as exploiting the analytical version of the Hessian, combine with other extensions of the algorithm, such as the ones proposed by Fan et al. (2019) and Fan et al. (2021). Furthermore, the analytical version of the Hessian of the volume of a polyhedron presented in this thesis needs to be corrected to account for appearing and disappearing facets upon support change, as suggested in the discussion of the Shape Determination Chapter.

In relation to the attitude estimation algorithm (AGMUKF), this thesis has demonstrated its potential as a competitor of the PF. However, as a new method in the field of AD, it has ample room for improvement. As indicated in the discussion of the Attitude Determination Chapter, some upgrades could be performing a wider Monte Carlo study of different scenarios, to better tune the splitting and merging thresholds of the algorithm, or the governing criteria themselves; incorporating split/merging steps around the propagation step, to support its non-linearities; implement advanced process noise models from the literature of PF-based AD solutions; initialize the filter with more complex GMs; or consider recursive splits when the non-linearity is especially high. Furthermore, it is necessary to pit the AGMUKF against the PF with a formal Monte Carlo study, in the same way as has been done in this thesis with the UKF. This would narrow down the advantages and disadvantages of one relative to the other more accurately, and could help suggest further improvements to either.

The next natural continuation for either of the two OC algorithms studied in this thesis is to validate them with real measurements from an optical telescope. This would help consolidate them and detect further paths for improvement.

When looking into the shape and attitude (and reflective properties) determination problem as a whole, the research carried out in this work could be expanded by coupling the two estimation problems in one, as done in *e.g.* Linares et al. (2018). This approach has the problem of weak observability of the full light curve inversion problem highlighted in the literature. However, it is worth exploring whether using multiple observers, combined types of sensors, several time-spaced light curves, or polarimetry and spectral information could mitigate it. Still within the line of coupled estimations, another potential upgrade to the developments of this thesis would be to incorporate them into an OD scheme. This would benefit the accuracy of derived SST services such as Collision Avoidance (CA) or Re-entry Analysis (RE). This would entail coupling the shape and attitude dynamic and measurement models with orbital dynamics.

Next, regarding OC-enabled use cases, this thesis has shown only some simple demonstrations, which prove the value that estimating RSO characteristics has for the space sector. However, there is work to be done to refine the concept of *e.g.* the health monitoring use case, optimize it and integrate it into an end-to-end pipeline, such as the one presented in Pedone et al. (2021). This applies to other SST use cases, too, such as OC for Active Debris Removal (ADR) or, as mentioned in the previous paragraph, OC-enhanced OD for CA and RE.

Moving on to higher-level future work, another potential way forward is to shift the focus

7. Conclusion and Outlook

toward less physical characteristics, such as *e.g.* owner identification, inference of the capabilities of the spacecraft being observed, or manoeuvre characterization. This could be approached by using raw sensor data as input, by processing more direct characteristics (shape or attitude), or by combining both types of data.

Machine Learning (ML) literature is proving that neural networks can handle OC-related classification problems. Thus, it is worth considering these type of techniques as a tool for further research. It could either be used as a complement or as an alternative to the shape and attitude determination algorithms presented here—or to estimate other, more abstract characteristics (such as *e.g.* bus type).

7.3. Final Word

Overall, OC is a wide and complex topic. With its contribution, this thesis has taken the field one small step further. However, there is still need for plenty of scientific work to bring OC and SST to the point where they need to be. This is, so that they can support, at their fullest potential, the worldwide effort of keeping the near-Earth space environment clean and under control. After all, this was the main reason that motivated this doctoral research to begin with.

Bibliography

- Abay, R., S. Gehly, S. Balage, M. Brown, and R. Boyce (2018). “Maneuver Detection of Space Objects Using Generative Adversarial Networks”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Aglietti, G. S. (2020). “From Space Debris to NEO, Some of the Major Challenges for the Space Sector”. In: *Frontiers in Space Technologies* 1. ISSN: 2673-5075.
- Ahn, H. T. and M. Shashkov (2008). “Geometric Algorithms for 3D Interface Reconstruction”. In: *Proceedings of the 16th International Meshing Roundtable*. Ed. by M. L. Brewer and D. Marcum. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 405–422. ISBN: 978-3-540-75103-8.
- Aksoy, A. and M. Martelli (2002). “Mixed Partial Derivatives and Fubini’s Theorem”. In: *The College Mathematics Journal* 33.2 (2), pp. 126–130. ISSN: 0746-8342. DOI: 10.1080/07468342.2002.11921930.
- Ashikhmin, M. and P. Shirley (2000). “An Anisotropic Phong BRDF Model”. In: *Journal of Graphics Tools* 5.2, pp. 25–32. DOI: 10.1080/10867651.2000.10487522. eprint: <https://doi.org/10.1080/10867651.2000.10487522>.
- Badura, G. P., C. R. Valenta, and B. Gunter (2022). “Convolutional Neural Networks for Inference of Space Object Attitude Status”. In: *The Journal of the Astronautical Sciences*. ISSN: 2195-0571. DOI: 10.1007/s40295-022-00309-z.
- Barber, C. B., D. P. Dobkin, and H. Huhdanpaa (1996). “The Quickhull Algorithm for Convex Hulls”. In: *ACM Transactions on Mathematical Software* 22.4 (4), pp. 469–483. ISSN: 0098-3500. DOI: 10.1145/235815.235821.
- Beamer, D. K., U. A. Abeywickrema, and P. P. Banerjee (2018). “Statistical Analysis of Polarization Vectors for Target Identification”. In: *Optical Engineering* 57.5 (5), p. 054110. ISSN: 0091-3286, 1560-2303. DOI: 10.1117/1.OE.57.5.054110.
- Bennett, D., D. Allen, J. Dank, M. Gartley, and D. Tyler (2014). “SSA Modeling and Simulation with DIRSIG”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*.
- Bernard, A. J. and D. K. Geller (2018). “A Comparison of the Bootstrap Particle Filter and the Extended Kalman Filter and the Effects of Non-linear Measurements in Spacecraft Attitude Determination Using Light Curves”. In: *Spacegrant*, p. 13.
- Bradley, B. K. and P. Axelrad (2014). “Lightcurve Inversion for Shape Estimation of Geo Objects from Space-Based Sensors”. In: *Univ. of Colorado. International Space Symposium for Flight Dynamics*.
- Burton, A. and C. Früh (2021). “Multi-Hypothesis Light Curve Inversion for Space Object Attitude Determination”. In: 8th European Conference on Space Debris. ESA Space Debris Office.
- Byrd, R. H., M. E. Hribar, and J. Nocedal (1999). “An Interior Point Algorithm for Large-Scale Nonlinear Programming”. In: *SIAM Journal on Optimization* 9.4 (4), pp. 877–900. DOI: 10.1137/S1052623497325107.

- Calef, B., J. Africano, B. Birge, D. Hall, and P. Kervin (2006). “Photometric Signature Inversion”. In: *Unconventional Imaging II*. Ed. by V. L. Gamiz, P. S. Idell, and M. S. Strojnik. Vol. 6307. International Society for Optics and Photonics. SPIE, pp. 141–150. DOI: 10.1117/12.683015.
- Cegarra Polo, M., I. Vaughn, T. Kamal, and A. Lambert (2019). “Characterisation of Geosynchronous Satellites through the Analysis of On-Sky Polarimetric Signatures Obtained with a Micropolariser Array Image Sensor”. In: *Imaging and Applied Optics 2019 (COSI, IS, MATH, pcAOP)*. Imaging and Applied Optics 2019 (COSI, IS, MATH, pcAOP). Optica Publishing Group, IM3B.3. DOI: 10.1364/ISA.2019.IM3B.3.
- Coder, R. D., R. Linares, and M. J. Holzinger (2015). “Improved Models for Attitude Estimation of Agile Space Objects”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*, p. 20.
- Coder, R. D., M. J. Holzinger, and R. Linares (2017). “Three-Degree-of-Freedom Estimation of Agile Space Objects Using Marginalized Particle Filters”. In: *Journal of Guidance, Control, and Dynamics* 41.2 (2), pp. 388–400. DOI: 10.2514/1.g001980.
- Coder, R. D., C. J. Wetterer, K. M. Hamada, M. J. Holzinger, and M. K. Jah (2018). “Inferring Active Control Mode of the Hubble Space Telescope Using Unresolved Imagery”. In: *Journal of Guidance, Control, and Dynamics* 41.1 (1), pp. 164–170. DOI: 10.2514/1.G002223.
- Collins, K. A., J. F. Kielkopf, K. G. Stassun, and F. V. Hessman (2017). “AstroImageJ: Image Processing and Photometric Extraction for Ultra-Precise Astronomical Light Curves (Expanded Edition)”. DOI: 10.3847/1538-3881/153/2/77. arXiv: 1701.04817 [astro-ph].
- Cook, R. and K. E. Torrance (1982). “A Reflectance Model for Computer Graphics”. In: *ACM Trans. Graph.* 1, pp. 7–24. DOI: 10.1145/965161.806819.
- Cox, A. P., C. K. Nebelecky, R. Rudnicki, W. A. Tagliaferri, J. L. Crassidis, and B. Smith (2016). “The Space Object Ontology”. In: *International Conference on Information Fusion (FUSION)*, pp. 146–153. ISSN: 9780996452748.
- Crassidis, J. L. and F. L. Markley (2003). “Unscented Filtering for Spacecraft Attitude Estimation”. In: *Journal of Guidance, Control, and Dynamics* 26.4 (4), pp. 536–542. DOI: 10.2514/2.5102. eprint: <https://doi.org/10.2514/2.5102>.
- Crassidis, J. L., F. L. Markley, and Y. Cheng (2007). “Survey of Nonlinear Attitude Estimation Methods”. In: *Journal of Guidance, Control, and Dynamics* 30.1, pp. 12–28. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/1.22452.
- Davis, R. J., R. C. Wells, and F. L. Whipple (1957). “On Determining the Orientation of a Cylindrical Artificial Earth Satellite”. In: *Astronautica Acta* 3, pp. 231–236.
- Dianetti, A. D., R. M. Weisman, and J. L. Crassidis (2018). “Observability Analysis for Improved Space Object Characterization”. In: *Journal of Guidance, Control, and Dynamics* 41.1 (1), pp. 137–148. DOI: 10.2514/1.G002229.
- Dianetti, A. D. and J. L. Crassidis (2019). “Space Object Material Determination from Polarized Light Curves”. In: *AIAA Scitech 2019 Forum*. AIAA Scitech 2019 Forum. San Diego, California: American Institute of Aeronautics and Astronautics. ISBN: 978-1-62410-578-4. DOI: 10.2514/6.2019-0377. eprint: <https://arc.aiaa.org/doi/pdf/10.2514/6.2019-0377>.
- Dianetti, A. D. and J. L. Crassidis (2020). “Space Object Attitude Observability from Multispectral Light Curves”. In: *AIAA Scitech 2020 Forum*. DOI: 10.2514/6.2020-1098.

- Dianetti, A. D. (2020). “Resident Space Object Characterization Using Polarized and Multispectral Light Curves”. Ph.D. thesis. University at Buffalo.
- Du, X., Y. Wang, H. Hu, R. Gou, and H. Liu (2018). “The Attitude Inversion Method of Geostationary Satellites Based on Unscented Particle Filter”. In: *Advances in Space Research* 61.8 (8), pp. 1984–1996. DOI: 10.1016/j.asr.2018.01.018.
- Eisele, R. (2008). *Proof: Quaternion from Two Vectors*. Computer Science and Machine Learning. URL: <https://www.xarg.org/proof/quaternion-from-two-vectors/> (visited on 08/17/2022).
- Elfring, J., E. Torta, and R. van de Molengraft (2021). “Particle Filters: A Hands-On Tutorial”. In: *Sensors* 21.2, p. 438. ISSN: 1424-8220. DOI: 10.3390/s21020438.
- Enomoto, T., T. Hanada, K. Fujita, and Y. Kitazawa (2016). “Monitoring the Health of Geosynchronous Spacecraft Using Photometric Observations”. In: *Proceedings of the International Astronautical Congress, IAC*. 67th International Astronautical Congress, IAC 2016. Vol. 0.
- ESA, S. D. O. (2019). *ESA’s Annual Space Environment Report*. GEN-DB-LOG-00271-OPS-SD. ESA.
- ESA (2022). *Space Surveillance and Tracking - SST Segment*. URL: https://www.esa.int/Space_Safety/Space_Surveillance_and_Tracking_-_SST_Segment (visited on 07/12/2022).
- EU SST (2021). *Service Portfolio*. URL: <https://portal.eusst.eu/portalng/public-download/ServicePortfolio> (visited on 07/15/2022).
- Fan, S., A. M. Friedman, and C. Früh (2019). “Satellite Shape Recovery from Light Curves with Noise”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*. Ed. by S. Ryan.
- Fan, S. and C. Früh (2020). “A Direct Light Curve Inversion Scheme in the Presence of Measurement Noise”. In: *The Journal of the Astronautical Sciences* 67.2 (2), pp. 740–761. ISSN: 2195-0571. DOI: 10.1007/s40295-019-00190-3.
- Fan, S. and C. Früh (2021). “Multi-Hypothesis Light Curve Inversion Scheme for Convex Objects with Minimal Observations”. In: 8th European Conference on Space Debris. ESA Space Debris Office.
- Fan, S. (2020). “The Light Curve Simulation and Its Inversion Problem for Human-Made Space Objects”. Ph.D. thesis. Purdue University Graduate School. DOI: 10.25394/PGS.12749570.v1.
- Friedman, A. M. and C. Früh (2017). “Determining Debris Characteristics from Observability Analysis of Artificial Near-Earth Objects”. In: 7th European Conference on Space Debris.
- Friedman, A. M. and C. Früh (2018). “Determining Characteristics of Artificial Near-Earth Objects Using Observability Analysis”. In: *Acta Astronautica* 144, pp. 405–421. ISSN: 0094-5765. DOI: 10.1016/j.actaastro.2017.12.028.
- Friedman, A. M., S. Fan, C. Früh, and T. Schildknecht (2019). “Observability of Light Curve Shape Inversion Based on Optical Data”. In: First International Orbital Debris Conference.
- Friedman, A. M. and C. Früh (2022). “Observability of Light Curve Inversion for Shape and Feature Determination Exemplified by a Case Analysis”. In: *Journal of the Astronautical Sciences* 69, pp. 537–569. ISSN: 0021-9142. DOI: 10.1007/s40295-021-00293-w.
- Früh, C. and M. K. Jah (2014). “Coupled Orbit–Attitude Motion of High Area-to-Mass Ratio (HAMR) Objects Including Efficient Self-Shadowing”. In: *Acta Astronautica* 95, pp. 227–241. DOI: 10.1016/j.actaastro.2013.11.017.

- Früh, C. (2015). “Observability of Non-cooperative Space Object’s Tracking and Characterization”. In: *25th International Symposium on Space Flight Dynamics*, p. 18.
- Fulcoly, D. O., K. I. Kalamaroff, and F. K. Chun (2012). “Determining Basic Satellite Shape from Photometric Light Curves”. In: *Journal of Spacecraft and Rockets* 49.1 (1), pp. 76–82. ISSN: 0022-4650. DOI: 10.2514/1.A32002.
- Furfaro, R., R. Linares, D. Gaylor, M. K. Jah, and R. Walls (2016). “Resident Space Object Characterization and Behavior Understanding via Machine Learning and Ontology-based Bayesian Networks”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Furfaro, R., R. Linares, and V. Reddy (2018). “Space Objects Classification via Light-Curve Measurements: Deep Convolutional Neural Networks and Model-based Transfer Learning”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Furfaro, R., R. Linares, and V. Reddy (2019). “Shape Identification of Space Objects via Light Curve Inversion Using Deep Learning”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Giese, R.-H. (1963). *Attitude Determination from Specular and Diffuse Reflection by Cylindrical Artificial Satellites*. Special Report 127. Massachusetts.
- Hall, D., J. Africano, D. Archambeault, B. Birge, D. Witte, and P. Kervin (2006). “AMOS Observations of NASA’s IMAGE Satellite”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Hall, D. and P. Kervin (2014). “Optical Characterization of Deep-Space Object Rotation States”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Hecht, E. (2002). *Optics*. 4th ed. Reading, Mass: Addison-Wesley. 698 pp. ISBN: 978-0-8053-8566-3.
- Henderson, L. S., P. Goyal, and K. Subbarao (2012). “Inverse Problem Formulation Coupled with Unscented Kalman Filtering for State and Shape Estimation of Space Objects”. In: *Advances in the Astronautical Sciences* 143, pp. 199–216.
- Henderson, L. S. (2014). “Modeling, Estimation, And Analysis Of Unresolved Space Object Tracking And Identification”. Ph.D. thesis. The University of Texas at Arlington. 187 pp.
- Hinks, J., R. Linares, and J. L. Crassidis (2013). “Attitude Observability from Light Curve Measurements”. In: *AIAA Guidance, Navigation, and Control (GNC) Conference*. AIAA Guidance, Navigation, and Control (GNC) Conference. Boston, MA: American Institute of Aeronautics and Astronautics. ISBN: 978-1-62410-224-0. DOI: 10.2514/6.2013-5005.
- Holzinger, M. J., K. T. Alfriend, C. J. Wetterer, K. K. Luu, C. Sabol, K. M. Hamada, and A. Harms (2012). *Attitude Estimation for Unresolved Agile Space Objects with Shape Model Uncertainty*. Air Force Research Laboratory Wright-Patterson AFB OH.
- Holzinger, M. J., K. T. Alfriend, C. J. Wetterer, K. K. Luu, C. Sabol, and K. M. Hamada (2014). “Photometric Attitude Estimation for Agile Space Objects with Shape Uncertainty”. In: *Journal of Guidance, Control, and Dynamics* 37.3 (3), pp. 921–932. DOI: 10.2514/1.58002. eprint: <https://doi.org/10.2514/1.58002>.
- Horn, B. K. P. (1984). “Extended Gaussian Images”. In: *IEEE* 72.12 (12), pp. 1671–1686.
- Jankovic, M. (2022). “Characterization and Capture of Space Debris Objects Using Domain Ontology and Optimal Control”. Ph.D. thesis. DOI: 10.26092/e1ib/1476.
- Jia, B., K. D. Pham, E. Blasch, Z. Wang, D. Shen, and G. Chen (2018). “Space Object Classification Using Deep Neural Networks”. In: *2018 IEEE Aerospace Conference*. Vol. 2018-March. IEEE, pp. 1–8. ISBN: 978-1-5386-2014-4. DOI: 10.1109/AERO.2018.8396567.

- Juillard, M. and J.-P. Kneib (2021). “Simulation Tool: Resources Management in High Performance Avionic for ADR Missions”. In: *2021 IEEE Aerospace Conference (50100)*. 2021 IEEE Aerospace Conference (50100), pp. 1–12. DOI: 10.1109/AERO50100.2021.9438138.
- Julier, S. J. and J. K. Uhlmann (2004). “Unscented Filtering and Nonlinear Estimation”. In: *Proceedings of the IEEE* 92.3 (3), pp. 401–422. ISSN: 1558-2256. DOI: 10.1109/JPROC.2003.823141.
- Kaasalainen, M. and J. Torppa (2001). “Optimization Methods for Asteroid Lightcurve Inversion. I. Shape Determination”. In: *Icarus* 153.1 (1), pp. 24–36. DOI: 10.1006/icar.2001.6673.
- Kaasalainen, M. (2016). “Asteroid Models from Photometry and Complementary Data Sources”. In: *AIP Conference Proceedings*. Vol. 1732. 020003. DOI: 10.1063/1.4948806. eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.4948806>.
- Kaplan, G. H. (2006). “The IAU Resolutions on Astronomical Reference Systems, Time Scales, and Earth Rotation Models”. arXiv: astro-ph/0602086.
- Kececy, T. M., D. Hall, K. M. Hamada, and D. Stocker (2007). “Satellite Maneuver Detection Using Two-Line Elements Data”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, pp. 175–190.
- Kesemen, O., B. U. Ğ. R. A. Tiryaki, E. D. A. Özkul, and Ö. Z. G. E. Tezel (2018). “Determination of the Confidence Intervals for Multimodal Probability Density Functions”. In: *Gazi University Journal of Science* 31.1, pp. 310–326.
- Kim, J. (2011). “Nonnegative Matrix and Tensor Factorizations, Least Squares Problems, and Applications”. Ph.D. thesis. Atlanta, Georgia, US: Georgia Institute of Technology. 145 pp.
- Kucharski, D. and M. K. Jah (2021). “Astrometric Data Engine for Attitude Anomaly Detection of Megaconstellation Satellites”. In: *7th Annual Space Traffic Management Conference*.
- Larsson, S. (1996). “Parameter Estimation in Epoch Folding Analysis”. In: *Astronomy and Astrophysics Supplement Series* 117.1 (1), pp. 197–201. ISSN: 0365-0138, 1286-4846. DOI: 10.1051/aas:1996150.
- Lawson, C. L. and R. J. Hanson (1995). *Solving Least Squares Problems*. Society for Industrial and Applied Mathematics. 351 pp. ISBN: 978-0-89871-356-5.
- Linares, R., M. K. Jah, and J. L. Crassidis (2012). “Inactive Space Object Shape Estimation via Astrometric and Photometric Data Fusion”. In: *Strengthening Cooperation in The 21st Century* 143, pp. 217–232.
- Linares, R., J. L. Crassidis, and M. K. Jah (2014a). “Particle Filtering Light Curve Based Attitude Estimation for Non-Resolved Space Objects”. In: *Advances in the Astronautical Sciences* 152, pp. 119–130.
- Linares, R., M. K. Jah, J. L. Crassidis, F. A. Leve, and T. M. Kececy (2014b). “Astrometric and Photometric Data Fusion for Inactive Space Object Mass and Area Estimation”. In: *Acta Astronautica* 99, pp. 1–15. ISSN: 0094-5765. DOI: 10.1016/j.actaastro.2013.10.018.
- Linares, R., M. K. Jah, J. L. Crassidis, and C. K. Nebelecky (2014c). “Space Object Shape Characterization and Tracking Using Light Curve and Angles Data”. In: *Journal of Guidance, Control, and Dynamics* 37.1 (1), pp. 13–25. DOI: 10.2514/1.62986.

Bibliography

- Linder, E., J. Šilha, T. Schildknecht, and M. Hager (2015). “Extraction of Spin Periods of Space Debris from Optical Light Curves”. In: *66th Astronautical Congress International*. Curran Associates, Inc.
- Linares, R. and J. L. Crassidis (2016a). “Dynamic Observability Analysis for Attitude, Angular Velocity, Shape, and Surface Parameters”. In: *26th AAS/AIAA Space Flight Mechanics Meeting*. Vol. 158, pp. 2016–515.
- Linares, R. and R. Furfaro (2016b). “Space Object Classification Using Deep Convolutional Neural Networks”. In: *2016 19th International Conference on Information Fusion (FUSION)*, pp. 1140–1146. ISBN: 978-0-9964527-4-8.
- Linares, R. and R. Furfaro (2017). “Space Objects Maneuvering Detection and Prediction via Inverse Reinforcement Learning”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Linares, R. and J. L. Crassidis (2018). “Space-Object Shape Inversion via Adaptive Hamiltonian Markov Chain Monte Carlo”. In: *Journal of Guidance, Control, and Dynamics* 41.1 (1), pp. 47–58. DOI: 10.2514/1.G002296.
- Liou, J. C. (2013). “Engineering and Technology Challenges for Active Debris Removal”. In: *Progress in Propulsion Physics*. Progress in Propulsion Physics. St. Petersburg, Russian: EDP Sciences, pp. 735–748. ISBN: 978-2-7598-0876-2. DOI: 10.1051/eucass/201304735.
- Little, J. J. (1985). “Recovering Shape and Determining Attitude from Extended Gaussian Images”. Ph.D. thesis. Vancouver, Canada: University of British Columbia. 93 pp. DOI: 10.14288/1.0051804.
- Liu, T. and K. U. Schreiber (2021). “Photometric Space Object Classification via Deep Learning Algorithms”. In: *Acta Astronautica* 185, pp. 161–169. ISSN: 0094-5765. DOI: 10.1016/j.actaastro.2021.05.008.
- Ly, S., H. Zhang, Y. Zhang, B. Ning, and R. Qi (2022). “Design of an Integrated Platform for Active Debris Removal”. In: *Aerospace* 9.7, p. 339. ISSN: 2226-4310. DOI: 10.3390/aerospace9070339.
- Magnusson, P. (1986). “Distribution of Spin Axes and Senses of Rotation for 20 Large Asteroids”. In: *Icarus* 68.1 (1), pp. 1–39. ISSN: 0019-1035. DOI: 10.1016/0019-1035(86)90072-2.
- Markley, F. L. and J. L. Crassidis (2014). *Fundamentals of Spacecraft Attitude Determination and Control*. Space Technology Library 33. New York: Springer. ISBN: 978-1-4939-0802-8.
- Matsushita, Y., R. Arakawa, Y. Yoshimura, and T. Hanada (2019). “Light Curve Analysis and Attitude Estimation of Space Objects Focusing on Glint”. In: p. 8.
- Minkowski, H. (1897). “Allgemeine Lehrsätze Über Die Convexen Polyeder (German) [General Theorems about Convex Polyhedra]”. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1897, pp. 198–220.
- Möller, T. and B. Trumbore (1997). “Fast, Minimum Storage Ray-Triangle Intersection”. In: *Journal of Graphics Tools* 2.1, pp. 21–28. DOI: 10.1080/10867651.1997.10487468. eprint: <https://doi.org/10.1080/10867651.1997.10487468>.
- Mol, M. (2022). *Determine If Two Triangles Overlap*. Rosetta Code. URL: https://rosettacode.org/wiki/Determine_if_two_triangles_overlap (visited on 06/01/2022).
- Nicodemus, F. E., J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis (1977). “Geometrical Considerations and Nomenclature for Reflectance”. In: *NBS monograph*, p. 4.

- Nicodemus, F. E. (1965). “Directional Reflectance and Emissivity of an Opaque Surface”. In: *Applied Optics* 4.7 (7), pp. 767–775. DOI: 10.1364/AO.4.000767.
- Nocedal, J. and S. J. Wright (2006). *Numerical Optimization*. 2nd ed. Springer Series in Operations Research. New York: Springer. 664 pp. ISBN: 978-0-387-30303-1.
- OECD (2019). *The Space Economy in Figures*, p. 200. DOI: <https://doi.org/10.1787/c5996201-en>.
- Pasqual, M. C. and K. L. Cahoy (2017). “Active Polarimetric Measurements for Identification and Characterization of Space Debris”. In: *IEEE Transactions on Aerospace and Electronic Systems* 53.6 (6), pp. 2706–2717. ISSN: 1557-9603. DOI: 10.1109/TAES.2017.2711718.
- Payne, T. E., S. A. Gregory, A. Dentamaro, M. Ernst, J. Hollon, A. Kruchten, A. B. Chaudhary, and P. D. Dao (2017). “Development and Evaluation of New Methods for Estimating Albedo-Area for Stable GEOs”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*, 24, p. 15.
- Pedone, G., D. Vallverdú Cabrera, M. G. Dimitrova Vesselinova, G. Cirillo, Y. Heinz, F. Schiemenz, and J. Utzmann (2021). “SPOOK: A Tool for Space Objects Catalogue Creation and Maintenance Supporting Space Safety and Sustainability”. In: *Acta Astronautica* 188, pp. 89–98. ISSN: 0094-5765. DOI: 10.1016/j.actaastro.2021.07.026.
- Piergentili, F., F. Santoni, and P. Seitzer (2017). “Attitude Determination of Orbiting Objects from Lightcurve Measurements”. In: *IEEE Transactions on Aerospace and Electronic Systems* 53.1 (1), pp. 81–90. DOI: 10.1109/TAES.2017.2649240.
- Planck, M. and M. Masius (1914). *The Theory of Heat Radiation*. In collab. with University of California Libraries. Philadelphia, P. Blakiston’s Son & Co. 256 pp.
- Powell, M. J. D. (1964). “An Efficient Method for Finding the Minimum of a Function of Several Variables without Calculating Derivatives”. In: *The Computer Journal* 7, pp. 155–162.
- Qin, J., H. Peng, and T. Dang (2018). “Space Object Attitude and Position Determination Based on Astrometry and Photometric Data Fusion”. In: *Procedia Computer Science* 131, pp. 327–335. ISSN: 0000000000. DOI: 10.1016/j.procs.2018.04.171.
- Renhorn, I. G. E., T. Hallberg, and G. D. Boreman (2015). “Efficient Polarimetric BRDF Model”. In: p. 21.
- Reyes, J. A., K. Fulford, E. Plis, R. Hoffmann, V. Murray, H. Cowardin, D. Cone, D. Ferguson, M. Bengtson, J. Shah, and D. Engelhart (2021). “Spectroscopic Behavior of Various Materials in a GEO Simulated Environment”. In: *Acta Astronautica* 189, pp. 576–583. DOI: 10.1016/j.actaastro.2021.09.014.
- Reyes, J. A. (2018). “Characterization of Spacecraft Materials Using Reflectance Spectroscopy”. In: p. 10.
- Richardson, M. and T. M. Kelecý (2018). “Mass Estimation Through Fusion of Astrometric and Photometric Data Collection with Applications to Orbital Debris Characterization”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, p. 16.
- Runnalls, A. (2007). “Kullback-Leibler Approach to Gaussian Mixture Reduction”. In: *IEEE Transactions on Aerospace and Electronic Systems* 43.3 (3), pp. 989–999. ISSN: 1557-9603. DOI: 10.1109/TAES.2007.4383588.
- Ryan, E. V. and W. H. Ryan (2007). “Space Object Characterization Studies and the Magdalena Ridge Observatory’s 2 . 4-Meter Telescope”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*. Vol. 48. Citeseer.

Bibliography

- Sarkar, J. (2017). *Computer Aided Design*. 1st edition. CRC Press. 740 pp. ISBN: 978-1-138-88544-8.
- SatCen (2023). *Space Situational Awareness (SSA)*. European Union Satellite Centre. URL: <https://www.satcen.europa.eu/page/ssa> (visited on 02/04/2023).
- Schiemenz, F., J. Utzmann, and H. Kayal (2020). “Adaptive Gaussian Mixture Based Orbit Determination with Combined Atmospheric Density Uncertainty Consideration”. In: *Advances in Space Research* 66.7 (7), pp. 1609–1634. ISSN: 0273-1177. DOI: 10.1016/j.asr.2020.05.042.
- Schott, J. R., S. D. Brown, R. V. Raqueño, H. N. Gross, and G. Robinson (1999). “An Advanced Synthetic Image Generation Model and Its Application to Multi/Hyperspectral Algorithm Development”. In: *Canadian Journal of Remote Sensing* 25.2, pp. 99–111. DOI: 10.1080/07038992.1999.10874709. eprint: <https://doi.org/10.1080/07038992.1999.10874709>.
- Schiemenz, F. (2021). “Covariance and Uncertainty Realism for Low Earth Orbiting Satellites via Quantification of Dominant Force Model Uncertainties”. Ph.D. thesis. Universität Würzburg. DOI: 10.25972/OPUS-24947.
- Shampine, L. F. and M. K. Gordon (1975). *Computer Solution of Ordinary Differential Equations: The Initial Value Problem*. W. H. Freeman. 318 pp. ISBN: 978-0-7167-0461-4.
- Shirley, P. and R. K. Morley (2003). *Realistic Ray Tracing*. 2nd edition. Natick, Massachusetts: A K Peters/CRC Press. 235 pp. ISBN: 978-1-56881-461-2.
- Shuster, M. D. et al. (1993). “A Survey of Attitude Representations”. In: *Navigation* 8.9, pp. 439–517.
- Shuster, M. D. (2003). “Uniform Attitude Probability Distributions”. In: *Journal of the Astronautical Sciences* 51.4 (4), pp. 451–475. DOI: 10.1007/BF03546294.
- Solà, J. (2017). “Quaternion Kinematics for the Error-State Kalman Filter”. In: *CoRR* abs/1711.02508. arXiv: 1711.02508.
- Somers, P. W. (2011). “Cylindrical RSO Signatures, Spin Axis Orientation and Rotation Period Determination”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Stryjewski, J., D. Hand, D. Tyler, S. Murali, M. Roggemann, and N. Peterson (2010). “Real Time Polarization Light Curves for Space Debris and Satellites”. In.
- Subbarao, K. and L. S. Henderson (2019). “Observability and Sensitivity Analysis of Lightcurve Measurement Models for Use in Space Situational Awareness”. In: *Inverse Problems in Science and Engineering* 27.10, pp. 1399–1424. ISSN: 1741-5977. DOI: 10.1080/17415977.2018.1554655.
- Tassoul, J. L. (2004). *A Concise History of Solar and Stellar Physics*. Princeton University Press. 314 pp. ISBN: 978-0-691-11711-9.
- Titterton, D., J. L. Weston, and J. Weston (2004). *Strapdown Inertial Navigation Technology*. IET. 578 pp. ISBN: 978-0-86341-358-2. Google Books: NN95yh2cVwYC.
- Tobergte, D. R. and S. Curtis (2013). *Handbook of CCD Astronomy, Second Edition*. Vol. 53. 1689 pp. ISBN: 978-85-7811-079-6. DOI: 10.1017/CB09781107415324.004.
- U.S. Government (2019). *Space-Track.Org*. URL: <https://www.space-track.org/> (visited on 06/26/2022).
- Utzmann, J., M. G. Dimitrova Vesselinova, and O. Rodriguez Fernandez (2019). “Airbus Robotic Telescope”. In: 1st NEO and Debris Detection Conference.

- Vallado, D. A. and W. D. McClain (2013). *Fundamentals of Astrodynamics and Applications*. 4. ed. Space Technology Library 21. Hawthorne, Calif: Microcosm Press. 1106 pp. ISBN: 978-1-881883-18-0.
- Vallverdú Cabrera, D., J. Uetzmann, and R. Förstner (2021a). “Integration of Attitude Characterization in a Space Debris Catalogue Using Light Curves”. In: 8th European Conference on Space Debris. ESA Space Debris Office.
- Vallverdú Cabrera, D., J. Uetzmann, and R. Förstner (2021b). “Inversion of the Shape of Space Debris from Non-Resolved Optical Measurements within SPOOK”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.
- Vallverdú Cabrera, D., J. Uetzmann, and R. Förstner (2023). “The Adaptive Gaussian Mixtures Unscented Kalman Filter for Attitude Determination Using Light Curves”. In: *Advances in Space Research* 71.6, pp. 2609–2628. ISSN: 02731177. DOI: 10.1016/j.asr.2022.11.027.
- Van der Merwe, R. and E. Wan (2001). “The Square-Root Unscented Kalman Filter for State and Parameter-Estimation”. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*. Vol. 6, pp. 3461–3464. DOI: 10.1109/ICASSP.2001.940586.
- VanderPlas, J. T. (2018). “Understanding the Lomb–Scargle Periodogram”. In: *The Astrophysical Journal Supplement Series* 236.1 (1). DOI: 10.3847/1538-4365/aab766. arXiv: 1703.09824.
- Vatti, B. R. (1992). “A Generic Solution to Polygon Clipping”. In: *Communications of the ACM* 35.7 (7), pp. 56–63.
- Viikinkoski, M., J. Hanuš, M. Kaasalainen, F. Marchis, and J. Ďurech (2017). “Adaptive Optics and Lightcurve Data of Asteroids: Twenty Shape Models and Information Content Analysis”. In: *Astronomy & Astrophysics* 607, A117.
- Vittaldev, V. and R. P. Russell (2016). “Multidirectional Gaussian Mixture Models for Nonlinear Uncertainty Propagation”. In: *Computer Modeling in Engineering & Sciences* 111.1 (1), pp. 83–117. ISSN: 1526-1506. DOI: 10.3970/cmescs.2016.111.083.
- Voight, J. (2021). *Quaternion Algebras*. Graduate Texts in Mathematics 288. Cham, Switzerland: Springer. 885 pp. ISBN: 978-3-030-56692-0. DOI: 10.1007/978-3-030-56694-4.
- Wang, Y., X. Du, J. Zhao, and R. Gou (2018). “Apparent Magnitude Calculation Method for Complex Shaped Space Objects”. In: *Advances in Space Research* 62.11 (11), pp. 2988–2997. DOI: 10.1016/j.asr.2018.07.033.
- Watanabe, K., T. Hanada, and K. Fujita (2017). “Dynamics Observation of Space Objects Using Adaptive Optics Simulation and Light Curve Analysis”. In: *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*. Ed. by S. Ryan, 123, p. 123.
- Wertz, J., D. Everett, and J. Puschell (2011). *Space Mission Engineering: The New SMAD*. Space Technology Library. Microcosm Press. ISBN: 978-1-881883-15-9.
- Wetterer, C. J. and M. K. Jah (2009). “Attitude Determination from Light Curves”. In: *Journal of Guidance, Control, and Dynamics* 32.5 (5), pp. 1648–1651. DOI: 10.2514/1.44254. eprint: <https://doi.org/10.2514/1.44254>.
- Wetterer, C. J., R. Linares, J. L. Crassidis, T. M. Kelecy, M. K. Ziebart, M. K. Jah, and P. J. Cefola (2014). “Refining Space Object Radiation Pressure Modeling with Bidirectional Reflectance Distribution Functions”. In: *Journal of Guidance, Control, and Dynamics* 37.1 (1), pp. 185–196. DOI: 10.2514/1.60577.
- White, R. A. and S. W. Stemwedel (1992). “The Quadrilateralized Spherical Cube and Quad-Tree for All Sky Data”. In: *Astronomical Data Analysis Software and Systems I*.

Bibliography

- Ed. by D. M. Worrall, C. Biemesderfer, and J. Barnes. Vol. 25. Astronomical Society of the Pacific Conference Series, p. 379.
- Wiesel, W. E. (2010). *Modern Orbit Determination*. 2nd ed. CreateSpace Independent Publishing Platform. 147 pp.
- Willison, A. and D. Bédard (2015). “Light Curve Simulation Using Spacecraft CAD Models and Empirical Material Spectral BRDFs”. In: DOI: 10.13140/RG.2.1.3654.8245.
- Willison, A. (2016). “A Novel Approach to Modeling Spacecraft Spectral Reflectance”. In: *Advances in Space Research*, p. 13.
- Williams, J. G. (1969). “The Determination of the Orientation of a Tumbling Cylinder from the Shape of the Light Curve”. In: *Symposium on Optical Properties of Orbiting Satellites*, pp. 31–39.
- Wu, Y., D. Hu, M. Wu, and X. Hu (2005). “Unscented Kalman Filtering for Additive Noise Case: Augmented versus Nonaugmented”. In: *IEEE Signal Processing Letters* 12.5, pp. 357–360. ISSN: 1070-9908. DOI: 10.1109/LSP.2005.845592.
- Wu, Y., H. Lu, Z. Zhang, and F. Zhao (2016). “Space Object Characterization with Infrared Signature from a Space-Based Platform”. In: *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*. IEEE, pp. 1242–1247. ISBN: 978-1-5090-3710-0. DOI: 10.1109/CISP-BMEI.2016.7852905.
- Yun, S. and R. Zanetti (2020). “Nonlinear Filtering of Light-Curve Data”. In: *Advances in Space Research* 66.7 (7), pp. 1672–1688. ISSN: 0273-1177. DOI: 10.1016/j.asr.2020.06.024.
- Zhu, C., R. H. Byrd, and J. Nocedal (1997). “L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN Routines for Large Scale Bound Constrained Optimization”. In: *ACM Transactions on Mathematical Software* 23.4 (4), pp. 550–560.
- Zigo, M., J. Šilha, and S. Krajcovic (2019). “BVRI Photometry of Space Debris Objects at the Astronomical and Geophysical Observatory in Modra”. In: *Advanced Maui Optical and Space Surveillance Technologies Conference*.

A. Framework

This appendix briefly describes the two main resources that have been used to support the scientific work presented in this thesis.

These are

- the Special Perturbations Orbit determination and Orbit analysis toolKit (SPOOK), a SST software framework; and
- the Airbus Robotic Telescope (ART), an optical telescope and main source of real measurements.

The following two sections describe these two systems in more detail.

A.1. Special Perturbations Orbit Determination and Orbit Analysis Toolkit

The Special Perturbations Orbit determination and Orbit analysis toolKit (SPOOK) is a collection of software tools developed by the Security in Space Team of the Future Programmes department of Airbus Defence & Space GmbH. It provides a wide range of SST capabilities, the most recent comprehensive summary of which can be found in Pedone et al. (2021).

Most notably it provides the functionality of an end-to-end SST pipeline that can be used to create and maintain a catalogue of RSOs using optical and Radar measurements. Furthermore, it can perform probabilistic conjunction assessment of catalogued RSOs (Schiemenz, 2021). As a result of this thesis, SPOOK is now equipped with attitude and shape characterization capabilities, too.

These activities are enabled by a lower level layer of functionalities that include high-fidelity orbit propagation, sensor simulation (extended to light curve generation within the scope of this work), observation planning, measurement correlation and Orbit Determination (OD).

SPOOK itself is composed of two major code-bases that interface with each other: one is written in Fortran; the other one, in Python. On the one hand, the Fortran part, or Fortran SPOOK performs the most computationally demanding tasks (*e.g.* light curve simulation and sequential Attitude Determination (AD)). As such, it is considered to be the core of SPOOK. On the other hand, the Python part, or Python SPOOK, complements and extends upon the Fortran core. It acts as a wrapper of the core functionality provided by Fortran SPOOK, but also includes other capabilities, entirely independent of the Fortran part (*e.g.* shape

A. Framework

determination and frequency analysis of light curves). Furthermore, it is the part responsible for plot generation, as well as to interface with other third-party tools.

SPOOK has been used to produce all the simulations, numerical analyses and plots of this thesis.

A.2. Airbus Robotic Telescope

The Airbus Robotic Telescope (ART) is an optical telescope located in Extremadura, Spain. Its aperture has a diameter of 40 cm, and its f-number is 2.4. It has a rectangular field of view of $2.15^\circ \times 1.43^\circ$. ART is equipped with a Charge-Coupled Device (CCD) detector whose nominal wavelength range spans from 300 nm to 1000 nm. The reader may refer to Utzmann et al. (2019) for a fully-detailed list of features. Fig. A.1 shows ART.



Figure A.1.: The Airbus Robotic Telescope

Within the scope of this thesis, real ART measurements have been used in Section 6.2. Furthermore, ART has been used as a simulation model both in the Shape Determination Chapter and in Section 6.3.

B. Summary of OC Methods from Light Curves in the Literature

Tables B.1 to B.5 synthesize the content of Section 2.2 of the Literature Review Chapter for a quick and compact overview.

Table B.1.: Summary of AD techniques

Characteristic	Input	Methods	Constraints
Apparent period	Light curve	Fourier-based; epoch folding	No single all-rounder
	SLR	Direct observation	Only collaborative targets
Phase plot	Apparent period; light curve	Modulus operation	
Attitude status	Light curve	CNN	Simulated data
Sidereal period & spin axis	Light curve	Fourier-based; epoch method	Only under constraining hypotheses; only spinning objects
Sidereal period & spin axis	Light curve	Constrained glint pattern	Only cylinder shapes
Full attitude state @ one instant	Light curve	Heuristic optimization	Known shape & albedo
Full attitude state time series	Light curve	Sequential estimation (EKF, UKF, PF, etc.)	Known shape & albedo; initial estimate necessary

Table B.2.: Summary of shape characterization techniques

Characteristic	Input	Methods	Constraints
Convex-equivalent shape	Light curve; attitude	EGI plus Minkowski minimization	Coupled with albedo
Average albedo-area	Phase plot	Phase plot comparison	Only stable GEO objects
Shape	Light curve	Lower bound intensity comparison	Only trained shapes
Shape	Light curve	Convolution Neural Network	Only trained shapes; higher success rate

B. Summary of OC Methods from Light Curves in the Literature

Table B.3.: Summary of material characterization techniques

Characteristic	Input	Methods	Constraints
Surface material	Polarized light curve	Multiple-model adaptive est.	
Surface material	Spectroscopic light curve		
Surface material classification	Polarized/spectroscopic light curve	ML	Not implemented in the literature

Table B.4.: Summary of combined characterization techniques

Characteristic	Input	Methods	Constraints
Attitude & orbit & reflective properties	Light curve	PF plus GM	Known shape
Attitude & shape	Light curve	Multi-hypothesis UKF	Shape must be among proposed candidates
Attitude & orbit & reflective properties & dimensions	Light curve	PF or batch-UKF	Shape topology restricted to a cuboid only
Attitude & shape	Light curve	Hamiltonian PF	Albedo and shape are coupled

Table B.5.: Summary of abstract characterization techniques

Characteristic	Input	Methods	Constraints
Object class (status, bus type)	Light curve	CNN	
Orbit change manoeuvre	Orbital ephemeris	Energy assessment; UKF-residual analysis; ML	
Attitude manoeuvre	Light curve; shape	Stochastic attitude model & PF	

C. Quaternion Algebra

This appendix is a summary of quaternion algebra, meant to serve as a quick reference. Since different quaternion conventions exist in the literature, its second purpose is to define which exact conventions are used in this thesis. Refer to relevant literature (e.g. Voight (2021)) for a fully detailed, formal derivation of quaternion algebra.

The quaternion algebra, also known as Hamilton space and denoted here by \mathbb{H} , deals with quaternions, represented as

$$\mathbf{q} = q_r + q_i \mathbf{i} + q_j \mathbf{j} + q_k \mathbf{k} = [1 \ \mathbf{i} \ \mathbf{j} \ \mathbf{k}] \mathbf{q}, \quad (\text{C.1})$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k} \in \mathbb{H}$ are the three independent hyper-imaginary numbers, or basic quaternions, and

$$\mathbf{q} = \begin{bmatrix} q_r \\ q_i \\ q_j \\ q_k \end{bmatrix} \in \mathbb{R}^4 \quad (\text{C.2})$$

is the vector representation of \mathbf{q} . This is the convention adopted in this document: quaternions are represented in bold font; their vector representation, in normal font.

Analogously, if a 3-dimensional vector $v \in \mathbb{R}^3$ appears in bold font, it represents a quaternion with 0 real part whose hyper-imaginary components correspond to each dimension of v :

$$\mathbf{v} = [\mathbf{i} \ \mathbf{j} \ \mathbf{k}] v. \quad (\text{C.3})$$

The real part of a quaternion is obtained with the map $\Re : \mathbb{H} \rightarrow \mathbb{R}$:

$$\Re(\mathbf{q}) = q_r; \quad (\text{C.4})$$

and its hyper-imaginary part, with $\Im : \mathbb{H} \rightarrow \mathbb{R}^3$:

$$\mathbf{q} := \Im(\mathbf{q}) = \begin{bmatrix} q_i \\ q_j \\ q_k \end{bmatrix}. \quad (\text{C.5})$$

Addition is performed using the same rules for real numbers applied to the real and to each of the three hyper-imaginary parts,

$$\mathbf{q} + \mathbf{p} = (q_r + p_r) + (q_i + p_i)\mathbf{i} + (q_j + p_j)\mathbf{j} + (q_k + p_k)\mathbf{k}. \quad (\text{C.6})$$

Multiplication is governed by the quaternion equation

$$\mathbf{i}\mathbf{i} = \mathbf{j}\mathbf{j} = \mathbf{k}\mathbf{k} = \mathbf{i}\mathbf{j}\mathbf{k} = -1. \quad (\text{C.7})$$

C. Quaternion Algebra

This is known as the Hamilton convention (Markley et al., 2014). Following these rules, one can multiply any two arbitrary quaternions by applying associative and distributive laws to each of the real and hyper-imaginary components. Note, however, that quaternion multiplication is not commutative. For the sake of clarity, the multiplication of two quaternions is indicated throughout this thesis with the symbol \odot , following the nomenclature for the Hamilton quaternion product used in Markley et al. (2014).¹ All in all, from Equation C.7, the quaternion product used in this thesis becomes

$$\begin{aligned} \mathbf{q} \odot \mathbf{p} = & q_r p_r - q_i p_i - q_j p_j - q_k p_k \\ & + (q_r p_i + q_i p_r + q_j p_k - q_k p_j) \mathbf{i} \\ & + (q_r p_j - q_i p_k + q_j p_r + q_k p_i) \mathbf{j} \\ & + (q_r p_k + q_i p_j - q_j p_i + q_k p_r) \mathbf{k}. \end{aligned} \quad (\text{C.8})$$

Note that $\mathbb{R} \in \mathbb{H}$. Thus, the multiplication of a real number a against a quaternion \mathbf{q} is just a particular case of quaternion multiplication, in which one of the two quaternions has zero hyper-complex parts. From Equation C.8, in this particular case commutativity is recovered, so that

$$\mathbf{q}a = a\mathbf{q} = aq_r + aq_i\mathbf{i} + aq_j\mathbf{j} + aq_k\mathbf{k}. \quad (\text{C.9})$$

Because of this, the \odot symbol is dropped for the scalar-quaternion multiplication, so as not to abuse it.

It is often most convenient to express this multiplication in vector notation, where it becomes a matrix-vector product.

$$q \odot p = [\mathbf{q}]_{\odot}^+ p = [\mathbf{p}]_{\odot}^- q. \quad (\text{C.10})$$

The maps $[\cdot]_{\odot}^+, [\cdot]_{\odot}^- : \mathbb{H} \rightarrow \mathbb{R}^{4 \times 4}$ are the quaternion multiplication matrices:

$$[\mathbf{q}]_{\odot}^+ = [q \quad \Theta(\mathbf{q})] \quad (\text{C.11})$$

and

$$[\mathbf{q}]_{\odot}^- = [q \quad \Xi(\mathbf{q})]. \quad (\text{C.12})$$

In turn, the maps $\Theta, \Xi : \mathbb{H} \rightarrow \mathbb{R}^{4 \times 3}$ are

$$\Theta(\mathbf{q}) = \begin{bmatrix} -\mathbf{q}^T \\ q_r I_3 + [\mathbf{q}]_{\times} \end{bmatrix} \quad (\text{C.13})$$

and

$$\Xi(\mathbf{q}) = \begin{bmatrix} -\mathbf{q}^T \\ q_r I_3 - [\mathbf{q}]_{\times} \end{bmatrix}, \quad (\text{C.14})$$

where I_n is the $n \times n$ identity matrix; $[\cdot]_{\times} : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 3}$ is the cross-product matrix, a matrix such that, given $v, w \in \mathbb{R}^3$,

$$v \times w = [v]_{\times} w. \quad (\text{C.15})$$

This is, for $v = [v_1 \ v_2 \ v_3]^T$,

$$[v]_{\times} = \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix}. \quad (\text{C.16})$$

¹In this thesis, the non-Hamiltonian quaternion product, indicated as \otimes in Markley et al. (2014), is not used at all.

Some properties of these matrices are that $[\mathbf{q}^*]_{\odot}^+ = ([\mathbf{q}]_{\odot}^+)^T$ and $[\mathbf{q}^*]_{\odot}^- = ([\mathbf{q}]_{\odot}^-)^T$. Note that $([v]_{\times})^T = -[v]_{\times}$.

The conjugate of a quaternion is obtained by reversing the sign of the hyper-complex part only:

$$\mathbf{q}^* = [1 \quad -i \quad -j \quad -k] \mathbf{q}. \quad (\text{C.17})$$

With this, the absolute value of the quaternion can be obtained as

$$|\mathbf{q}| = \sqrt{\mathbf{q} \odot \mathbf{q}^*} = \sqrt{q_r^2 + q_i^2 + q_j^2 + q_k^2}. \quad (\text{C.18})$$

The inverse is such that

$$\mathbf{q} \odot \mathbf{q}^{-1} = \mathbf{q}^{-1} \odot \mathbf{q} = 1. \quad (\text{C.19})$$

Using the above definitions, it is trivial to show that

$$\mathbf{q}^{-1} = \frac{\mathbf{q}^*}{|\mathbf{q}|^2}. \quad (\text{C.20})$$

D. Extended Attitude Concepts

This appendix gathers attitude definitions and formulas used throughout this thesis. Some of these formulas vary slightly (w.r.t. other sources in the literature) depending on the conventions used. In these cases, the key steps necessary to reach the particular variants used in these dissertation have been reproduced below, too.

D.1. Attitude Representations

Attitude is a parametrization with three Degrees of Freedom (DoFs) that can be represented in different formats, including quaternions, Euler-angle representation, Direction Cosine Matrix (DCM) and others. This section describes briefly these mathematical entities, as well as the kinematics and dynamics of the attitude of a rigid solid. Attitude in this thesis is represented by the transformation of a vector in some frame A to some other frame C . This is, it is $f_A^C : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that, given any vector $v \in \mathbb{R}^3$,

$$v^C = f_A^C (v^A) . \quad (\text{D.1})$$

For vectors (*e.g.* v) the super-script determines in which frame the given vector is represented. For attitude representations (*e.g.* f), the transformation goes from the frame in the subscript to the frame in the superscript. These frame-specification rules are observed throughout the thesis wherever frames and frame transformations are involved. They shall be omitted wherever they are implicit from context, to avoid unnecessary clutter.

Refer to *e.g.* Shuster et al. (1993) and Solà (2017) for a formal proof of the equations shown below. However, note that the conventions here do not match those taken therein.

D.1.1. Attitude as a Quaternion

The standard attitude representation used in this thesis is the quaternion. Whenever attitude states need to be propagated over time, the quaternion form is used.

The quaternions needed to represent attitude are only those at the surface of a unit hypersphere in \mathbb{H} with positive real part. This is the subset

$$\mathbb{A} = \{ \mathbf{q} \in \mathbb{H} \mid \Re(\mathbf{q}) \geq 0 \wedge |\mathbf{q}| = 1 \} . \quad (\text{D.2})$$

Given two 3-dimensional coordinate frames A and C , related to each other by a rotation, the operation that transforms a vector $v \in \mathbb{R}^3$ from frame A (denoted as v^A) to frame C (v^C) can be expressed as

$$\mathbf{v}^C = (\mathbf{q}_A^C)^* \odot \mathbf{v}^A \odot \mathbf{q}_A^C , \quad (\text{D.3})$$

D. Extended Attitude Concepts

where $\mathbf{q}_A^C \in \mathbb{A}$.

Although quaternions are the preferred attitude representation in this thesis, other formats have been used at specific points throughout this document. The following subsections describe them.

D.1.2. Attitude as an Euler Axis-Angle Pair

An attitude quaternion is related to the Euler-angle pair by

$$\mathbf{q}_A^C = \left[\cos\left(\frac{\psi}{2}\right) \quad e^T \sin\left(\frac{\psi}{2}\right) \right]^T, \quad (\text{D.4})$$

where $e \in \{u \in \mathbb{R}^3 \mid \|u\| = 1\}$ is the common rotation axis around which frame A needs to rotate ψ radians to align itself with frame C . From this, it is trivial that $\mathbf{q}_C^A = (\mathbf{q}_A^C)^*$.

D.1.3. Attitude as a Direction Cosine Matrix

The same frame transformation can be expressed using a DCM. This is, the matrix R_A^C that transforms some vector from frame A to frame C :

$$v^C = R_A^C v^A. \quad (\text{D.5})$$

It is related to the attitude quaternion as:

$$R_A^C = R(\mathbf{q}_A^C) = \Theta^T(\mathbf{q}_A^C) \Xi(\mathbf{q}_A^C) = \Xi^T(\mathbf{q}_C^A) \Theta(\mathbf{q}_C^A). \quad (\text{D.6})$$

This follows from the quaternion multiplication rules in Appendix C, where Θ and Ξ are defined (Equations C.13 and C.14).

D.1.4. Attitude as Generalized Rodrigues Parameters

One last attitude representation used in this dissertation is the Generalized Rodrigues Parameters (GRP), obtained from the quaternion with the map $p_{a,f} : \mathbb{A} \rightarrow \mathbb{R}^3$, $\forall a \in \mathbb{R}_{\geq 0}$ and $\forall f \in \mathbb{R}_{>0}$, defined as (Dianetti et al., 2018)

$$p_{a,f}(\mathbf{q}) = f \frac{\mathbf{q}}{a + q_r}. \quad (\text{D.7})$$

The quaternion can be recovered with the inverse map $p_{a,f}^{-1} : \mathbb{R}^3 \rightarrow \mathbb{A}$, whose equations are

$$q_r = \frac{-a \|p\|^2 + f \sqrt{f^2 + (1 - a^2) \|p\|^2}}{f^2 + \|p\|^2} \quad \text{and} \quad \mathbf{q} = \frac{a + q_r}{f} p. \quad (\text{D.8})$$

D.2. Attitude Quaternion Derivative

The attitude quaternion time derivative has a well known equation as a function of the quaternion itself and the angular rate (Solà, 2017). However, this thesis uses an uncommon convention for attitude: the quaternion transforms from body to inertial frame, instead of the more common direction from inertial to body frame. Thus, this section provides the proof of the attitude quaternion time derivative for this particular convention, based on Solà (2017).

Theorem 4. *The time-derivative of attitude quaternion $\mathbf{q}_B^I \in \mathbb{A}$ is*

$$\dot{\mathbf{q}}_B^I = -\frac{1}{2}\boldsymbol{\omega}_{B/I}^B \odot \mathbf{q}_B^I, \quad (\text{D.9})$$

where $\boldsymbol{\omega}_{B/I}^B \in \mathbb{R}^3$ is angular velocity of frame B relative to frame I , expressed in B coordinates.

Proof. Let

$$\mathbf{q}_B^{B^+} = \begin{bmatrix} \cos\left(\frac{\Delta\psi}{2}\right) \\ e \sin\left(\frac{\Delta\psi}{2}\right) \end{bmatrix} = \begin{bmatrix} 1 \\ e \frac{\Delta\psi}{2} \end{bmatrix} + O(\Delta\psi^2) \quad (\text{D.10})$$

be the quaternion that transforms from frame B at some time t to frame B^+ at some time $t + \Delta t$, where $e \in \mathbb{R}^3$ and $\Delta\psi$ are the axis-angle Euler pair representation of $\mathbf{q}_B^{B^+}$.

One can obtain the attitude at time $t + \Delta t$ from the one at time t by composing rotations (Equation 3.4):

$$\mathbf{q}_{B^+}^I = \left(\mathbf{q}_B^{B^+}\right)^* \odot \mathbf{q}_B^I. \quad (\text{D.11})$$

The limit definition of the derivative is

$$\begin{aligned} \dot{\mathbf{q}}_B^I &= \lim_{\Delta t \rightarrow 0} \frac{\left(\mathbf{q}_B^{B^+}\right)^* \odot \mathbf{q}_B^I - \mathbf{q}_B^I}{\Delta t} = \\ &= \lim_{\Delta t \rightarrow 0} \frac{\cancel{(1-1)}\mathbf{q}_B^I - \left(e \frac{\Delta\psi}{2}\right) \odot \mathbf{q}_B^I + (1 + \mathbf{i} + \mathbf{j} + \mathbf{k})O(\Delta\psi^2)}{\Delta t}. \end{aligned} \quad (\text{D.12})$$

From Solà (2017),

$$\boldsymbol{\omega}_{B/I}^B := \lim_{\Delta t \rightarrow 0} \frac{e \Delta\psi}{\Delta t}. \quad (\text{D.13})$$

Thus, if the body inertia is non-zero, the angular velocity is finite, so the above limit resolves to

$$\dot{\mathbf{q}}_B^I = -\frac{1}{2}\boldsymbol{\omega}_{B/I}^B \odot \mathbf{q}_B^I + \lim_{\Delta t \rightarrow 0} \cancel{(1 + \mathbf{i} + \mathbf{j} + \mathbf{k})O(\Delta t^2 \|\boldsymbol{\omega}_{B/I}^B\|^2)}. \quad (\text{D.14})$$

□

D.3. Attitude Quaternion Given Two Vectors

Let $v \in \{x \in \mathbb{R}^3 \mid \|x\| = 1\}$ be a unit vector in some reference frame. Let $w \in \mathbb{R}^3 \in \{x \in \mathbb{R}^3 \mid \|x\| = 1\}$ be some other unit vector. The goal is to find some frame transformation f such that $w = f(v)$. This can be done by finding some attitude quaternion q that satisfies

$$w = q^* \odot v \odot q. \quad (\text{D.15})$$

There exist infinite solutions, because not enough information is available to determine the orientation of the axis in the new frame that are orthogonal to w . Nonetheless, it is just necessary to obtain one such transformation that complies with the above, it does not need to be unique.

Theorem 5. *One such transformation can be realized with the quaternion whose vector form is*

$$q = \frac{1}{\|v + w\|} \begin{bmatrix} 1 + v^T w \\ -v \times w \end{bmatrix}, \quad (\text{D.16})$$

as long as $v \neq -w$, which would result in a 0/0 indetermination.

Proof. The proof reproduced here is a variant from Eisele (2008), where it is shown for vector rotations, instead of frame transformations.

First, the dot product $v^T w$ is equal to $\cos \psi$, being $\psi = \angle(v, w)$ the angle between the two vectors. Using trigonometric transformations, $1 + \cos \psi$ becomes $2 \cos^2(\psi/2)$.

Second, the norm of the cross product $v \times w$ is equal to $|\sin \psi|$. Thus, it can be rewritten as

$$v \times w = \frac{v \times w}{\|v \times w\|} |\sin \psi|. \quad (\text{D.17})$$

Using again trigonometric transformations, this becomes

$$v \times w = 2 \frac{v \times w}{\|v \times w\|} \cos\left(\frac{\psi}{2}\right) \sin\left(\frac{\psi}{2}\right). \quad (\text{D.18})$$

Note how the absolute value disappears because $\psi/2$ can be minimum 0 and maximum $\pi/2$ —within this range both sin and cos return a non-negative value.

Third, the Euclidean norm of the sum $v + w$ can be expanded to

$$\begin{aligned} \|v + w\| &= \sqrt{(u_1 + v_1)^2 + (u_2 + v_2)^2 + (u_3 + v_3)^2} = \sqrt{\|v\|^2 + \|w\|^2 + 2v^T w} = \\ &= \sqrt{2 + 2v^T w} = 2 \cos\left(\frac{\psi}{2}\right). \end{aligned} \quad (\text{D.19})$$

The negative root would make the real part of the quaternion become negative later on, which is not allowed as per the definition of attitude quaternion given in this thesis. Thus, the positive one is taken.

Putting these three statements together into Equation D.16 yields

$$q = \begin{bmatrix} \cos\left(\frac{\psi}{2}\right) \\ -\frac{v \times w}{\|v \times w\|} \sin\left(\frac{\psi}{2}\right) \end{bmatrix}. \quad (\text{D.20})$$

Finally, by comparing this with Equation D.4, it can be seen that $e = -v \times w / \|v \times w\|$ is the axis around which the original frame needs to rotate ψ degrees for v to become w : it is orthogonal to the plane that contains both vectors, and indicates that the frame needs to rotate opposite to the direction from v to w per the right-hand rule. In other words, ψ and e form the Euler axis-angle pair that represent the desired transformation. \square

Corollary 2. *For the particular case when $v = -w$, $\Re(\mathbf{q}) = 0$ while $\Im(\mathbf{q})$ can be any unit vector orthogonal to v .*

Proof. The proof is trivial from Equation D.4. Since $\psi = \pi$, $\cos(\psi/2) = 0$ and $\sin(\psi/2) = 1$. Then, e must be any axis orthogonal to v so that the frame rotates exactly 180° and w becomes antiparallel to v . \square

D.4. Generalized Rodrigues Parameters Derivations

D.4.1. Taylor Expansion of the GRP

Given

$$p_{a,f} = f \frac{\mathbf{q}}{a + q_r}, \quad (\text{D.21})$$

Taylor expansion around ψ (see Equation D.4) yields

$$p_{a,f} = f \frac{e\psi}{2(a+1)} + O(|\psi|^3). \quad (\text{D.22})$$

Thus, if $f = 2(a+1)$, it cancels with the denominator and makes p equal to $e\psi$ up to second order approximation.

D.4.2. Change of the Reference Quaternion

Let $\tilde{\mathbf{q}}$ be a reference attitude quaternion and two other quaternions

$$\mathbf{q} = \delta\mathbf{q}^* \odot \tilde{\mathbf{q}} \quad (\text{D.23})$$

and

$$\tilde{\mathbf{q}}' = (\delta\tilde{\mathbf{q}}')^* \odot \tilde{\mathbf{q}} \quad (\text{D.24})$$

whose GRP error representations are, then, $\delta p = p(\delta\mathbf{q})$ and $\delta\tilde{p}' = p(\delta\tilde{\mathbf{q}}')$. The goal is to find $\delta p'$ that represents the same quaternion \mathbf{q} , but using $\tilde{\mathbf{q}}'$ as reference quaternion. This is, it is necessary to find $\delta p' = p(\delta\mathbf{q}')$ such that

$$\delta\mathbf{q}^* \odot \tilde{\mathbf{q}} = (\delta\mathbf{q}')^* \odot \tilde{\mathbf{q}}' \Rightarrow \delta\mathbf{q}^* \odot \tilde{\mathbf{q}} = (\delta\mathbf{q}')^* \odot (\delta\tilde{\mathbf{q}}')^* \odot \tilde{\mathbf{q}}, \quad (\text{D.25})$$

which yields

$$\delta\mathbf{q}' = (\delta\tilde{\mathbf{q}}')^* \odot \delta\mathbf{q} \Rightarrow \delta p' = p \left((p^{-1}(\delta\tilde{p}'))^* \odot p^{-1}(\delta p) \right). \quad (\text{D.26})$$

D. Extended Attitude Concepts

This expression can be expanded within first order as

$$\delta p' (\delta \tilde{p}', \delta p) = \delta p'(0,0) + \left. \frac{\partial \delta p'}{\partial \delta \tilde{p}'} \right|_{0,0} (\delta \tilde{p}' - 0) + \left. \frac{\partial \delta p'}{\partial \delta p} \right|_{0,0} (\delta p - 0) + O\left(\|\delta \tilde{p}'\|^2 + \|\delta p\|^2\right). \quad (\text{D.27})$$

The Jacobians involved can be obtained by applying the derivative chain rule, which yields

$$\frac{\partial \delta p'}{\partial \delta \tilde{p}'} = \frac{\partial \delta p'}{\partial \delta q'} \frac{\partial \delta q'}{\partial \delta \tilde{q}'} \frac{\partial \delta \tilde{q}'}{\partial \delta \tilde{p}'} \quad (\text{D.28})$$

and

$$\frac{\partial \delta p'}{\partial \delta p} = \frac{\partial \delta p'}{\partial \delta q'} \frac{\partial \delta q'}{\partial \delta q} \frac{\partial \delta q}{\partial \delta p}. \quad (\text{D.29})$$

Starting with the first term in both derivatives, which is the same, one can compute it by differentiating Equation D.21 with respect to q , as a function of q , which yields

$$\frac{\partial p}{\partial q}(q) = \frac{1}{a + q_r} \begin{bmatrix} -p(q) & f I_3 \end{bmatrix}, \quad (\text{D.30})$$

where I_i is the identity matrix in $\mathbb{R}^{i \times i}$. The last term can likewise be obtained by finding the inverse derivative, by differentiating Equation D.8:

$$\frac{\partial q}{\partial p}(p) = \begin{bmatrix} g(p) \\ \frac{1}{f} p g(p) + \frac{a + q_r(p)}{f} I_3 \end{bmatrix}, \quad (\text{D.31})$$

where

$$g(p) = \frac{\partial q_r}{\partial p}(p) = \left(\frac{-2a + f(1 - a^2) \left(f^2 + (1 - a^2) \|p\|^2 \right)^{-\frac{1}{2}}}{f^2 + \|p\|^2} - 2 \frac{-a \|p\|^2 + f \sqrt{f^2 + (1 - a^2) \|p\|^2}}{\left(f^2 + \|p\|^2 \right)^2} \right) p^T. \quad (\text{D.32})$$

Finally, the middle terms are

$$\frac{\partial \delta q'}{\partial \delta \tilde{q}'} = \begin{bmatrix} \delta q_r & \delta \mathbf{q}^T \\ \delta \mathbf{q} & -\delta q_r I_3 + [\delta \mathbf{q}]_{\times} \end{bmatrix} \quad (\text{D.33})$$

and

$$\frac{\partial \delta q'}{\partial \delta q} = \begin{bmatrix} \delta \tilde{q}_r & \delta \tilde{\mathbf{q}}^T \\ -\delta \tilde{\mathbf{q}} & \delta \tilde{q}_r I_3 - [\delta \tilde{\mathbf{q}}]_{\times} \end{bmatrix}, \quad (\text{D.34})$$

which come from differentiating the multiplication rules defined by Equation C.7, applied to the vector form of $\delta \mathbf{q}'$ as defined in Equation D.26.

Evaluating these derivatives at $\delta \tilde{p}' = 0$ and $\delta p = 0$ and substituting in Equation D.27 results in

$$\delta p'(0,0) = 0, \quad (\text{D.35a})$$

D.4. Generalized Rodrigues Parameters Derivations

$$\left. \frac{\partial \delta p'}{\partial \delta \tilde{p}'} \right|_{0,0} = -I_3 \text{ and} \tag{D.35b}$$

$$\left. \frac{\partial \delta p'}{\partial \delta p} \right|_{0,0} = I_3; \tag{D.35c}$$

thus

$$\delta p' = \delta p - \delta \tilde{p}' + O\left(\|\delta \tilde{p}'\|^2 + \|\delta p\|^2\right). \tag{D.36}$$

Note that this result is reached regardless of the values of a and f .

E. Three Dimensional Objects

E.1. Extended Gaussian Images

From Horn (1984), an Extended Gaussian Image (EGI), is

Definition 8. Let the a convex 3D object be represented by the finite volume $O \in \mathbb{R}^3$, whose outer closed surface is $\delta O \subset O$. Let S be a Gaussian sphere $S = \{n \in \mathbb{R}^3 \mid \|n\| = 1\}$. Given some point $x \in \delta O$, whose surface unit normal is n , one can map it to the point on S that has the same surface normal n , which is in fact n itself: $n = S(x)$. If $S(x)$ is invertible, any point on δO can be obtained as $x = \delta O(n) = S^{-1}(n)$.

The **Extended Gaussian Image (EGI)** of δO is the variation of area on δO with the area of S . This is

$$\text{EGI}(n) = \frac{d}{dS}(\delta O). \quad (\text{E.1})$$

It can be understood as a measure of area density on the surface of δO , and is the inverse of the Gaussian curvature of δO .

Minkowski's theorem (Minkowski, 1897) states that the EGI exists if δO is convex. In particular:

Theorem 6. Minkowski's theorem states that if

$$\int_S \frac{d}{dS}(\delta O) n dS = 0, \quad (\text{E.2})$$

then the EGI is a unique representation of a convex closed surface. For a polyhedron, this is

$$\sum_{i=1}^{N_f} A_i n_i = 0. \quad (\text{E.3})$$

Proof. Refer to Horn (1984) and Little (1985). □

The general definition of EGI can be particularized to a polyhedron. From Horn (1984) and Little (1985):

Lemma 1. Let $\mathcal{P} \in \mathbb{R}^3$ be a convex polyhedron with N_f flat facets. Each facet i has a certain area A_i and normal n_i , for $i \in 1 \dots N_f$. The EGI of (the surface of) \mathcal{P} can be expressed as

$$\text{EGI}(n) = \sum_{i=1}^{N_f} A_i \delta(\|n - n_i\|), \quad (\text{E.4})$$

E. Three Dimensional Objects

where $\delta : \mathbb{R} \rightarrow \mathbb{R}$ is the impulse function (or Dirac delta). A compact alternative to represent this equation is to use the set of vectors

$$\text{EGI}_{\mathcal{P}} = \{A_i n_i \mid i \in 1 \dots N_f\} . \quad (\text{E.5})$$

Proof. The proof is trivial. Let $\Omega \subset S$ be some arbitrary region such that $n_i \notin \Omega$, for all $i \in 1 \dots N_f$. Applying the fundamental theorem of calculus to Equation E.1, the integral of the EGI of \mathcal{P} along Ω must be 0. This is,

$$0 = \int_{\Omega} \frac{d}{dS}(\delta O) dS , \quad (\text{E.6})$$

This is because there is no surface of \mathcal{P} orthogonal to any of the n directions in Ω .

Next, given any facet $i \in 1 \dots N_f$, let $\Omega_i \subset S$ be some other region such that $n_i \in \Omega_i$, but $n_j \notin \Omega_i$, for all $j \in 1 \dots N_f$ and $i \neq j$. Note that Ω_i is only realizable if N_f is finite—*i.e.* there is always some $n \in \Omega$ on the shortest arch that joins n_i and n_j . Applying again the fundamental theorem of calculus to Equation E.1, the integral of the EGI along any such region Ω_i must equal the area of facet i . This is

$$A_i = \int_{\Omega_i} \frac{d}{dS}(\delta O) dS . \quad (\text{E.7})$$

The above two cases must be true for any possible realization of Ω and Ω_i , for all $i \in 1 \dots N_f$. Furthermore, the union of all realizable Ω and Ω_i cover the whole of S . Thus, the only function that satisfies these conditions is the sum of impulses from Equation E.4. \square

E.2. Sphere Tessellation

This section addresses the fact that the exact set of n_i directions that represent the true polyhedron is not known. Instead, it is necessary to approximate it.

The EGI of a polyhedron can be approximated with a fixed set of N_a directions $\{n_j \mid j \in 1 \dots N_a\}$ on the unit sphere, which may not necessarily coincide with the set of the polyhedron facet normals. This is done by associating the out-facing normals of \mathcal{P} to the closest of each of these directions. As a set, this is

$$\widehat{\text{EGI}}_{\mathcal{P}} = \left\{ \sum_{i=1}^{N_f} \varrho_j(n_i) A_i n_j \mid j \in 1 \dots N_a \right\} , \quad (\text{E.8})$$

where $\varrho_j(n)$ is 1 if n_j is the closest approximation to $n/\|n\|$ and 0 otherwise. Depending on how $\varrho_j(n)$ is implemented, different definitions of *closest* may be used. Whatever the implementation, it is necessary that $\sum_j \varrho_j(n) = 1$ for any $n \in \mathbb{R}^3$.

There are multiple ways to implement $\varrho_j(n)$. An effective way to decide a good set of $\{n_j\}$ directions is to work with spherical tessellations. This is, one subdivides the surface S with non-intersecting geometrical shapes, a.k.a. cells, so that they cover it as homogeneously as possible. According to Horn (1984), one can define homogeneity in the sense that:

1. each cell is as close in shape to a circle as possible,
2. any two cells have an area as close to each other as possible,
3. the distance between any two adjacent cells varies as little as possible.

It is impossible to realize this for the surface of a sphere (Little, 1985). However, one can use tessellations based on regular polyhedrons, such as the icosahedron or the dodecahedron, and subsequent geodesic polyhedrons generated from these. Regardless of the tessellation chosen, $\varrho_j(n)$ can then be implemented by setting it to 1 if, and only if, n falls within cell j ; n_j would be then the centroid of cell j .

Another option for $\varrho_j(n)$ could be, given a set of $\{n_j\}$ as homogeneously distributed as possible,

$$\varrho_j(n) = \begin{cases} 1 & \text{if } j = \arg \min_j \|n - n_j\|, \\ 0 & \text{else.} \end{cases} \quad (\text{E.9})$$

Whatever the implementation of $\varrho_j(n)$, if the distribution of cells (or $\{n_j\}$) is homogeneous enough, when $N_a \rightarrow \infty$, then $\widehat{\text{EGI}}_{\mathcal{P}} \rightarrow \text{EGI}_{\mathcal{P}}$. Thus, at the end it is important to choose a method that is computationally not too expensive. For this particular reason, the tessellation used to generate an EGI in this dissertation is the so called quad-cube tessellation (White et al., 1992). It is the one used in Fan et al. (2019), Fan et al. (2020), Fan et al. (2021), and Friedman et al. (2019), too. To construct the quad-cube tessellation, each edge of the cube becomes a great circle arc with 90° span on the sphere, and each of its faces becomes a spherical square delimited by four of these great circle arcs. Each cube face is then subdivided into smaller cells that are as similar as possible to each other (in the homogeneity sense above).

To do so, first, each of the cube faces is mapped onto the domain $\gamma \times \beta$, for all $\gamma, \beta \in [-45^\circ, 45^\circ]$. Figure E.1a illustrates this for the cube face that intersects the x axis—the x , y and z axes here represent the body frame. In the figure, cube edges transformed to great circle arcs are shown as solid arcs—those of the example face are shown in thicker style. The dashed lines are great circle arcs that indicate constant γ (arcs that intersect with z) and β (those that intersect with y). As shown in the figure, the γ coordinate corresponds to the more traditional azimuth or longitude (φ), while spherical trigonometry can be used to express elevation or latitude (ϑ) as a function of both γ and β .

This can be achieved by following the steps shown in Figure E.2. This figure highlights the spherical triangles that relate γ , β and ϑ . First from Figure E.2a, the cosine theorem for spherical trigonometry applied on β yields

$$\cos \beta = \cos \frac{\pi}{2} \cos \frac{\pi}{2} + \sin \frac{\pi}{2} \sin \frac{\pi}{2} \cos b, \quad (\text{E.10})$$

while the sine theorem gives

$$\frac{\sin \beta}{\sin b} = \frac{\sin \frac{\pi}{2}}{\sin \frac{\pi}{2}} = 1. \quad (\text{E.11})$$

This is,

$$b = \beta. \quad (\text{E.12})$$

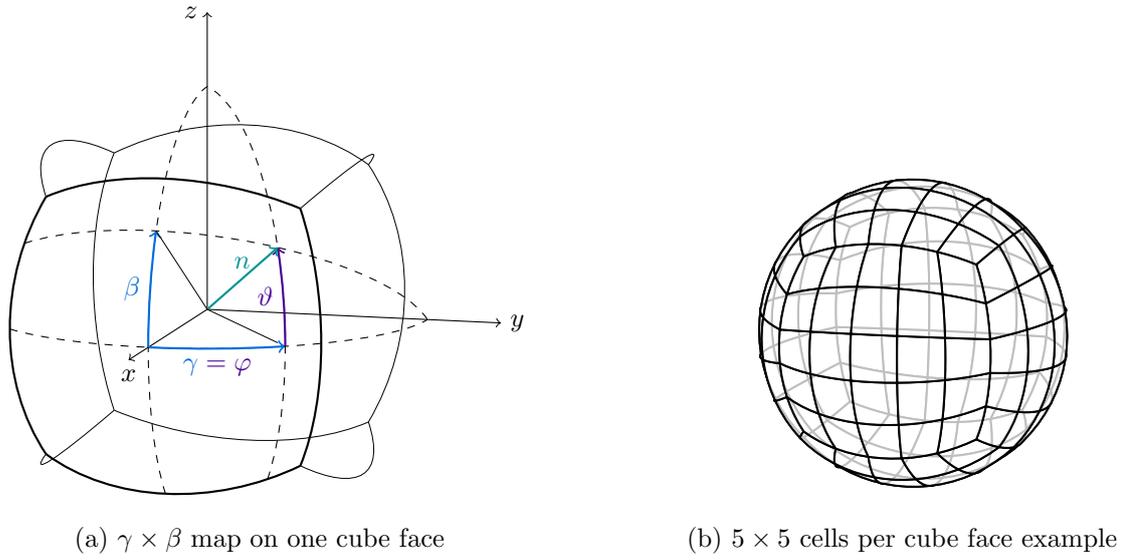


Figure E.1.: The quad-cube tessellation

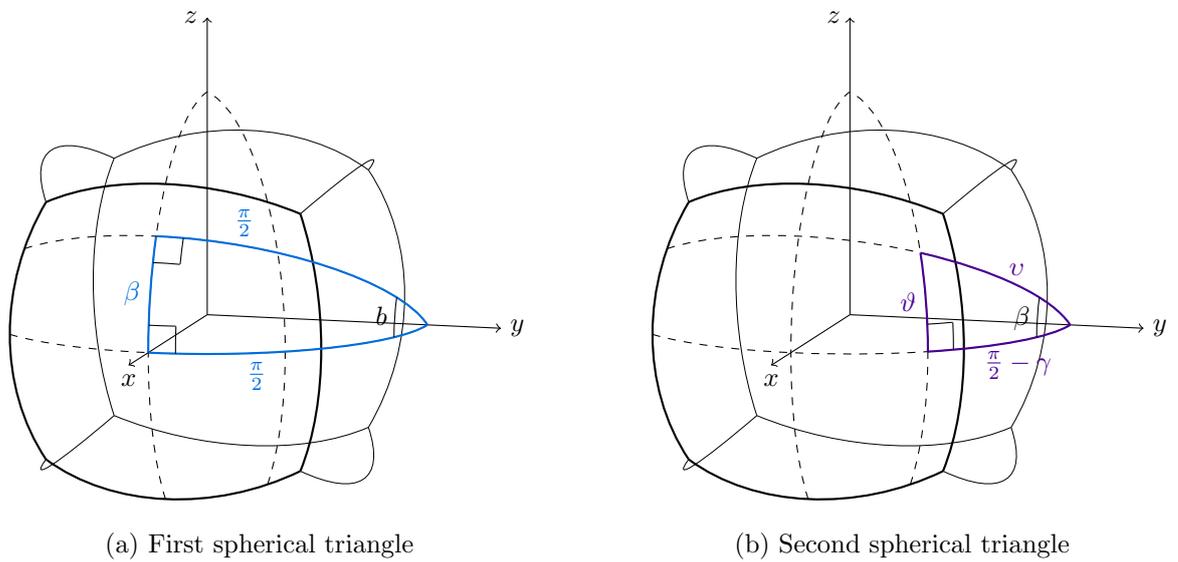


Figure E.2.: Spherical triangles involved in the quad-cube tessellation coordinate system

Using this result, the sine theorem applied to Figure E.2b gives

$$\frac{\sin v}{\sin \frac{\pi}{2}} = \frac{\sin \vartheta}{\sin \beta}. \quad (\text{E.13})$$

Then, one applies the spherical cosine theorem on Figure E.2b twice, obtaining

$$\cos v = \sin \gamma \cos \vartheta + \cos \gamma \sin \vartheta \cos \frac{\pi}{2} \quad (\text{E.14})$$

and

$$\cos \vartheta = \sin \gamma \cos v + \cos \gamma \sin v \cos \beta. \quad (\text{E.15})$$

Note that the equalities $\cos(\pi/2 - \gamma) = \sin \gamma$ and $\sin(\pi/2 - \gamma) = \cos \gamma$ have been used.

Now, substituting Equations E.13 and E.14 into Equation E.15 results in

$$\cos \vartheta = \sin^2 \gamma \cos \vartheta + \cos \gamma \frac{\sin \vartheta}{\sin \beta} \cos \beta. \quad (\text{E.16})$$

Then, dividing by $\cos \vartheta$ and isolating ϑ , one gets

$$\tan \vartheta = \frac{1 - \sin^2 \gamma}{\cos \gamma} \tan \beta. \quad (\text{E.17})$$

Finally, simplifying with $\cos^2 \gamma = 1 - \sin^2 \gamma$ gives the desired result:

$$\vartheta(\gamma, \beta) = \tan^{-1}(\cos \gamma \tan \beta). \quad (\text{E.18})$$

The Cartesian position n on the unit sphere is, then,

$$n(\gamma, \beta) = \begin{bmatrix} \cos \varphi \cos \vartheta \\ \sin \varphi \cos \vartheta \\ \sin \vartheta \end{bmatrix} = \frac{1}{\sqrt{\cos^2 \gamma \tan^2 \beta + 1}} \begin{bmatrix} \cos \gamma \\ \sin \gamma \\ \cos \gamma \tan \beta \end{bmatrix}. \quad (\text{E.19})$$

By combining 90° rotations around the x , y and z axes, the same $\gamma \times \beta$ map can be applied to each of the other five cube faces. The grid is then created by intersecting great circle arcs at regular γ and β intervals. Figure E.1b shows an example of this tessellation for the case where each cube face is subdivided into a 5×5 grid of cells, which are delimited by great circle arcs at $\gamma, \beta \in \{-45^\circ, -27^\circ, -9^\circ, 9^\circ, 27^\circ, 45^\circ\}$. Generalizing, a quad-cube tessellation whose faces are divided into $d \times d$ cells each has a total of $N_a = 6d^2$ cells.

Furthermore, when using the quad-cube tessellation to build an EGI, the direction assigned to each cell is the centre of the cell in the $\gamma \times \beta$ domain. This is, if a cell is delimited by all the possible combinations of $\gamma_1 < \gamma_2$, $\beta_1 < \beta_2$, the direction it represents is $\left(\frac{\gamma_1 + \gamma_2}{2}, \frac{\beta_1 + \beta_2}{2}\right)$. Following this notion, ϱ_j is defined as

$$\varrho_j(n) = \begin{cases} 1 & \text{if } \gamma_1 \leq \gamma(n) < \gamma_2 \text{ and } \beta_1 \leq \beta(n) < \beta_2 \text{ falls within cell } j, \\ 0 & \text{else.} \end{cases} \quad (\text{E.20})$$

It is trivial to obtain $(\gamma, \beta)(n)$ by reversing Equations E.18 and E.19.

F. The Unscented Transform

This section provides a compact reproduction of the Unscented Transform (UT) equations, for the reader's reference (Julier et al., 2004).

Let $X \in \mathbb{R}^n$ be a random variable that follows a statistical distribution whose first (mean) and second (covariance) moments are μ_x and P_x , respectively. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a map that applies a non-linear transformation to X , such that a new random variable $Y \in \mathbb{R}^m$ can be obtained as

$$Y = f(X) + V, \quad (\text{F.1})$$

where $V \in \mathbb{R}^m$ is the random variable that represents the uncertainty associated to the physical model represented by f , a.k.a as model noise. Its associated distribution has 0 mean and covariance P_v .

The goal of the UT is to obtain the mean and covariance of the distribution associated to Y . It uses $\alpha, \beta, \kappa \in \mathbb{R}$ as tunable parameters.

First, $2n + 1$ sigma points are computed with

$$\mathcal{X}(0) = \mu_x, \quad (\text{F.2a})$$

$$\mathcal{X}(i) = \mathcal{X}(0) + \sqrt{n + \lambda} S_i \text{ and} \quad (\text{F.2b})$$

$$\mathcal{X}(i + n) = \mathcal{X}(0) - \sqrt{n + \lambda} S_i, \quad (\text{F.2c})$$

for all $i \in 1 \dots n$, where S_i is the i -th column of the Cholesky decomposition of $P_x = SS^T$ and

$$\lambda = \alpha^2(n + \kappa) - n. \quad (\text{F.3})$$

Then, the sigma cloud is transformed under the map f , obtaining the transformed sigma cloud

$$\mathcal{Y}(i) = f(\mathcal{X}(i)), \forall i \in 0 \dots 2n. \quad (\text{F.4})$$

Finally, the mean of the transformed distribution is estimated as

$$\hat{\mu}_y = \sum_{i=0}^{2n} w_m(i) \mathcal{Y}(i) \quad (\text{F.5})$$

and its covariance as

$$P_y = \sum_{i=0}^{2n} w_c(i) (\mathcal{Y}(i) - \hat{\mu}_y) (\mathcal{Y}(i) - \hat{\mu}_y)^T + P_v. \quad (\text{F.6})$$

F. The Unscented Transform

The weights w_m and w_c are obtained as

$$w_m(0) = \frac{\lambda}{n + \lambda}, w_c(0) = w_m(0) + 1 - \alpha^2 + \beta, w_m(i) = w_c(i) = \frac{1}{2(n + \lambda)}, \quad (\text{F.7})$$

for all $i \in 1 \dots 2n$.

Likewise, the cross covariance of X and Y may be estimated using the UT as

$$\text{cov}(X, Y) = \sum_{i=0}^{2n} w_c(i) (\mathcal{X}(i) - \mu_x) (\mathcal{Y}(i) - \hat{\mu}_y)^T. \quad (\text{F.8})$$

The UT preserves the first and second statistical moments of any transformed distribution. Additionally, proper tuning of the β parameter can guarantee the third and fourth moments when dealing with Gaussian distributions. The parameters α and κ control the spread of the sigma points.