

Uncovering Peculiar Rainbows

Automated Detection and Characterisation of Luminous Bands in Saturn's E Ring

Arnaud Mathieu

Uncovering Peculiar Rainbows

Automated Detection and Characterisation of Luminous Bands
in Saturn's E Ring

by

Arnaud Mathieu

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday June 4, 2026.

Student number: 4860136

Project duration: 7 October 2025 – 4 June 2026

Thesis committee: Dr. S. Paardekooper, TU Delft, Chair
Dr. J.J.D. Loicq, TU Delft, Responsible supervisor
Dr. S.M. Cazaux, TU Delft, Supervisor
Dr.ir. E.J.O. Schrama, TU Delft, External examiner

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

Where to begin. Maybe the start? I would like to thank Neil deGrasse Tyson. Dropping the single most inspiring documentary, *Cosmos: A Spacetime Odyssey*, right when I was starting to form my interests and hobbies was something he really did not have to do. The very first book I ever read that made me the space nerd I am today also deserves to be cited in this thesis, in my opinion. So, if you would like to know where the foundations of this document come from, I refer to *Ons zonnestelsel* (Éditions Atlas 2006).

On a more serious note, thanking my family and all of its extensions for all the support throughout the years, not only during the making of this thesis but also all the way from the start of this adventure abroad, cannot be done enough, nor put into words. Thank you for everything, and I hope you are eager to read what it all has led to!

One of the main reasons that I enjoyed working on this thesis, and I have shared this with my surroundings often, is that I was supervised so well during this period. The biggest *merci* to my two thesis parents, Stéphanie and Jérôme, for keeping me motivated week after week, for providing insights into seemingly impossible topics to understand, and to turn my nightmares into easy-to-follow to-dos.

The only group of people left to thank are my dearest friends. Thank you, Loes and Frederick. To Lore, Evert, and of course Kasper: the fact that my office space was also the living room and kitchen made work-life balance a difficult exercise, but you made sure it all worked out. Warm hugs. Thank you, Niels, for the years of friendship, and also for the indirect opportunity you gave me by continuing your thesis topic. Please give me some time to adapt back to calling you Niels again instead of Rubbrecht (et al.). Thank you, Oscar and Kato, for providing me a second home here in Delft purely out of *gezelligheid* and the goodness of your hearts.

Arnaud Mathieu
Delft, 4 June 2026

Bibliography of the Preface

Éditions Atlas (2006). *Ons zonnestelsel*. Atlas voor de jeugd: Wetenschap en ontdekkingen. Translation by De Vertaalfabriek. ISBN 978-2-8302-2313-2. Amsterdam: Éditions Atlas / Literary Rights International.

Contents

About This Document	3
I Literature Study	5
1 Historical Context and Discovery of Enceladus	6
2 Formation and Evolution Models of the Saturnian Moons	8
2.1 Contextualising regular moon formation	9
2.2 Orbital evolution, tidal dissipation, and age constraints	10
2.3 Detailed scenarios for Saturn’s regular moons	12
2.3.1 Late formation via disruption and re-accretion	12
2.3.2 Formation from a circumplanetary disk (CPD)	12
2.3.3 Formation/migration from a massive ring (Tidal Disk Model)	14
2.4 Origin of irregular moons	16
2.5 A tentative dual formation scenario for Saturn’s system	17
3 Plumes and the E Ring	19
3.1 Time variability of the Enceladus plumes	19
3.1.1 Diurnal forcing (1.37 days)	19
3.1.2 Multi-year modulation (~ 4 and ~ 11 years)	19
3.1.3 Week-to-month variability (stochastic jets)	21
3.1.4 Dust-gas co-variation and velocities	21
3.1.5 Implications for the E ring	21
3.1.6 Imaging approach and particle-size sensitivity	22
3.1.7 Vertical structure and equatorial depletion	23
3.1.8 Global north-south warp of the mid-plane	23
3.1.9 Longitudinal (hour-angle) variations	23

3.2	A system sculpted by non-gravitational forces	24
3.3	New findings	26
3.3.1	The E ring grating effect	26
3.3.2	Enceladus–Saturn electromagnetic coupling	27
4	Cassini Instrument Deep Dive: ISS and VIMS	28
4.1	Scene physics: from the Sun to the E ring	28
4.1.1	Illumination and observation geometry	28
4.1.2	Forward scattering in the E ring	28
4.2	ISS imaging pathway (E17)	29
4.3	VIMS spectro-imaging pathway (E17)	30
4.4	What an image means here	30
4.5	ISS polarimetry	31
4.5.1	Polariser sets and orientations	31
4.5.2	Optical construction and performance	32
4.5.3	How ISS uses polarisers in practice	32
4.5.4	Status of on-orbit ISS polarimetry and current evidence	33
4.5.5	Outlook and potential of ISS polarimetry	33
5	Optical Properties and Scattering Framework	34
5.1	Physical principles of light-matter interaction	35
5.1.1	Complex refractive index and scattering regimes	35
5.1.2	Mueller matrix formalism and degree of linear polarisation	35
5.1.3	Periodic density structures as diffraction gratings	36
5.1.4	Stokes parameter retrieval from ISS polariser triplets	36
5.2	Optical constants of water ice	36
5.2.1	Hexagonal ice (I_h): broad-spectrum constants	36
5.2.2	Phase differentiation: amorphous versus crystalline ice	37
5.3	Application to E ring modelling	40
5.3.1	Simulation methodologies	40
5.3.2	Applicability to the present study	41
6	Thesis Outlook	42
6.1	Scientific objectives and research questions	42
6.1.1	Main research question	42

6.1.2	Sub questions	42
II	Scientific Article	55
	Automated Detection and Characterisation of Luminous Bands in Saturn's E Ring	56
III	Research Chapters	75
8	PyISS: From Archive to Image Pairs	78
8.1	Querying the archive	78
8.1.1	Query builder	80
8.2	Reconstructing observation sets	81
8.2.1	Temporal grouping	81
8.3	Building image pairs	82
8.4	Software	83
9	Data Preparation Pipeline	84
9.1	The residual image and why it reveals bands	84
9.2	Download and format conversion	85
9.3	Enceladus disc masking	87
9.4	Quality control	87
10	Detecting and Characterising Luminous Bands	90
10.1	Why frequency-domain detection?	90
10.1.1	Starting from a gray-level residual image	90
10.1.2	The 2D Fourier transform and the angular spectrum	92
10.1.3	The precision problem: scalloping loss	93
10.1.4	Zero-padding: achieving sub-degree precision	95
10.1.5	Statistical significance: the null ensemble	96
10.2	Image preparation	98
10.3	Pipeline parameters	100
10.4	Survey results	100
10.4.1	Detection geometry	101
10.4.2	Example detections and the doubtful category	101

10.4.3 Band contrast as a function of observing geometry 104

References (Research Chapters) 107

About This Document

Enceladus, the small icy moon of Saturn, continuously vents a plume of water vapour and ice grains into space, feeding Saturn’s diffuse E ring with a constant supply of fresh material. The E ring is a broad, thin disc of micron-sized water-ice particles that would otherwise drain away within decades. In high-phase-angle imaging, where forward-scattered light from fine grains dominates the scene, the E ring occasionally reveals faint quasi-periodic brightness features known as luminous bands: inclined stripe-like structures whose orientation and contrast carry information about the ring’s physical state and the electromagnetic environment that shapes its particle population. Rubbrecht et al. (2025) reported these features in Cassini ISS and VIMS observations and noted a chromatic character: the apparent position of a band appears to shift with wavelength, which, if confirmed, would distinguish luminous bands from the spectrally neutral ring background and point toward a wavelength-dependent physical mechanism.




Figure 1: Luminous bands in Saturn’s E ring, revealed by image processing. Enceladus and its south-polar plume are visible at the lower edge.

This thesis presents an automated, mission-wide survey for luminous bands in the Cassini ISS archive, the first to operate without pre-selection by flyby identity or

prior knowledge of where bands are expected to appear. It is organised in three parts.

Part I is the literature study. Understanding what luminous bands are, and whether the chromatic hypothesis is physically plausible, required reviewing four bodies of knowledge: the geology and plume activity of Enceladus, which set the source conditions for E ring material; the dynamics and electromagnetic environment of the ring, which determine how particles are organised once injected; the optical properties of water-ice grains, which govern how scattering varies with wavelength; and the characteristics of the ISS instrument, which determine what the data can and cannot reveal. Part I closes with the research questions that motivated the work.

Part II is the central contribution of this thesis. It presents the scientific results as a self-contained journal article that can be read independently of the surrounding chapters. The article reports the detection catalogue, characterises band orientations and contrasts as a function of observing geometry and wavelength, and connects the survey findings to the chromatic hypothesis and to the broader electromagnetic environment of the Saturn–Enceladus system. Figure 1 is a product of the image processing pipeline described in Part III. Readers interested primarily in the scientific findings are encouraged to read Parts I and II and treat Part III as optional reference material.

Part III documents the technical infrastructure that made the survey possible. Chronologically, it was built before any result could exist: without the archive query tool, the calibration pipeline, and the detection algorithm described in Chapters 7–9, there would be nothing to report in Part II. It is presented here, after the findings, so that the scientific case is made first and the methodology can be read with that context in mind. Part III serves as a complete technical record for reproducibility: anyone wishing to understand how the catalogue was constructed, verify an individual detection, or extend the survey to new data will find the necessary detail there.

Part I

Literature Study

Chapter 1

Historical context and discovery of Enceladus

Enceladus was discovered in 1789 by William Herschel. In his book, *On the Satellites of Saturn*, Herschel describes how "protuberant and lucid points on the arms of the rings of Saturn have been noticed in previous observations of the Saturnian System" (S. W. Herschel, 1790, p. 1), which led him to perform more focused investigations in search to explain these phenomena. Herschel successfully calculated the distances and orbital periods of the newly discovered moons, which were previously known only by number. His son, John F. W. Herschel, later proposed the mythological names still in use today, including Enceladus, in his *Results of Astronomical Observations* (J. F. W. Herschel, 1847, p. 415).

After William and John Herschel's early observations, the next significant contribution came more than a century later. The early telescopic era provided little new information on Enceladus beyond more accurately determining its orbital parameters. However, in 1914, Slipher (1914) reported a temporary increase in Enceladus' brightness based on measurements made with the 0.6-m telescope at Lowell Observatory. Although these data were limited by the observational techniques of the time, they noted for the first time that Enceladus' surface or illumination might vary over time. As summarised by Dougherty et al. (2018), this observation was largely overlooked for decades, and no follow-up studies were able to confirm or explain the reported brightening. Only with the Voyager flybys in the early 1980s did Enceladus again emerge as an object of significant scientific interest.




Figure 1.1: Voyager 2 mosaic of Enceladus, made from images taken through the clear, violet, and green filters Aug. 25, 1981. Credit: NASA/JPL-Caltech; PIA01394

Figure 1.1 shows the highest-resolution image of Enceladus obtained during the Voyager encounters. The Voyager missions provided the first detailed views of the moon's surface, revealing a strikingly young and geologically diverse terrain compared to Saturn's other icy moons (Smith et al., 1982). Even in Figure 1.1, the contrast between cratered and smooth plains is evident, as Smith et al. (1982) describes: "Enceladus' surface ranges from old, densely cratered terrain to relatively young, uncratered plains crossed by grooves and faults" (1982, para. 1). Although no active phenomena were directly observed at the time, the contrast between heavily cratered regions and smoother plains hinted at a complex resurfacing history. These findings reestablished Enceladus as a body of considerable scientific interest and laid the groundwork for future exploration.

Prior to Cassini's arrival, the circumstantial case for Enceladus as an active body had already been building through indirect observations of the E ring. Showalter et al. (1991) documented that the ring's particles are concentrated in an unusually narrow size range around 1 μm and that the ring's brightness peaks precisely at Enceladus' orbit, a distribution incompatible with passive collisional processes and indicative of active resupply. Hamilton and Burns (1994) demonstrated dynamically that micron-sized grains launched from Enceladus would naturally populate the elliptical orbits reproducing the ring's observed structure, providing a quantitative link between the moon and its surroundings.

The arrival of the Cassini spacecraft in 2004 transformed Enceladus from a geologically intriguing moon into one of the most active worlds in the Solar System. During three close flybys in 2005, Cassini's instruments delivered a coherent picture: the Imaging Science Subsystem (ISS) directly imaged plumes of water ice towering above the south polar terrain and identified four fractures, the tiger stripes, as their source (Porco et al., 2006); the Composite Infrared Spectrometer (CIRS) detected a 3–7 GW endogenic thermal anomaly along those fractures (Spencer et al., 2006); the Ion and Neutral Mass Spectrometer (INMS) characterised the plume gas as predominantly water vapour with trace organics (J. Waite et al., 2006); and the Cosmic Dust Analyser (CDA) confirmed in situ that the plume supplies the E ring (Spahn et al., 2006). These results were published as a coordinated package in *Science* in 2006, establishing Enceladus as only the third solar system body, after Earth and Io, with confirmed internally driven geological activity.

Subsequent Cassini observations progressively built the case for a habitable interior. Gravity measurements confirmed a subsurface liquid water reservoir beneath the south polar ice shell (Iess et al., 2014), later shown to be a global ocean through detection of a forced physical libration of the ice shell (Thomas et al., 2016). Silica nanoparticles found in Saturn's inner system, originating from plume grains, provided the first evidence of ongoing high-temperature water-rock reactions on Enceladus' ocean floor (Hsu et al., 2015). Molecular hydrogen detected in the plume during the final targeted flyby in 2015 then confirmed active serpentinisation, creating chemical conditions that could sustain methanogenesis (J. H. Waite et al., 2017).

In retrospect, the photographic evidence for the plumes had existed far earlier than recognised. A reanalysis of the Voyager 1 Wide Angle Camera images from 13 November 1980 by Stryk (2017), made possible only after knowing what to look for, identified faint plume structures extending from Enceladus' south pole, as shown in Figures 1.2 and 1.3.



Figure 1.2: Reprocessed Voyager 1 Wide Angle Camera image from 13 November 1980, showing Saturn with Enceladus visible above the rings. Credit: NASA/JPL-Caltech/Ted Stryk (Stryk, 2017).



Figure 1.3: Close-up from the same dataset, showing Enceladus above Saturn's limb. The plume is visible at the bottom of the crescent, illuminated by sunlight and Saturn-shine. Credit: NASA/JPL-Caltech/Ted Stryk (Stryk, 2017).

Chapter 2

Formation and evolution models of the Saturnian moons

The origin of Saturn’s diverse satellite system, from the giant, atmosphere-bearing Titan to the icy, mid-sized moons and the enigmatic rings, remains a fundamental and complex question in planetary science. Any successful formation theory must account for the system’s current orbital architecture, the varied compositions of the moons, and the powerful tidal interactions that govern their evolution. This challenge is heightened by Cassini-era discoveries, such as the unexpectedly rapid outward migration of the moons (Lainey et al., 2020), which places new and stringent constraints on their ages and formation timescales.

This chapter reviews the primary scenarios currently debated in the scientific community. The four main contexts for moon formation are outlined in Section 2.1, followed by the critical roles of orbital evolution and tidal dissipation (Section 2.2). The three leading, and competing, models for Saturn’s regular moons are then examined: late formation via disruption, in-situ formation from a circumplanetary disk, and the tidal disk model involving a massive ancient ring (Section 2.3). The origin of the irregular moons (Section 2.4) and a proposed dual-formation scenario that attempts to unify these different concepts (Section 2.5) are considered in closing.

2.1 Contextualising regular moon formation

The origin of Saturn’s mid-sized, regular moons, which orbit in prograde, near-equatorial paths, is a core area of planetary science research. As Blanc et al. (2025) discusses, any successful theory regarding the formation of these satellites must reconcile with their estimated surface ages, which suggest an antiquity of approximately four billion years, provided that most impact features originate from heliocentric projectiles.

The scientific community generally considers four primary scenarios for the creation of regular moon systems, as illustrated conceptually in Figure 2.1 (Blanc et al., 2025):

1. **Accretion Following a Catastrophic Impact:** This model posits that a massive impact or similar disruptive event generates a temporary debris disk, from which the moons subsequently accrete. This mechanism is the favoured explanation for Earth’s Moon (Canup et al., 2023) and has been adapted for systems like the prograde moons of Uranus, where the re-accretion of debris from a disrupted, primordial satellite system is believed to have followed the giant impacts that likely tilted the planet’s spin axis. However, this scenario requires a specific set of initial conditions and assumptions.
2. **In Situ Formation from a Primordial Circumplanetary Disk (CPD):** Considered the primary mechanism for the Jovian system, this generic model involves the formation of regular moons directly within a natal CPD that orbits near the host planet’s equatorial plane. While highly applicable to Jupiter’s Galilean moons, it is unclear if this scenario is universally applicable to all giant planets, particularly ice giants.
3. **Evolution and Outward Migration from a Massive Ring System (Tidal Disk):** Originally proposed to explain the Saturnian system’s highly dissipative nature (i.e. its strong tidal coupling between planet and moons), this hypothesis involves a massive body being stripped within the Roche limit, forming a large ring of debris. Moons then form from the outer edge of this ring and migrate outward. This concept has been generalised to ice giants and Earth’s moon.

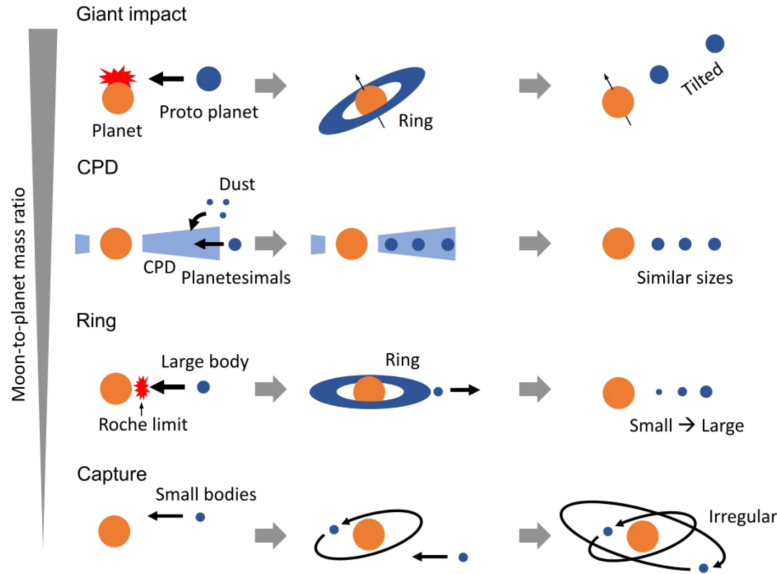


Figure 2.1: A schematic of the four main mechanisms for moon formation. From top to bottom: Giant Impact, Circumplanetary Disk (CPD), a spreading massive Ring (Tidal Disk), and Capture. These scenarios produce systems with different typical moon-to-planet mass ratios and orbital characteristics. From Blanc et al. (2025).

4. Gravitational Capture of Small Bodies: This mechanism is reserved for explaining the populations of irregular moons around all four giant planets, which exhibit highly inclined and often retrograde orbits. Capture typically involves a three-body interaction, such as collisions during the Solar Nebula phase or close encounters with a third body, as featured in the Nice model (Nesvorný et al., 2007), or the tidal disruption of a binary Trans-Neptunian Object (TNO), as proposed for Neptune’s moon Triton.

Of these scenarios, the tidal-disk model (scenario 3) is currently the leading hypothesis for Saturn’s regular moons, including Enceladus. The key constraint ruling out early in-situ formation is the unexpectedly rapid outward migration of the moons, confirmed by Lainey et al. (2020) using the full Cassini dataset. These fast migration rates imply the moons formed much closer to Saturn and at a geologically recent epoch, inconsistent with a primordial CPD origin. The physical basis for these migration rates, tidal dissipation within Saturn and the resonance-locking mechanism, is discussed in Section 2.2, and the competing late-formation scenarios are reviewed in detail in Section 2.3.

2.2 Orbital evolution, tidal dissipation, and age constraints

Saturn is known to be a highly dissipative host planet, meaning that tidal energy raised in the planet by its moons is converted to heat within the planet at an exceptionally high rate, making the migration rate of its moons a critical parameter for constraining their formation time and location. The orbital migration of a moon is often assessed using the quality factor, or Q , associated with the tidal interaction mode responsible for orbital change. Under the classical tidal theory assumption of constant Q , the migration rate decreases rapidly with increasing radial distance from the planet, implying that distant satellites experience minimal migration over the age of the Solar System.

However, the assumption of constant Q was strongly challenged by Fuller et al. (2016), who proposed an alternative mechanism involving resonance locking between the moon’s orbit and the planet’s internal oscillation modes (e.g., gravity, inertial, or fundamental modes). In this mechanism, the moon and the planet’s internal modes maintain phase synchrony, enabling efficient and sustained energy transfer. This resonant lock can produce a migration regime that is less dependent on radial distance, making it equally applicable to distant moons. The profound difference between these two migration models is visualized in Figure 2.2 (Blanc et al., 2025).

Migration timescales derived from the classical model, which assumes a constant Q (dashed curves in Figure 2.2), suggest that if their current locations are assumed, the distant moon Titan must have formed near its

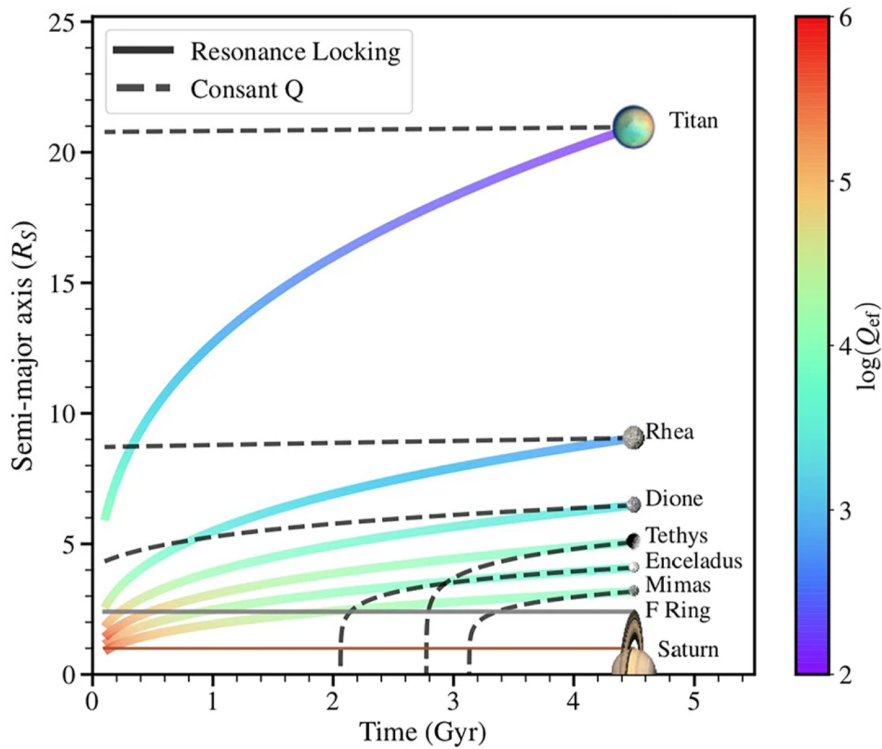


Figure 2.2: A comparison of orbital migration trajectories for Saturn’s moons over the age of the Solar System. The "Constant Q" model (dashed curves) predicts that distant moons, such as Titan, barely migrate. In contrast, the "Resonance Locking" model (solid curves) predicts rapid migration for all moons, including Titan, suggesting they formed much closer to the planet. The colour of each solid curve indicates $\log(Q_{\text{ef}})$, the effective tidal quality factor experienced by that moon under resonance locking; because each moon locks onto a different internal oscillation mode of Saturn, its effective Q differs from that of its neighbours. From Blanc et al. (2025), adapted from Fuller et al. (2016).

present orbit, whereas inner moons might have formed considerably later than the Solar System’s inception. In contrast, the resonant locking model (solid curves in Figure 2.2) predicts that all mid-sized moons migrate at comparable, relatively swift rates (Blanc et al., 2025).

Crucially, observations by Lainey et al. (2020), utilising the full Cassini dataset, tended to validate the predictions of the resonant-locking migration regime, confirming that the moons are migrating outwards at an unexpectedly fast rate. For instance, Titan’s orbit is expanding approximately ten times faster than predicted by the classical model. This confirmation of fast migration has significant implications for the plausible formation models.

2.3 Detailed scenarios for Saturn’s regular moons

The confirmation of rapid tidal migration, as discussed in Section 2.2, fundamentally reshapes the constraints on satellite formation. It implies that moons, including Titan, likely formed much closer to Saturn. This section reviews the three leading scenarios proposed to explain the origin of Saturn’s regular moons in light of these new dynamical models.

2.3.1 Late formation via disruption and re-accretion

This scenario proposes that the current moon system is the product of a collisional disruption and subsequent re-accretion of a previous, older satellite system (scenario 1 in Figure 2.1). Asphaug and Reufer (2013) first explored this idea, starting from the premise that Saturn’s primordial CPD initially gave rise to Galilean-like satellites, which were later disrupted. Their analysis, utilising 3D Smooth Particle Hydrodynamics, is illustrated in Figure 2.3. This simulation showed that both Titan and the mid-sized moons could be simultaneously reproduced as products of this catastrophic event. While a merit of the model is its comprehensive outcome, it does not inherently account for the current mass-distance relationship or the circularized orbits of the newly formed moons.

A variation, the “young moon” hypothesis by Čuk et al. (2016), focuses on the role of orbital resonances in the dynamical evolution of the mid-sized moons. The authors investigated two potential resonance crossings in the system’s past: the 3:2 mean-motion resonance (MMR) of Tethys and Dione, and a later 5:3 MMR of Dione and Rhea. They found that only the Dione–Rhea crossing yields final eccentricities consistent with the moons’ present-day values. Given the rapid tidal migration rates observed by Lainey et al. (2012), this suggests that the formation of Dione and Rhea must have occurred between the two potential resonance crossings, implying a remarkably recent origin, possibly as young as 100 Myr.

Čuk et al. (2016) therefore propose that the mid-sized moons formed late from a disk of solids created by the catastrophic disruption of an earlier moon system, potentially related to the rings themselves. The validation of this scenario hinges on several future observations: a more precise measure of tidal heating in Enceladus, a definitive post-Cassini validation of the fast tidal migration rates which would prevent resonance crossing in typical constant- Q models (Lainey et al., 2020), and a careful analysis of the cratering records to verify surface age-dating.

2.3.2 Formation from a circumplanetary disk (CPD)

Regular satellites are generally thought to have originated from a CPD, mirroring the planet formation process within a protoplanetary disk (PPD). CPDs are byproducts of the gas accretion phase of planets, forming as the growing planet carves a gap in the PPD. Moons are thought to form relatively close to the centre of the CPD, within a region of approximately 0.03 to 0.1 Hill radii (R_H), where the Hill radius is approximated by:

$$R_H \approx a \left(\frac{m_2}{3(m_1 + m_2)} \right)^{\frac{1}{3}} \quad (2.1)$$

Here, a is the semi-major axis, m_1 is the planet mass, and m_2 is the stellar mass.

Two main models characterise CPD formation:

1. The Minimum Mass Model (Segregated Disk): This model, proposed by Mosqueira and Estrada (2003a, 2003b), assumes that the CPD is segregated from the PPD and that the total solid mass available for

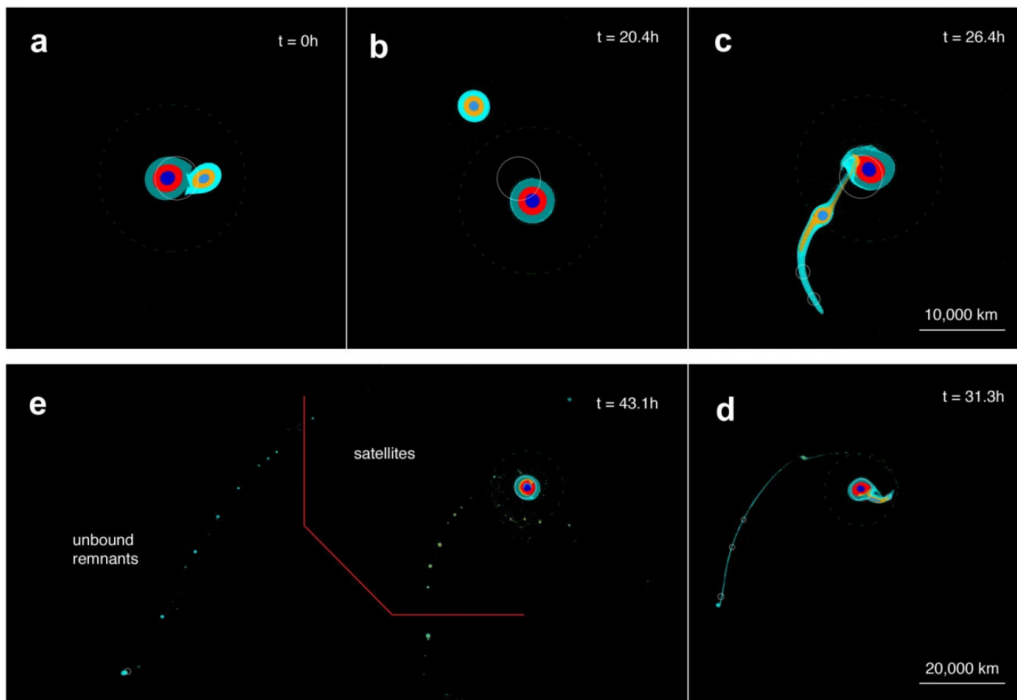


Figure 2.3: Five consecutive SPH simulation snapshots (clockwise, panels a–e) by Asphaug and Reufer (2013) of a 1:3 mass-ratio collision between a Titan-sized satellite and a smaller impactor. Panel (a) shows the initial approach ($t = 0$ h); panel (b) the first close encounter ($t = 20.4$ h); panel (c) the onset of tidal disruption and stretching of material ($t = 26.4$ h); panel (d) the formation of clumps along the debris stream ($t = 31.3$ h); and panel (e) the final state ($t = 43.1$ h), in which the red line separates the clumps that remain in Saturn-centred orbit and may form mid-sized moons (right) from the unbound remnants that escape the system (left). The merged central body is the analogue of Titan. From Blanc et al. (2025).

moon formation is approximately equal to the current moon mass (e.g., $M_s \approx 10^{-4}M_p$, where M_s is the total satellite mass and M_p is the planet mass). The mass discrepancy between Titan and Iapetus is explained by the distribution of solids, which are concentrated closer to Titan’s present orbital location.

2. The Gas-Starved Disk Model: This model suggests that gas accretion onto the CPD persists during moon formation, but at a rate low enough to keep temperatures below the ice sublimation point (Canup & Ward, 2002). In this scenario, solid material is continuously delivered from the parent PPD. Recent numerical simulations lend support to this by showing that moderate gas accretion continues even after the embedded planet has cleared an orbital gap (Haffert et al., 2019; Szulágyi et al., 2022).

A critical challenge for these models is the efficient formation of satellitesimals, the initial building blocks of moons. Dust grains supplied to the CPD must grow in size through collision and adhesion but tend to drift inward towards the planet as centimetre- to metre-sized pebbles before reaching the size of satellitesimals (Shibaike et al., 2017; Zhu et al., 2018). Potential solutions involve:

- Outward flows in specific CPD regions facilitating satellitesimal formation via streaming instability (Batygin & Morbidelli, 2020; Cilibrasi et al., 2018; Drażkowska & Szulágyi, 2018).
- Planetesimals formed in the PPD being gravitationally captured by the CPD and acting as satellitesimals (Lyra et al., 2009; Shibaike & Alibert, 2020).

Despite these formation challenges, the satellitesimal accretion model of Sasaki et al. (2010) successfully reproduces the observed mass-distance relationship and rock-to-ice ratios for both Saturn’s and Jupiter’s regular moon systems, as shown in the top panel of Figure 2.4.

An alternative to satellitesimal accretion is pebble accretion, a mechanism, initially developed for planet formation in PPDs (Lambrechts & Johansen, 2012; Ormel & Klahr, 2010). Pebbles are objects large enough to partially decouple from the gas, but small enough to still be significantly affected by gas friction. Analogously, in CPDs, moon embryos accrete these drifting pebbles as they move towards the planet. Figure 2.4 shows the two models side by side, though on different axes: the top panel shows the final moon mass as a function of orbital distance predicted by the satellitesimal model (Sasaki et al., 2010), while the bottom panel shows the CPD thermodynamic structure (temperature and surface density) from the pebble accretion model (Ronnet & Johansen, 2020a), with the observed positions of Titan and the Galilean moons overlaid. However, two issues complicate pebble accretion for Saturn’s moons:

1. The ablated material from which Saturnian moons accrete may lack ice, contradicting their observed low bulk densities. This is mitigated if hydrated minerals are abundant, as the subsequent dehydration and outward diffusion of water vapour could facilitate the recondensation of substantial amounts of ice (Ronnet & Johansen, 2020b).
2. The necessity of inner cavities to impede rapid inward migration (Ronnet & Johansen, 2020a) was likely not met in young Saturn, posing a challenge to Titan’s formation through this mechanism.

2.3.3 Formation/migration from a massive ring (Tidal Disk Model)

This model suggests that moons form from the spreading of a dense, massive tidal disk located within the planet’s Roche limit. A tidal disk is dynamically cold, but its particles are prevented from coalescing by the planet’s strong tidal forces. As the disk spreads beyond the Roche limit, the tidal forces weaken, allowing icy particles to gravitationally cohere.

The mechanism of outward migration is crucial because of Keplerian shear: inner ring particles orbit faster and outer ones slower than the newly formed moonlet. These differential velocities create asymmetric gravitational torques that transfer angular momentum, resulting in a net outward migration of the moonlet. This physical interaction, where a moonlet gravitationally sculpts the ring edge and gains angular momentum, is directly observed in Saturn’s rings, as shown in Figure 2.5.

The migration speed and the disk spreading rate determine the outcome: either one large moon (e.g., Earth’s Moon) or numerous smaller ones (e.g., Saturn’s moons). If the migration rate decreases with distance, the resulting collisions and mergers result in a mass increase with increasing orbital distance. This process has been modelled around Saturn’s present rings and successfully reproduces the small inner moons (Janus, Epimetheus, Prometheus, Pandora, and Atlas) (Charnoz et al., 2010).

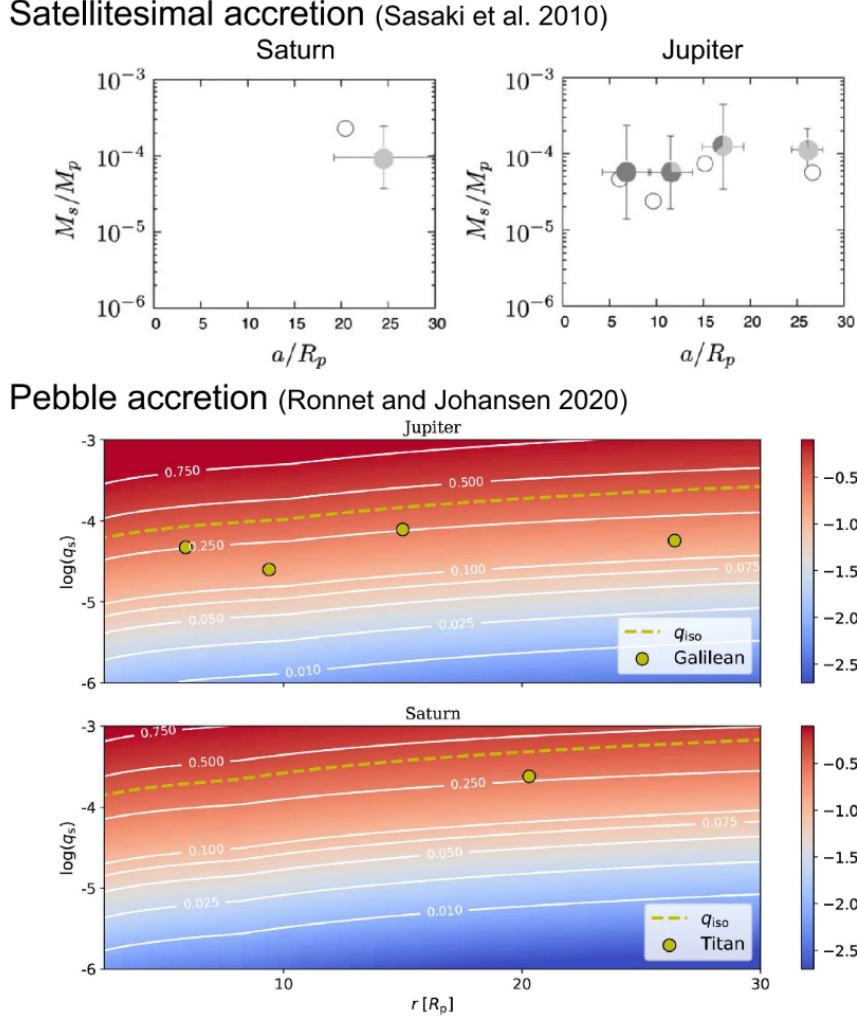


Figure 2.4: A comparison of moon formation models. **Top:** The satellitesimal accretion model (Sasaki et al., 2010). The y-axis shows the satellite-to-planet mass ratio M_s/M_p as a function of normalised orbital distance; the grey shading of each data point encodes the rock-to-ice ratio, with darker points being more rocky and lighter points more ice-rich. The model reproduces the observed mass-distance relationship and compositions for both Saturn and Jupiter. **Bottom:** The pebble accretion model (Ronnet & Johansen, 2020a), showing the CPD thermodynamic structure (temperature and surface density) as a function of orbital distance. The current masses of the Galilean moons and Titan (yellow circles) are below their pebble isolation mass (dashed lines; the mass at which a growing body generates a pressure bump that halts further pebble accretion), making this a viable formation pathway. From Blanc et al. (2025).




Figure 2.5: The moonlet Daphnis orbiting within the Keeler Gap of Saturn’s A ring. Its gravity creates waves on the gap edges, a visible manifestation of the angular momentum exchange that drives the outward migration of moons formed from a tidal disk. Credit: NASA/JPL/Space Science Institute; PIA07809

The model has been generalised to explain all of Saturn’s mid-sized moons. The rings may have been primordial and substantially more massive in the past, losing mass to the mid-sized moons that formed beyond the Roche radius. The strong internal tidal dissipation within Saturn (Lainey et al., 2020) is essential, allowing the mid-sized moons to migrate to their current orbits from the Roche radius.

The analytical description by Crida and Charnoz (2012) demonstrates the mass–distance relationship for moons forming from a spreading tidal disk that does not yield a single large satellite. The mass q , expressed in units of the central planet’s mass M_p , relates to the normalized distance Δ from the Roche radius ($\Delta = (R - R_{\text{Roche}})/R_{\text{Roche}}$) by:

$$q \propto \begin{cases} \left(\frac{\Delta}{\Delta_{2:1}}\right)^{9/5} & \text{for } \Delta < \Delta_{2:1} \\ \left(\frac{\Delta+1}{\Delta_{2:1}+1}\right)^{3.9} & \text{for } \Delta > \Delta_{2:1} \end{cases} \quad (2.2)$$

where $\Delta_{2:1} \approx 0.58$ corresponds to the location of the 2:1 mean motion resonance with the Roche radius. This relation provides a good fit for moons ranging from Pandora out to Rhea, and even Titan.

This model is inconsistent with the present-day system if tidal dissipation (Q) is constant, but the fast, resonant-locking migration proposed by Lainey et al. (2020) makes the concept viable. Furthermore, the model suggests that this formation mechanism may also apply to the regular moons of Uranus and Neptune, implying that these ice giants also possessed massive rings in the past (Blanc et al., 2025).

2.4 Origin of irregular moons

The irregular moons of Saturn are characterised by their large inclinations and often retrograde orbits, clearly differentiating them from the regular moons. These bodies are widely believed to be captured small, Sun-orbiting bodies through various mechanisms:

- Pull-down capture by Saturn as the body crossed the planet’s Hill sphere, facilitated by gas drag during the Solar Nebula phase (Pollack et al., 1979).
- Collisions between Sun-orbiting planetesimals within the Solar Nebula, producing fragments that could be captured by Saturn’s gravity (Colombo & Franklin, 1971).
- Three-body interactions occurring during the dynamical instability of the outer planets, as described by the Nice model (Nesvorný et al., 2007).

A more detailed analysis of the orbital and physical characteristics of these moons is required to differentiate between these formation scenarios (Blanc et al., 2025).

2.5 A tentative dual formation scenario for Saturn’s system

Blanc et al. (2025) propose a dual formation scenario for the Saturnian and Jovian systems, based on an assessment of the critical radii of young Saturn, including the planet radius (R_p), the Roche limit (R_{Roche}), and the size of the primordial magnetospheric cavity (R_{cav}).

The Roche radius for a moon with density ρ_m orbiting a planet with density ρ_p and radius R_p is defined as:

$$R_{\text{Roche}} = 2.456 \left(\frac{\rho_p}{\rho_m} \right)^{\frac{1}{3}} R_p \quad (2.3)$$

Furthermore, the radius of the inner cavity carved by the planet’s magnetic field is defined by the magnetospheric accretion model:

$$R_{\text{cav}} = \left(\frac{\mathcal{M}^4}{4GM_p \dot{M}_g^2} \right)^{\frac{1}{7}} \quad (2.4)$$

Where \mathcal{M} is the planet’s magnetic moment, G is the gravitational constant, M_p is the planet’s mass, and \dot{M}_g is the mass accretion rate of the gas onto the planet (Takasao et al., 2022). Analysis suggests that during early times, young Saturn likely lacked a wide inner cavity, meaning the CPD was directly connected to the planet’s surface.

This dual scenario proposes that the CPD gave rise to at least two large satellites (here temporarily denoted Moon 1 and Moon 2). These moons likely had masses similar to those of the current Titan, constrained by Type I migration timescales or the pebble isolation mass.

For clarity, the inner moon (Moon 1) is referred to as Gaea and the outer moon (Moon 2) as proto-Titan for the remainder of this discussion, the latter name immediately conveying its eventual fate.*

1. Origin of the Rings and Mid-Sized Moons (Gaea): The inner satellite, Gaea (Moon 1), migrating inward due to Type I migration in the CPD, eventually crossed the Roche limit. The moon’s ice-rich mantle would have been tidally stripped to form a massive ice ring, while its rocky core would have fallen into the planet. The subsequent viscous expansion of this massive ring beyond the Roche limit would have given rise to the mid-sized moons (as described in the tidal disk model). This scenario aligns with Canup (2010) proposal that the rings originated from the tidal stripping of an inward-migrating, Titan-mass satellite.
2. Origin of Titan (proto-Titan): The outer satellite, proto-Titan (Moon 2), is the direct precursor of present-day Titan. Unlike Gaea, it survived without disruption and must have either:
 - Formed late, just before the CPD dissipated, halting its Type I inward migration before falling inside the Roche limit.
 - Been captured and subsequently pushed outward by the powerful resonant-locking tides from Saturn. The rapid migration, as confirmed by Lainey et al. (2020), makes this outward capture and migration plausible.

This dual framework also allows for a third object, such as Iapetus, to have been born further out in the CPD, where its smaller mass would have protected it from significant Type I migration. A key remaining constraint for this scenario is reconciling the rings’ age with the need for them to contain sufficient silicates from Gaea’s mantle to provide the necessary material to the mid-sized moons. The proposed sequence of events for the Saturnian system is summarised in Figure 2.6.

*These names are not used in Blanc et al. (2025); they are proposed here by the author for clarity of exposition.




Figure 2.6: A schematic of the tentative dual formation scenario for Saturn. Top to bottom: (1) Two large moons (1 and 2) form in the CPD. (2) Moon 1 migrates inward via Type I migration and is tidally disrupted at the Roche limit. (3) This disruption forms a massive ring. (4) The ring spawns the mid-sized moons, while Moon 2 (Titan) survives, and all moons migrate outward via resonance locking. From Blanc et al. (2025).

Chapter 3

Plumes and the E ring

Saturn’s E-ring is a faint, dust-rich structure distinct from the main rings, spanning from the neighbourhood of Mimas to at least Titan’s orbit and composed predominantly of sub-micron to micron-sized water-ice grains (Hedman et al., 2012). Voyager first revealed a vertically integrated brightness peak close to Enceladus (Showalter et al., 1991), while Earth-based observations refined the picture by showing that the ring thickens with increasing radius exterior to Enceladus (de Pater et al., 2004; Nicholson et al., 1996; Showalter et al., 1991). Cassini established Enceladus as the E ring’s source, with cryovolcanic venting near the south pole, injecting gas and ice into orbit (Spencer et al., 2009). Beyond identifying the source, Cassini’s imaging enables a three-dimensional description of the ring and a dynamical link to the plume physics and Saturn’s magnetospheric environment (Hedman et al., 2012).

The cryovolcanic jets of Enceladus, first resolved by Cassini in 2005 (Porco et al., 2006), originate from a set of four sub-parallel fractures near the south pole known as the tiger stripes. These fissures eject water vapour, micron-sized ice particles, and organic compounds into space; more than 30 individual jets have been identified in a single Cassini image (Porco et al., 2006), as shown in Figure 3.1. The ejected material partially escapes Enceladus’s gravity and is the primary source of the E ring, as established above. The plume composition, which also includes molecular hydrogen and silica nanoparticles (Postberg et al., 2018; J. H. Waite et al., 2017), points to ongoing hydrothermal activity at the ocean floor, making the plumes a direct window into the interior.

3.1 Time variability of the Enceladus plumes

Enceladus’ plumes provide a direct sample of the subsurface ocean, so their variability probes how tidal and surface stresses modulate exchange with the interior. Using Cassini ISS observations from 2005–2017, Ingersoll and Ewald (2020) identify three characteristic timescales and related behaviours. This orbital-period variability is clearly captured in Figure 3.2, which plots the measured plume slab density against Enceladus’s orbital position (mean anomaly).

3.1.1 Diurnal forcing (1.37 days)

Plume brightness varies with Enceladus’ orbital period, peaking near apoapsis where tensile tidal stresses tend to open the tiger-stripe fractures. The phase curve is stable across the mission. A distinct secondary maximum appears ahead of the main peak at mean anomaly $M \approx 25^\circ\text{--}60^\circ$, with two peaks near 36° and 51° . This two-peak structure is particularly well-defined by the 2017 data (the black points in Figure 3.2), and the solid trend line, which is a model fit to all 2017 observations, clearly highlights their locations.

3.1.2 Multi-year modulation (~ 4 and ~ 11 years)

After removing the diurnal cycle with a template, a longer-term modulation in overall plume output becomes apparent. These patterns are shown in Figure 3.3, where dividing the data by the template (the trend line from Figure 3.2) removes the orbital variation and clearly separates the data by year. Maxima in $\sim 2007\text{--}2009$ (cyan points) and $2016\text{--}2017$ (black points) are visibly higher than the trough during $2010\text{--}2015$ (orange and red points).

These align with the 3.9- and 11.1-year components of the Enceladus–Dione resonance that modulates the orbital eccentricity (e) by $\sim 5\%$. This change in eccentricity alters the degree to which Enceladus’s orbit is




Figure 3.1: Cassini mosaic of Enceladus's south polar plumes, composed of two high-resolution images acquired on 21 November 2009 at a distance of approximately 14 000 km. More than 30 individual jets of varying sizes are visible, erupting from the tiger-stripe fractures. Credit: NASA/JPL/Space Science Institute.




Figure 3.2: Slab density vs. mean anomaly with a trend line superimposed on the 2017 data (Ingersoll & Ewald, 2020)




Figure 3.3: Slab density normalised by the diurnal template. This view highlights the multi-year modulation, with 2009 (cyan) data being high and 2015 (red) data being low. (Ingersoll & Ewald, 2020)

elliptical. A more squashed (more eccentric) orbit increases the magnitude of the tidal stresses Enceladus experiences from Saturn, which in turn are thought to force the plume-sourcing fissures open wider. A near-linear relation between the decadal brightness factor and $\Delta r = r_{\text{apo}} - r_{\text{peri}} = 2ae$ implies a factor-of-two change in particulate mass flux over the cycle (Ingersoll & Ewald, 2020; Ingersoll & Ewald, 2017; Porco et al., 2018). This correlation strongly suggests that the $\sim 5\%$ orbital wobble, driven by the Dione resonance, is the underlying mechanism responsible for the factor-of-two breathing of the plumes over a decadal timescale.

3.1.3 Week-to-month variability (stochastic jets)

Individual vents switch on and off on \sim month timescales, largely independent of tidal phase, consistent with local processes such as transient blockages or evolving conduit geometry. This stochastic behaviour is clearly visible in near-repeat Cassini ISS images, as shown in Figure 3.4, where individual jets seen in one observation (panel a) are absent or different in another (panel b).

Near-repeat ISS image pairs show specific jets vanishing or appearing between revisits, while the global plume output remains comparatively steady due to averaging over many sources (Ingersoll & Ewald, 2020; Yeoh et al., 2017).

3.1.4 Dust-gas co-variation and velocities

Solid grains and water vapour vary together on diurnal and decadal timescales once data sets are put on a common scale, indicating a nearly constant dust-to-vapour ratio (cf. Gao et al., 2016; Hansen et al., 2011; Ingersoll & Ewald, 2020). The fall-off of slab density with altitude follows $S \propto z^{-\alpha}$ between 50 and 200 km. This exponent α is a simple measure of how quickly the plume thins out with height; a larger α indicates that the density drops off faster, suggesting that the particles had lower launch speeds and could not travel as high. Typical particle speeds at the vent are $\sim 140\text{--}150 \text{ ms}^{-1}$. Near apoapsis, when brightness and mass flux are highest, α is larger and the inferred speeds are lower. This provides a key physical insight: when the plume is densest (at apoapsis), the denser gas exerts more drag on the solid ice grains, slowing them down.

3.1.5 Implications for the E ring

These source-side variations set time-dependent boundary conditions for the E ring: mass-flux and launch-speed distributions feed the size-frequency and orbital-element distributions of freshly ejected grains. Combined with

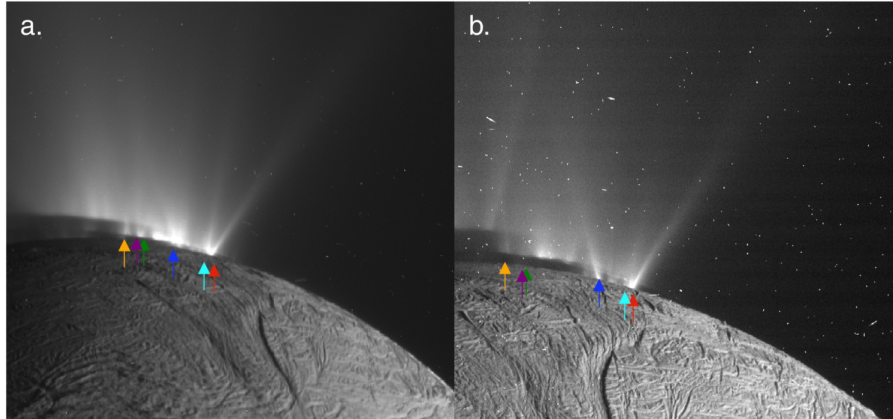


Figure 3.4: Example of stochastic jet variability from Cassini ISS. Panel (a) shows a set of active jets, while panel (b), taken at a different time, shows a different configuration (e.g., the jet indicated by the green arrow in 'a' is missing). This variability is independent of the main orbital (diurnal) cycle. Adapted from Ingersoll and Ewald (2020).




Figure 3.5: A global, near-edge-on mosaic of Saturn and the E ring from Cassini ISS. The labels (Morning, Evening, Sub-Solar, Anti-Solar) indicate the geometry relative to the Sun (Hedman et al., 2012).

non-gravitational perturbations (solar radiation pressure and Lorentz forces), they help explain the features in Section 3.2, including the equatorial depletion near Enceladus' orbit, the smooth south-to-north warp of the mid-plane, and day-night asymmetries organised by solar longitude.

Saturn's E ring is a highly tenuous, dust-dominated structure extending from the interior to Mimas to at least the orbit of Titan. The grains are mainly water-ice with radii below a few microns and are continually replenished by Enceladus' plumes (Hedman et al., 2012; Showalter et al., 1991; Spencer et al., 2009). Cassini ISS observations in 2005–2006 (solar elevation $\sim 16\text{--}20^\circ$) provide a global, near-edge-on view that constrains the ring's three-dimensional geometry in cylindrical coordinates (ρ, θ, z) (Hedman et al., 2012). A global mosaic of these observations is shown in Figure 3.5, providing a contextual map for the ring's overall structure and the geometries discussed in the following subsections. At visible wavelengths near $0.63, \mu\text{m}$, the scattered-light signal is dominated by grains with $s \sim 0.5\text{--}1, \mu\text{m}$ (transition between geometric and Rayleigh regimes), so the morphology retrieved here primarily reflects that size range (Hedman et al., 2012).

3.1.6 Imaging approach and particle-size sensitivity

Calibrated ISS data are used (CISSCAL; Porco et al., 2004) with appropriate SPICE kernels, selecting high-phase, low-opening-angle views for vertical profiles and shadow-edge mosaics for longitudinal trends. Because the scattering cross-section scales approximately as $\sigma \propto s^2$ in the geometric regime and $\sigma \propto s^6$ in the Rayleigh limit, images at $\lambda \approx 0.63, \mu\text{m}$ emphasise $s \sim \lambda/2\pi$ -scale grains (Hedman et al., 2012). Nanometre-scale particles are known in the plumes but contribute little to the ISS brightness at these wavelengths.




Figure 3.6: The vertical brightness profile of the E ring (black data points) near Enceladus’s orbit, showing the characteristic central deficit or dip at $z = 0$. This double-banded structure is accurately modelled by a composite Lorentzian Peak + Gaussian Dip (red line) (Hedman et al., 2012).

3.1.7 Vertical structure and equatorial depletion

Near-edge-on profiles show that the E-ring’s vertical full width at half maximum (FWHM) is smallest at Enceladus’ orbit ($\rho \simeq 2.40 \times 10^5$, km), where $\text{FWHM} \approx 4,000\text{--}5,000$, km, and increases roughly linearly with $|\rho - a_E|$ to $\sim 10,000$, km by the orbits of Mimas and Tethys (Hedman et al., 2012). A single Gaussian or Lorentzian does not reasonably describe the vertical profile. Between $\rho \simeq (2.30\text{--}2.80) \times 10^5$, km a double-banded morphology appears: a central deficit near $z = 0$ superposed on a broader core. This complex vertical structure is clearly shown in Figure 3.6, where the data (black line) cannot be fit by a single function but is well-described by a composite model (red line) that includes a Gaussian dip at the centre.

A composite model (Lorentzian core plus Gaussian dip) reproduces the data markedly better than either function alone (Hedman et al., 2012). The depletion is most substantial around Enceladus and has $\text{FWHM} \sim 2,000$, km, comparable to Enceladus’ Hill sphere, consistent with simulations of finite plume ejection speeds and close-encounter scattering (Juhász & Horányi, 2007; Kempf et al., 2010).

3.1.8 Global north–south warp of the mid–plane

The peak brightness is systematically displaced from the equatorial plane, southward ($z < 0$) inside Enceladus’ orbit and northward ($z > 0$) outside. Typical offsets are $\sim -1,000$ to $-2,000$, km near Mimas, crossing $z = 0$ at $\rho \approx a_E$ and reaching $\sim +1,000$, km near Tethys (Hedman et al., 2012). This smooth, large-scale warp implies that the brightest grains occupy eccentric, inclined orbits with arguments of pericentre $\omega \approx \pm 90^\circ$. Figure 3.7 provides a conceptual model of how such an eccentric and inclined orbit (dashed ellipse) translates into the observed vertical warp (solid grey area) relative to the equatorial plane (thick dashed line).

Hence, particles are south of the plane near periapsis and north near apoapsis. The longitudinal independence of the warp indicates a broad distribution of ϖ and Ω .

3.1.9 Longitudinal (hour–angle) variations

Large-scale asymmetries are organised with respect to the Sun rather than Enceladus. The global mosaic in Figure 3.5 illustrates these different viewing geometries, labelling the Sub-Solar, Anti-Solar, Morning, and Evening ansae. Comparing sub-solar and anti-solar ansae under matched geometries, the inner E-ring ($\rho \sim 1.8\text{--}2.5 \times 10^5$ km) is up to $\sim 50\%$ brighter on the day side. Furthermore, the ring’s morphology changes: the brightness peaks directly at Enceladus’ orbit on the sub-solar side but is displaced ~ 1000 km further out on the anti-solar side. The ring also appears to grow in vertical thickness more rapidly on the sub-solar side (Hedman et al., 2012). These specific differences are quantified in Figure 3.8, which plots five properties versus cylindrical radial distance ρ (in units of 10^3 km from Saturn’s rotation axis) for both the sub-solar (black lines) and anti-solar (green lines) ansae, derived from Cassini ISS clear-filter images. Note that the global mosaic in




Figure 3.7: Conceptual diagram of the E ring’s vertical structure (Hedman et al., 2012). The horizontal dashed line marks Saturn’s equatorial plane. The grey shaded bands show the ring’s vertical extent at different longitudes, illustrating both the symmetric north-south warp and the outward flare. The dashed ellipses are representative particle orbits with arguments of pericenter $\omega \approx -90^\circ$: such orbits cross south of the equatorial plane near periapsis and north of it near apoapsis, collectively producing the observed warp. Planet and rings are not shown to scale.

Figure 3.5 may give the visual impression that the anti-solar side is brighter; this reflects the specific phase angle of that single image rather than an intrinsic difference. The calibrated, matched-geometry comparison in Figure 3.8 shows that the inner ring is systematically brighter on the sub-solar side. In contrast, regions exterior to $\sim 2.5 \times 10^5$ km are nearly symmetric (Hedman et al., 2012; Ingersoll & Ewald, 2017).

High-phase images also show a sharper inner edge and a bluer colour on the morning ansa, consistent with a fundamental redistribution of grain sizes and densities. Shadow-edge mosaics further reveal that brightness contrasts change across the tip of Saturn’s shadow in the equatorial plane. This effect is dramatic, with some observations suggesting the ring is an order of magnitude fainter on the shadow’s morning side (Showalter et al., 1991), ultimately implicating shadow-induced charging as a driver of the asymmetry (Hedman et al., 2012).

3.2 A system sculpted by non-gravitational forces

The observed warp and hour-angle asymmetries agree quantitatively with the analytical framework of Hamilton (1993). For small, charged dust, non-gravitational forces, namely solar radiation pressure (the push from sunlight) and Lorentz forces (the push from Saturn’s magnetic field on a charged particle) couple eccentricity and inclination, driving $\omega \rightarrow \pm 90^\circ$. This is the underlying physical mechanism that creates the stable, warped orbital geometry, as conceptually shown in Figure 3.7. From the measured vertical offsets $z_{\text{off}}(\rho)$ one may define an effective inclination,

$$\tan i_{\text{eff}}(\rho) = \frac{|z_{\text{off}}(\rho)|}{\rho}, \quad (3.1)$$

which, together with the radial location of brightness peaks, yields effective orbital elements for the dominant grain population: $a_{\text{eff}} \approx a_E$, $e_{\text{eff}} \sim 0.02\text{--}0.04$, and $i_{\text{eff}} \sim 0.2^\circ$ increasing both inward and outward (Hedman et al., 2012). The data exhibit a near-linear coupling $\sin i_{\text{eff}} \simeq \alpha e_{\text{eff}}$ with $\alpha \approx 1.2 \times 10^{-2}$, as expected for Hamilton’s forced-equilibrium solutions. Matching the trends favours $s \sim 0.5\text{--}1, \mu\text{m}$, grain potentials of order -2 , V and $\rho_g \sim 1, \text{g cm}^{-3}$ (Hedman et al., 2012).

Hour-angle asymmetries further require time-dependent forcing from Saturn’s shadow. In sunlit regions, solar UV photons eject electrons from grain surfaces via the photoelectric effect, leaving grains with a net positive charge. Inside Saturn’s shadow this photoelectric current is cut off and grains instead accumulate ambient plasma electrons, acquiring a net negative charge. This change in charge sign and magnitude alters the Lorentz force ($\mathbf{F} = q\mathbf{v} \times \mathbf{B}$) acting on each grain, which in the Hamilton framework couples differently to inclination (i) and eccentricity (e), shifting material between the vertical and radial orbital reservoirs. Using the observed pattern and Hamilton’s expressions, Hedman et al. (2012) infer an apsidal precession rate $\dot{\omega}_x \approx 0.42^\circ \text{ day}^{-1}$ (period ~ 4 year) for the affected grains, which explains why day-side material may not yet have crossed the shadow whereas night-side grains have. Together, gravity, sunlight and electromagnetic torques produce a dynamic, time-variable dust population that is continuously replenished at Enceladus and redistributed throughout the ring.




Figure 3.8: E ring vertical structure at the sub-solar (black) and anti-solar (green) ansae, derived from Cassini ISS clear-filter images (E105PHASE observations) (Hedman et al., 2012). The x-axis is cylindrical radial distance ρ in units of 10^3 km; the dotted vertical line marks Enceladus's orbit. From top to bottom: (1) vertically integrated brightness; (2) normal I/F from onion-peeled maps; (3) peak brightness; (4) vertical offset of the ring mid-plane; (5) full width at half maximum (FWHM). The inner ring is systematically brighter on the sub-solar side, while the vertical offset is similar on both sides but the ring is clearly thicker on the sub-solar side.




Figure 3.9: A conceptual visualisation of the hypothesis for the observed stripe features. The phenomenon is thought to originate from a periodic structure (acting as a reflection grating) within the E ring. Rubbrecht et al. (2025) propose two possible origins for this structure: (left) periodic density waves or (middle) ordered particle aggregates.

3.3 New findings

3.3.1 The E ring grating effect

Building on the established understanding of plume-ring dynamics, recent analysis of high-phase-angle Cassini VIMS and ISS data has revealed a new, unexplained optical phenomenon. Rubbrecht et al. (2025) report the discovery of parallel, stripe-like features in the E ring that exhibit strong chromatic dispersion, with their position shifting in a way that is consistent with a diffraction grating.

This phenomenon is interpreted as a mm-scale periodic structure within the E ring, acting as a reflection grating. The material in this structure appears to be fresh, crystalline ice with larger particles and a trace of solid CO₂, suggesting it is recently sourced from the plumes. However, the origin of this periodic structure, whether it be density waves, particle aggregates, or another mechanism, and its unusual 43° inclination to the ring plane remain open questions. This hypothesis is visualised in Figure 3.9. This newly identified phenomenon provides a direct starting point for the research in this thesis.

Caveats

Remote-sensing brightness probes only the size range efficient at the observed wavelengths and geometries, so absolute densities and nanograin populations are less constrained. Three independent measurement approaches yield inconsistent vertical thickness estimates: the ISS remote-sensing values discussed throughout this section (Hedman et al., 2012), in situ dust impact rates from the Cosmic Dust Analyser (e.g. Kempf et al., 2008), and earlier ground-based optical observations (de Pater et al., 2004; Nicholson et al., 1996). The disagreement is not unexpected: each technique is sensitive to a different grain size range, and ISS visible-wavelength imaging emphasises micron-scale forward-scattering particles that are not the same population sampled by CDA impact detection or by ground-based data at different wavelengths. Temporal variability in the ring's structure may further contribute to the spread.

Summary of key E ring phenomena

(i) the E ring mid-plane is smoothly warped from south to north across Enceladus' orbit; (ii) a central equatorial depletion near a_E produces a double-banded vertical profile; (iii) strong day-night (hour-angle) asymmetries, including colour and inner-edge sharpness, point to shadow-modulated electromagnetic forcing; (iv) Cassini data provide direct observational confirmation of the Hamilton (1993) forced-equilibrium picture for micron-sized, charged dust in Saturn's magnetosphere; and (v) recent analysis has revealed a new, unexplained diffraction grating effect, likely caused by mm-scale periodic structures (such as density waves or aggregates) of fresh plume material.

3.3.2 Enceladus–Saturn electromagnetic coupling

Enceladus' electromagnetic interaction with Saturn's magnetosphere has long been understood to involve Alfvén wings, structures generated when Saturn's co-rotating plasma overtakes the moon at super-Alfvénic relative velocity, driving field-aligned currents that propagate toward Saturn's ionosphere (Neubauer, 1980; Saur et al., 2007). Early theoretical work predicted such coupling, and Cassini's magnetometer confirmed field-line draping and bending consistent with Alfvén wing formation during the first targeted flybys (Dougherty et al., 2006; Jia et al., 2011). Subsequent observations from the UVIS instrument detected a localised auroral spot in Saturn's northern atmosphere magnetically connected to Enceladus, providing the first direct remote-sensing evidence of a current system linking the moon to the ionosphere (Pryor et al., 2011). More recently, Hadid et al. (2026) conducted a comprehensive multi-instrument survey of all available Cassini data, identifying 36 events with clear Alfvénic signatures, including 13 newly reported detections from non-flyby orbits, and demonstrating that the interaction region extends longitudinally to at least $\sim 120^\circ$ ($\sim 2,000$ moon radii) downstream of Enceladus. Their analysis reveals that the main Alfvén wing (MAW) undergoes repeated reflections at both Saturn's ionosphere and the boundary of Enceladus' plasma torus, producing a system of reflected Alfvén wings (RAWs) whose total power decays exponentially with an e-folding distance of $\sim 7.4^\circ$ (i.e., the power drops to $1/e \approx 37\%$ of its value every 7.4° of longitude downstream), consistent with a reflection coefficient of 0.82–0.87. A systematic filamentation of the large-scale Alfvénic perturbations is observed at all distances from the moon, and a north–south asymmetry in small-scale wave activity is attributed to additional mass loading and conductivity enhancement by the southward-directed water plumes.

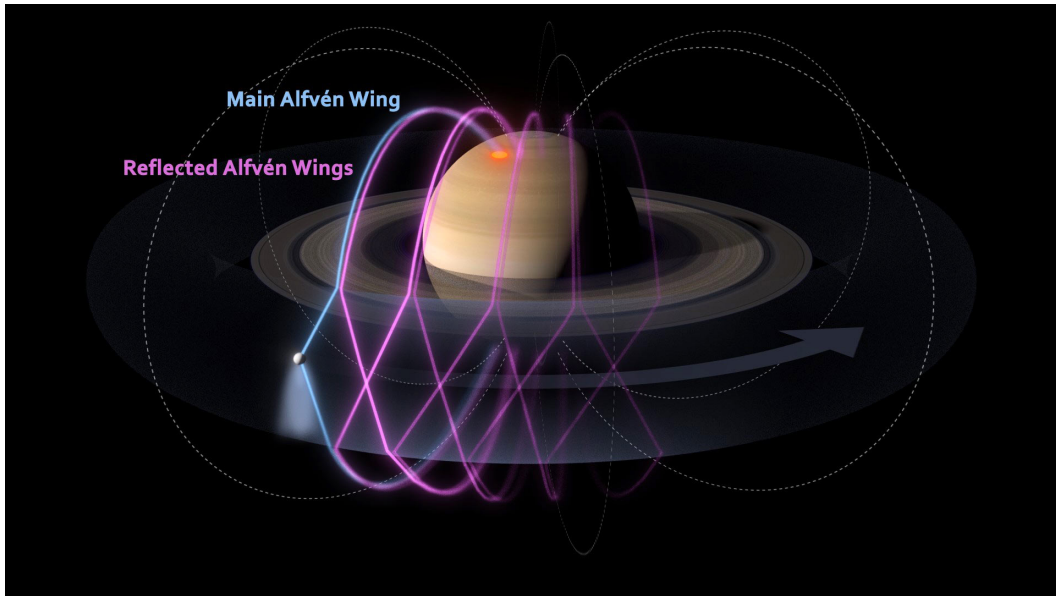


Figure 3.10: Schematic of the Alfvén wing system driven by Enceladus's super-Alfvénic interaction with Saturn's co-rotating magnetospheric plasma. Field-aligned currents carried by the main Alfvén wing (MAW, blue) propagate toward Saturn's ionosphere; partial reflection at the ionosphere and at the torus boundary produces a downstream chain of reflected Alfvén wings (RAWs, magenta). The arrow marks the plasma corotation direction. Image credit: Fabrice Etifier, École Polytechnique.

Chapter 4

Cassini instrument deep dive: ISS and VIMS

Understanding the optical phenomena discussed in this study, especially the faint structures in Saturn’s E ring, requires a clear view of how measurements are formed. Cassini’s remote-sensing instruments provided the raw data, but the final images and spectra result from an end-to-end chain: photons enter the instrument apertures, travel through optics and filters or dispersive elements, are converted to electrons at the detector, and finally become calibrated data numbers on Earth (Brown et al., 2004; Porco et al., 2004). Those photons have two primary origins: solar radiation that has scattered or been absorbed and re-emitted by E ring dust and plume material, and thermal emission from the moons and ring grains themselves, which radiate as blackbodies at their local equilibrium temperature. In the visible and near-infrared, solar scattering dominates; at wavelengths beyond roughly $3\ \mu\text{m}$, thermal emission becomes increasingly significant and must be accounted for in the calibration (Brown et al., 2004).

This chapter narrates that measurement chain. First, the scene physics is set by solar illumination at Saturn and the high-phase geometry that makes forward-scattered structures visible. Then the photon is traced through the two primary instruments used here:

1. Imaging Science Subsystem (ISS): a framing camera system acquiring two-dimensional images through selectable spectral filters (Porco et al., 2004).
2. Visual and Infrared Mapping Spectrometer (VIMS): an imaging spectrometer that builds three-dimensional data cubes by scanning, with separate VIS and IR channels (Brown et al., 2004).

The chapter closes with the calibration logic that converts raw data numbers to radiance and to I/F .

4.1 Scene Physics: From the Sun to the E Ring

Before tracing the photon’s path inside the instruments, the scene itself must first be defined. The appearance of any astronomical target is a product of the light source (the Sun), the properties of the target (the E ring), and the specific viewing geometry. For E ring observations of faint inclined structures, the high-phase forward-scattering geometry is the single most important factor in revealing these features (Rubbrecht et al., 2025).

4.1.1 Illumination and observation geometry

Solar spectral irradiance decreases as R^{-2} . At Saturn’s heliocentric distance, the available flux is about one percent of that at 1 AU. Enceladus flyby sequences targeting the E ring are typically conducted at high phase angles ($\alpha \gtrsim 150^\circ$), where Cassini views nearly toward the Sun through long lines of sight within the ring, and at times the inner rings lie in the background of the field (Rubbrecht et al., 2025). This forward-scattering geometry is deliberately chosen because small grains scatter strongly toward the observer at these angles, greatly enhancing the apparent brightness of otherwise tenuous structures.

Narrow Angle Camera




Figure 4.1: A diagram of the Cassini ISS Narrow Angle Camera (NAC) Head Assembly, showing the light path through the reflective telescope optics, baffles, and filter wheel to the CCD. From Porco et al. (2004).

4.1.2 Forward scattering in the E ring

The E ring consists mainly of micron and sub-micron water-ice grains sourced by Enceladus' plumes. At high phase angles, small particles scatter strongly in the forward direction, boosting the apparent brightness of otherwise tenuous structures. This geometry can reveal faint inclined stripe features; in ISS these are enhanced by subtracting a temporally adjacent CLEAR image from each filtered frame, while in VIMS they strengthen at longer wavelengths, first appearing near $2.2 \mu\text{m}$ and becoming most prominent beyond $4 \mu\text{m}$ (Rubbrecht et al., 2025).

4.2 ISS imaging pathway

4.2.1 Instrument description

The Imaging Science Subsystem Narrow Angle Camera (NAC) forms two-dimensional images through selectable spectral filters. For E ring high-phase observations, typical bands include broadband visible and near-infrared filters (for example, IR3 near $0.93 \mu\text{m}$) along with CLEAR frames used as references (Porco et al., 2004). The optical path of this instrument, a 2-meter focal length reflecting telescope, is shown in Figure 4.1. Photons forward-scattered by E ring dust enter the NAC baffle and reflective telescope, pass the chosen filter, and are blurred by the system point-spread function before sampling at the CCD. The detector converts incident photons into electrons according to the quantum efficiency as a function of wavelength; electronics then map electrons into data numbers with a known gain and read noise (Porco et al., 2004).

4.2.2 Calibration and image reduction

The radiometric calibration follows the published sequence in Porco et al. (2004): decompression and bit weighting, bias and dark subtraction (including residual image effects where applicable), removal of pattern noise or banding, non-linearity correction where required, flat-fielding, conversion from data numbers to electrons, and division by the passband integral of the system response.

For faint, spectrally structured features such as E ring stripe structures, a subtraction of a temporally adjacent CLEAR frame from each filtered image is applied *after* the standard radiometric steps described above. To account for differences in exposure time and passband throughput, the CLEAR frame was scaled by

$$k = \frac{t_X}{t_{\text{CLEAR}}} \frac{\int T_X(\lambda) \text{QE}(\lambda) d\lambda}{\int T_{\text{CLEAR}}(\lambda) \text{QE}(\lambda) d\lambda}, \quad (4.1)$$

and the difference image was then formed as

$$I_{\text{diff}} = I_X - k I_{\text{CLEAR}}. \quad (4.2)$$

This procedure suppresses the smooth scene background and unpolarised stray light while enhancing weak, wavelength-dependent morphology in the target region (Rubbrecht et al., 2025). The propagated uncertainty follows

$$\sigma_{\text{diff}}^2 = \sigma_X^2 + k^2 \sigma_{\text{CLEAR}}^2. \quad (4.3)$$

If required, I_{diff} may subsequently be converted to I/F using the X -band F_{band} (i.e., applying the filter X passband in the expression below) so that contrasts are reported consistently.

The conversion to I/F uses the exact passband weighting for the solar spectrum at Saturn, i.e.

$$F_{\text{band}} = \frac{\int F_1(\lambda) T_0(\lambda) T_1(\lambda) T_2(\lambda) \text{QE}(\lambda) d\lambda}{\pi R^2 \int T_0(\lambda) T_1(\lambda) T_2(\lambda) \text{QE}(\lambda) d\lambda}, \quad I/F = \frac{\pi L_{\text{band}}}{F_{\text{band}}}, \quad (4.4)$$

where F_1 is the solar spectral flux at 1 AU, R is the heliocentric distance in AU, and $T_0 T_1 T_2$ with QE represent the instrument passband weighting (Porco et al., 2004, p. 122).

Porco et al. (2004) also provide a direct calibrated relation between electrons per pixel, exposure time, and target brightness,

$$S = \omega_0 (I/F) t_{\text{exp}}, \quad (4.5)$$

with tabulated ω_0 factors for each NAC filter or filter combination (their Tables for ω_0 ; Porco et al., 2004). These values are the direct link between planned exposure times and expected signal levels for any ISS observation.

4.3 VIMS spectro-imaging pathway

4.3.1 Instrument description

The Visual and Infrared Mapping Spectrometer acquires three-dimensional cubes by scanning the scene. As shown in the instrument block diagram in Figure 4.2, VIMS is composed of two physically separate channels that are co-aligned to observe the same target: the VIS channel for visible light and the IR channel for infrared (Brown et al., 2004).

The VIS channel covers roughly 0.35–1.05 μm with a frame-transfer CCD and slit scanning; the IR channel covers roughly 0.85–5.1 μm using a fixed, triply-blazed grating that disperses light onto a 1×256 InSb linear array (Brown et al., 2004). The photon path is as follows: forward-scattered light from the target passes the fore-optics and scan mirror into the spectrometer entrance; the IR channel grating distributes efficiency across 0.85–5.1 μm onto the cooled InSb array; VIS/IR timing is coordinated so that the VIS lines correspond to the same spatial line as the IR fast scan, and the VIS samples are summed to produce square pixels of approximately 0.5 mrad with spectral bins of approximately 7 nm (Brown et al., 2004).

4.3.2 Calibration and data reduction

The shutter, positioned immediately before the slit, is periodically closed to acquire closed-shutter frames within each fast scan; these measure thermal background and detector dark current and are subtracted from open-shutter measurements as part of calibration (Brown et al., 2004). After background subtraction, channel-dependent responsivities convert to radiance, and radiance is converted to I/F per channel using the solar spectrum at Saturn; additional details of grating blaze distribution, order sorting, and spectral sampling are

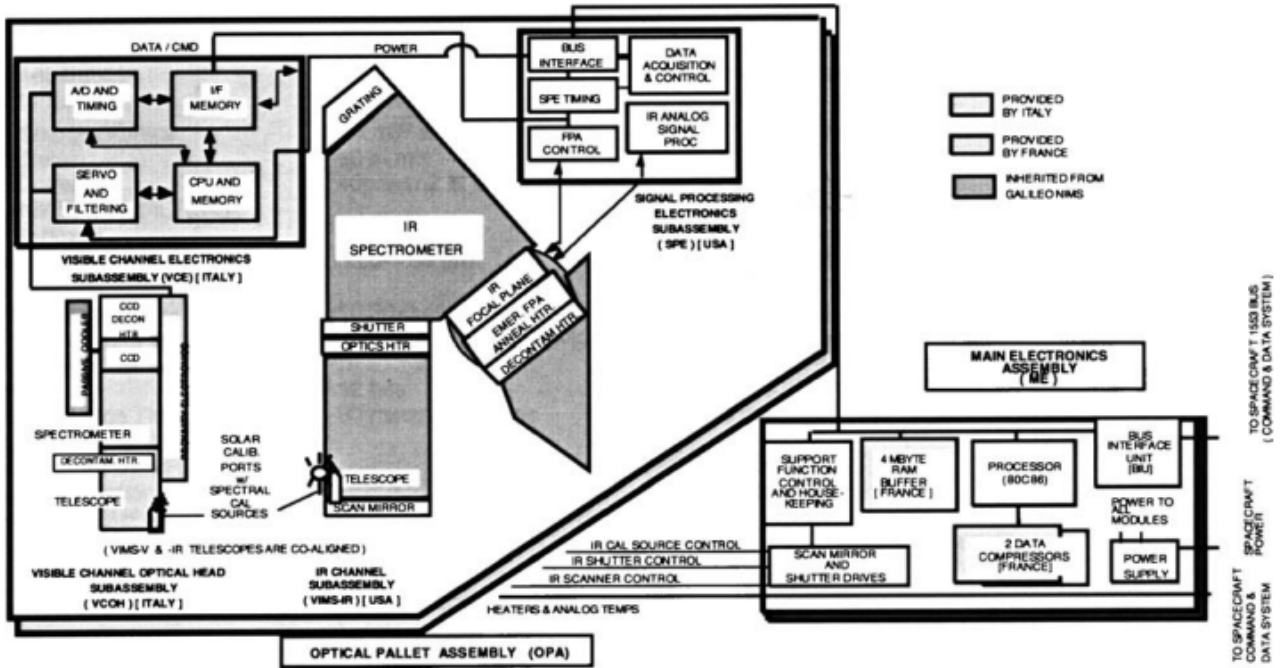


Figure 4.2: A functional block diagram of the integrated VIMS instrument, showing the separate optical heads for the VIS (left) and IR (center) channels, which are co-aligned on the optical pallet assembly. From Brown et al. (2004).

given in the instrument paper (Brown et al., 2004). In high-phase E ring observations, the resulting cubes show the stripe morphology evolving with wavelength: stripes strengthen beyond about $2.2\ \mu\text{m}$ and a single bright band becomes dominant beyond about $4.26\ \mu\text{m}$ (Rubbrecht et al., 2025).

4.4 What an image means here


Each pixel in an ISS frame or VIMS cube encodes the product of scene radiance, instrument throughput, detector conversion, and calibration. The spatial resolution is limited by the point-spread function of the optics and the detector sampling, which smears fine angular structure across neighbouring pixels. The photometric value further depends on the observation geometry: the phase angle between the Sun, the target, and Cassini sets which part of the grain scattering phase function is sampled, while the line-of-sight column length through the ring determines how much material contributes to each pixel’s signal. In the high-phase E ring case, the dominant scene term is forward scattering by small icy grains along long lines of sight, which explains why stripes become clearer with longer wavelengths in VIMS and why subtracting CLEAR helps isolate faint structure in ISS. The spectral overlap of ISS and VIMS around $0.9\text{--}1.0\ \mu\text{m}$ provides a cross-instrument consistency check, while the divergence at $> 3\ \mu\text{m}$ in VIMS reflects both grain-scattering physics and the spectrometer’s thermal-emission contribution (Brown et al., 2004; Porco et al., 2004; Rubbrecht et al., 2025).

4.5 ISS polarimetry

Cassini ISS was designed with linear-polarimetric capability to study scattering by atmospheres, rings, and surfaces. The hardware and its use are documented in the ISS instrument paper, which provides the filter layouts, polariser transmission functions, and the calibration links needed for quantitative work (Porco et al., 2004).

4.5.1 Polariser sets and orientations

A single polariser measurement records only the intensity of light aligned with one transmission axis; it does not distinguish between unpolarised and partially polarised light, and it gives no information about the orientation of polarisation. To fully characterise the linear polarisation state of a scene, at least three images at different



NAC Filters

(Wavelengths, in nm, are central wavelengths computed using the full system transmission function.)

Filter Wheel #1	Filter Wheel #2
1) CL1 (611W)	1) CL2 (611W)
2) Red (650W)	2) GRN (568W)
3) BL1 (451W)	3) UV3 (338W)
4) UV2 (298W)	4) BL2 (440M)
5) UV1 (258W)	5) MT2 (727N)
6) IRP0 (746W)	6) CB2 (750N)
7) P120 (617W)	7) MT3 (889N)
8) P60 (617W)	8) CB3 (938N)
9) P0 (617W)	9) MT1 (619N)
10) HAL (656N)	10) CB1 (619N)
11) IR4 (1002LP)	11) IR3 (930W)
12) IR2 (862W)	12) IR1 (752W)

Figure 4.3: The layout of the Cassini ISS spectral filters on the two filter wheels for both the NAC (top) and WAC (bottom). The locations of the polariser filters are explicitly labeled (P0, P60, P120, IRP0 on the NAC; IRP0, IRP90 on the WAC). From Porco et al. (2004).

polariser orientations are required. This is because linear polarisation is described by three independent quantities: total intensity I , and the two Stokes parameters Q and U , which together encode both the degree of linear polarisation ($\text{DoLP} = \sqrt{Q^2 + U^2}/I$) and the angle of polarisation relative to a reference direction. Three measurements at known, distinct orientations provide a system of equations from which I , Q , and U can be solved unambiguously, regardless of how the spacecraft is oriented in roll.

The Narrow Angle Camera (NAC) therefore features three visible linear polarisers with principal axes separated by 60° , labeled P0, P60, and P120. This triad is sufficient to retrieve I , Q , and U , and hence the full linear polarisation state, independent of spacecraft roll. The Wide Angle Camera (WAC) does not include visible polarisers due to filter-wheel slot constraints, but it carries two orthogonal infrared polarisers, IRP0 and IRP90, which provide I and the Stokes parameter Q referenced to their axes; when the axes are aligned parallel or perpendicular to the scattering plane, Q alone is typically as informative as a full three-angle set (Porco et al., 2004). The layout of these polarisers on their respective filter wheels is shown in Figure 4.3. The physical reference for the visible polariser angles in the NAC is the line in the image plane perpendicular to the spacecraft Z-axis (Porco et al., 2004).

The NAC includes a single infrared polariser (IRP0); by combining a measurement through IRP0 with one through CLEAR, Q can be estimated in the near-IR, albeit with less accuracy than with two orthogonal IR polarisers (Porco et al., 2004).

4.5.2 Optical construction and performance

The visible NAC polarisers are thin-film polymer elements sandwiched between fused silica plates. Their best performance is between 450 and 650 nm, where transmission along the principal axis is approximately 0.45–0.65 and the orthogonal transmission is less than 1%; toward the UV, they transmit too little of either polarization, and toward the near-IR, they leak more orthogonal light. The usable range extends from the UV3 band (~ 350 nm) to CB2 (750 nm) (Porco et al., 2004).

The infrared polarisers used in the IR for both cameras are Polarcor glass (borosilicate impregnated with sub-wavelength metallic wires) cemented between BK7-G18 slabs. The wire-grid mechanism provides efficient extinction across 700–1100 nm, a regime where thin-film polymer polarisers fail due to increasing leakage of the

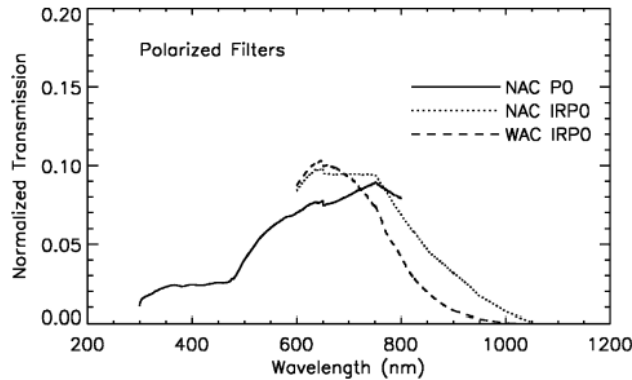


Figure 4.4: The system transmission functions for the NAC visible polariser (P0, solid line), the NAC infrared polariser (IRP0, dotted line), and the WAC infrared polariser (IRP0, dashed line). From Porco et al. (2004).

orthogonal component; the BK7 slabs protect the fragile Polarcor layer and provide polished, optically matched surfaces suitable for a space environment. Over 700–1100 nm the principal transmission is >0.9 and the orthogonal transmission is $\leq 10^{-3}$ (Porco et al., 2004). The system transmission functions for the key polarisers are plotted in Figure 4.4.

4.5.3 How ISS uses polarisers in practice

The underlying principle of polarimetric remote sensing is that unpolarised sunlight becomes partially linearly polarised when it scatters off particles, with the degree and angle of polarisation depending on particle size, shape, and refractive index in ways that total intensity alone cannot reveal. Photometry tells you how much light is scattered; polarimetry tells you by what kind of particle. For atmospheres, the degree of linear polarisation as a function of scattering angle (the polarisation phase curve) separates molecular Rayleigh scattering from aerosol scattering and constrains aerosol microphysics. For rings, the same phase curve carries signatures of particle size distribution and surface roughness, and polarimetry helps isolate singly-scattered from multiply-scattered contributions in optically thin structures. For icy satellites, high albedo surfaces are dominated by multiple scattering, which depolarises light, while single scattering from surface grains preserves polarisation; the balance constrains grain size and packing in a way complementary to reflectance spectra. ISS was designed to exploit all three regimes, with the polariser sets and spectral filters chosen to cover the relevant wavelength ranges for each target type (Porco et al., 2004).

Operationally, the polarisers are designed to be combined with spectral filters on the opposite wheel (for example, methane bands and their continua), enabling wavelength-dependent polarization studies in both NAC and WAC (Porco et al., 2004). The polariser entries also appear in the calibrated throughput and sensitivity tables: the ω_0 factors tabulate the proportionality (Equation (4.5)) for many filter–polariser combinations in NAC and WAC, providing the direct link between measured electrons per pixel, exposure time, and scene I/F for polarimetric setups (Porco et al., 2004). These tables enable the planning and interpretation of polarimetric image sets without introducing additional instrument physics formulas in this chapter.

4.5.4 Status of on-orbit ISS polarimetry and current evidence

A recent end-of-mission reanalysis shows that substantial ISS polarimetric data exist for Enceladus’ surface. From 1,398 NAC polarised images grouped into 466 triplets at 0° , 60° , and 120° , a stringent culling yielded 98 high-quality triplets (46 GRN, 30 UV3, 22 MT2). The reduction accounts for non-ideal transmissions and axis misalignments and was validated on Titan benchmarks. Disk-averaged degrees of linear polarisation (DoLP) are generally low ($\lesssim 5\%$), consistent with multiple scattering on bright icy terrains, and preliminary fits indicate non-spherical grains with characteristic sizes of order $\sim 2 \mu\text{m}$. (Morello et al., 2025)

However, Morello et al. (2025) emphasises that no polarimetric data have yet been analysed for the plumes or for the E ring itself, leaving the scattering properties of these tenuous environments essentially unconstrained by direct polarimetric measurements.

4.5.5 Outlook and potential of ISS polarimetry

ISS’s linear–polarimetric modes were conceived for precisely the kind of faint, structured targets represented by Saturn’s diffuse rings and icy satellites. In the specific context of the faint inclined stripe features identified by Rubbrecht et al. (2025) from simultaneous ISS–VIMS imaging in high-phase Enceladus flybys, polarimetry provides a complementary diagnostic dimension. Polarimetry provides a complementary diagnostic dimension by constraining the degree and orientation of linear polarisation, thereby offering enhanced sensitivity to particle size, morphology, and alignment, while simultaneously stabilising photometric interpretations in high–phase, forward–scattering geometries.

In practical terms, this entails revisiting targeted ISS sequences not only for morphology and photometry, but also for the degree and angle of linear polarisation as functions of wavelength and viewing geometry. Since no polarimetric analysis has yet been performed for either the plumes or the E ring material (Morello et al., 2025), these observations delineate an unexamined observational regime. The existence of a dedicated NAC polarimetric image set for Enceladus and adjacent ring regions within the Cassini archive presents a clear opportunity to extend the stripe analysis with polarimetric constraints, testing whether the inclined structure exhibits a distinctive polarisation phase curve relative to the ambient E ring background.

Should such measurements be undertaken, two practical limitations must be anticipated. First, at the high phase angles characteristic of E ring observations ($\alpha \gtrsim 150^\circ$), the degree of linear polarisation of forward-scattering particles is inherently low: the polarisation signal peaks near 90° scattering angle and falls toward both forward and backward directions, so the expected DoLP in this geometry is modest and may require careful noise control to detect. Second, polariser filter combinations have lower throughput than broadband filters, requiring longer exposure times for equivalent SNR, which increases the risk of scene change or spacecraft motion between the three orientations of a polarimetric triplet.

Chapter 5

Optical properties and scattering framework for the E ring environment

Interpreting the scattering phenomena observed in the E ring, such as the diffraction stripes described in Chapter 3 or the polarimetric signatures discussed in Chapter 4, requires understanding both the physical principles that govern how light interacts with matter and the specific material properties of the grains. This chapter is organised accordingly. Section 5.1 introduces the general physical framework: the complex refractive index, the Mueller matrix formalism, the grating equation, and the Stokes retrieval methodology. Section 5.2 presents the measured optical constants of water ice in its hexagonal and amorphous phases. Section 5.3 combines these two foundations to identify the most appropriate simulation strategy for the E ring context.

5.1 Physical principles of light-matter interaction

5.1.1 Complex refractive index and scattering regimes

The interaction between sunlight and a particle is controlled by two material quantities: the real part $n(\lambda)$ of the complex refractive index $m(\lambda) = n(\lambda) + ik(\lambda)$, which determines the phase velocity and thus refraction, and the imaginary part $k(\lambda)$, which governs absorption. The relative importance of diffraction, refraction, and absorption depends on the size parameter

$$x = \frac{2\pi a}{\lambda}, \quad (5.1)$$

where a is the particle radius and λ is the wavelength. For $x \ll 1$ (Rayleigh regime) scattering is isotropic and scales as λ^{-4} ; for $x \gg 1$ (geometric optics regime) diffraction is confined to a narrow forward lobe and ray tracing applies; for $x \sim 1$ (Mie regime) the full angular dependence must be computed from Maxwell's equations. Sub-micron to micron-sized E ring grains at Cassini wavelengths ($\sim 0.3\text{--}5\ \mu\text{m}$) span the Mie and Rayleigh-to-Mie transition, making the Mie regime the appropriate starting point.

5.1.2 Mueller matrix formalism and degree of linear polarisation

For a collection of randomly oriented particles the scattered Stokes vector \mathbf{S}_{scat} is related to the incident Stokes vector \mathbf{S}_{inc} through the 4×4 Mueller scattering matrix $\mathbf{M}(\alpha, \lambda)$, where α is the solar phase angle. The four Stokes parameters I , Q , U , and V fully describe the polarisation state of an electromagnetic wave. The degree of linear polarisation is defined as

$$\text{DoLP} = \frac{\sqrt{Q^2 + U^2}}{I}. \quad (5.2)$$

For unpolarised incident sunlight, the Mueller relation simplifies to

$$\text{DoLP} = -\frac{S_{12}}{S_{11}}, \quad (5.3)$$

where S_{11} is the total scattered intensity (the phase function) and S_{12} encodes the linear polarisation (Morello et al., 2025). The element S_{12} depends sensitively on the imaginary part $k(\lambda)$ and on particle asphericity. The real part $n(\lambda)$ controls the forward-scattering peak contrast, which is particularly relevant at the large phase

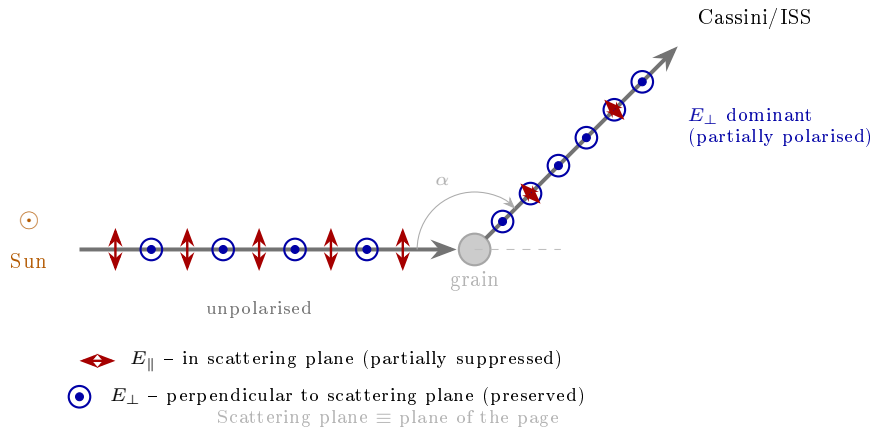


Figure 5.1: Schematic of polarisation generation by single-particle scattering. Unpolarised sunlight arrives from the left and strikes an icy grain. The electric field has two components: E_{\parallel} (red arrows), oscillating in the scattering plane (the plane of the page), and E_{\perp} (blue circles), oscillating perpendicular to it (out of the page). Upon scattering, E_{\parallel} is partially suppressed while E_{\perp} is preserved, so the scattered beam emerges partially linearly polarised perpendicular to the scattering plane. The degree of this asymmetry is quantified by $\text{DoLP} = -S_{12}/S_{11}$ and varies with phase angle α ; suppression is maximum near $\alpha = 90^\circ$ and reduces at the large phase angles ($\alpha \gtrsim 150^\circ$) typical of ISS E ring observations.




Figure 5.2: Geometry of the grating equation. An incident beam arrives at angle θ_i from the grating normal. A periodic structure with spacing d diffracts the beam into orders m ; the first order ($m = 1$) emerges at angle θ_m satisfying $m\lambda = d(\sin \theta_m - \sin \theta_i)$. The zeroth order ($m = 0$, dashed) undergoes specular reflection, emerging at the same angle as the incident beam.

angles ($\alpha \gtrsim 150^\circ$) typical of ISS E ring sequences. Figure 5.1 illustrates the physical origin of this asymmetry between E_{\parallel} and E_{\perp} upon scattering.

While spheres (Mie theory) offer a tractable first-order model, plume-fed grains are expected to be partially crystalline, porous, and non-spherical. Such deviations affect primarily S_{12} and therefore the polarisation behaviour. More general approaches, such as T-matrix calculations for spheroids or the distribution of hollow spheres (DHS), provide a more realistic basis for interpreting observed brightness and polarisation variations.

5.1.3 Periodic density structures as diffraction gratings

One hypothesis for the origin of the luminous bands reported by Rubbrecht et al. (2025) is that the E ring contains a quasi-periodic spatial modulation of grain density or optical depth. Under this hypothesis, such a structure would act as a diffraction grating, and the standard grating equation would apply. In the simplest scalar approximation, the angular position of the m th diffraction order is given by

$$m\lambda = d(\sin \theta_m - \sin \theta_i), \quad (5.4)$$

where d is the spacing of the periodic planes, θ_i is the incidence angle, and θ_m is the emergence angle of order m . Figure 5.2 illustrates this geometry.

Under the grating hypothesis, a three-dimensional density structure inclined with respect to Saturn's equatorial plane would project onto the image plane as a family of parallel, inclined bands. The separation between adjacent bands would scale with wavelength as $\Delta\theta \propto \lambda/d$, which would constitute an observable chromatic

signature of the underlying period d . This chromatic prediction has not yet been confirmed observationally and remains a key test of the hypothesis.

Because the E ring is optically thin, photons observed by ISS typically undergo only one scattering event before reaching the camera. A periodic modulation in grain density would therefore modulate the scattered intensity S_{11} and, through S_{12} , the polarisation state. If the observed stripe-like patterns are the observational manifestation of such a grating, the angular position of the bands relative to the ring plane encodes d through the equation above. This hypothesis motivates two specific analysis strategies applied in Part II of this thesis: the zero-padded Fourier angular-spectrum method used in Chapter 9 to extract the dominant band orientation from ISS residual images, and simultaneous two-wavelength ISS imaging, which would provide a direct test of the chromatic band-spacing prediction of the grating equation.

5.1.4 Stokes parameter retrieval from ISS polariser triplets

The Cassini ISS Narrow Angle Camera records polarimetric data using sets of three images taken through linear polarisers oriented at 0° , 60° , and 120° . These triplets provide the information needed to determine the total intensity I and the orientation of the electric field vector in the camera frame.

Following the end-of-mission calibration of the ISS polarimetric response, Morello et al. (2025) express the total intensity and the polarisation angle in terms of the measured intensities I_0 , I_{60} , and I_{120} and the transmission coefficients T_1 and T_2 of the polariser assembly as

$$I = I_u + I_p = \frac{2(I_0 + I_{60} + I_{120})}{3(T_1 + T_2)}, \quad (5.5)$$

$$\theta_p = \frac{1}{2} \arctan\left(\frac{\sqrt{3}(I_{120} - I_{60})}{-2I_0 + I_{60} + I_{120}}\right). \quad (5.6)$$

Here I_u and I_p denote the unpolarised and linearly polarised components of the disc-averaged intensity, and θ_p is the angle of the electric field vector measured in the NAC image plane (Morello et al., 2025).

The three Malus-law relations underlying these expressions form a linear system that can be inverted analytically to obtain the linear Stokes parameters I , Q , and U for each triplet. These yield the DoLP and the polarisation angle following the definitions given in the preceding subsection. The circular Stokes parameter V cannot be retrieved with ISS, because the camera contains only linear analysers and no retarders.

5.2 Optical constants of water ice

The optical constants required as inputs to the scattering calculations above are not universal constants but are measured, material-specific quantities. For E ring grains, two regimes are critical: the baseline constants for hexagonal ice (I_h) across a broad spectrum, and the phase-dependent behaviour (amorphous versus crystalline) relevant to the cryogenic environment of Enceladus. The interaction of light with water ice is governed by its complex refractive index $m = n + ik$ (Warren & Brandt, 2008), whose values must be measured in the laboratory and compiled into databases that simulations can query.

5.2.1 Hexagonal ice (I_h): broad-spectrum constants

Water ice exists in several crystalline and amorphous phases depending on temperature and deposition history. This section focuses on hexagonal ice (I_h), the thermodynamically stable crystalline phase, for which Warren and Brandt (2008) provide the most complete broadband compilation. Metastable phases such as cubic ice (I_c), which can form as an intermediate between amorphous and hexagonal ice around 130–200 K, are not characterised separately here; their optical properties are bounded by the I_h and amorphous end-members treated in Sections 5.2.1 and 5.2.2.

The standard reference for the optical constants of hexagonal ice (I_h) was updated by Warren and Brandt (2008), superseding the 1984 compilation with improved measurements from the ultraviolet (UV) through to the microwave (Warren & Brandt, 2008). These values are typically reported for a nominal temperature of $T = 266$ K, near the melting point (Warren & Brandt, 2008).




Figure 5.3: The imaginary part of the refractive index (m_{im}) of ice in the ultraviolet and visible regions. The solid line represents the revised compilation by Warren and Brandt (2008), illustrating the significantly weaker absorption compared to the previous values from Warren (1984) (open circles).

5.2.1.1 Ultraviolet and visible transparency

A major revision in the 2008 compilation concerns the transparency of ice in the visible and near UV (200–600 nm). Previous estimates of the absorption coefficient were found to be too high, likely due to scattering effects in laboratory samples (Warren & Brandt, 2008). New data indicate that the absorption minimum occurs near 390 nm, though this value is likely an upper limit set by impurities rather than the ice itself. The true minimum for pure ice may lie between 200 and 300 nm (Warren & Brandt, 2008).

For practical modelling of E ring grains in the visible spectrum, the imaginary part of the refractive index drops to $m_{im} \approx 2 \times 10^{-11}$ (Warren & Brandt, 2008). This implies absorption lengths greater than 1 km, meaning that micron-sized E ring grains act as effectively non-absorbing, purely scattering bodies in this regime (Warren & Brandt, 2008).

5.2.1.2 Near infrared absorption bands

In the near infrared (1.4–7.8 μm), the absorption coefficient k deviates by as much as a factor of two compared to earlier literature (Warren & Brandt, 2008). This spectral region contains vibrational absorption bands critical for remote sensing, particularly the strong 3 μm fundamental band (Warren & Brandt, 2008).

Because forward scattering and absorption are coupled in radiative transfer models, these updates directly impact the retrieval of grain sizes from phase curves.

5.2.1.3 Temperature dependence

While the compilation is anchored at 266 K, the E ring operates at significantly lower temperatures (typically 70–90 K). In weakly absorbing regions, the absorption coefficient k generally increases with temperature by approximately $1\% \text{K}^{-1}$ (Warren & Brandt, 2008). However, this linear relationship breaks down in the far infrared and microwave regions, where phonon bands (lattice vibrations) near 45 and 65 μm shift in position and strength as the ice cools (Warren & Brandt, 2008). Consequently, while Warren and Brandt (2008) provide a robust baseline, specific low-temperature datasets are required for high-fidelity spectral modelling of Enceladus’s output.

5.2.2 Phase differentiation: amorphous versus crystalline ice

The grains ejected by Enceladus may freeze rapidly into amorphous ice (I_a) or anneal into crystalline phases (I_c or I_h) depending on their thermal history. Differentiating these phases requires spectrally resolved data in




Figure 5.4: The imaginary part of the refractive index (m_{im}) of ice in the near infrared spectral region (1.2–2.8 μm). The revised values (solid line) show distinct deviations from the earlier compilation by Warren (1984) (open circles), particularly in the absorption minima windows. From Warren and Brandt (2008).

the near infrared, where the positions and shapes of absorption bands are diagnostic of the phase. Mastrapa et al. (2008) provide optical constants for both phases over the temperature range 20–150 K.

5.2.2.1 Spectral signatures of phase

The infrared spectra of crystalline and amorphous ices are distinct. Crystalline ice features bands that are stronger, sharper, and shifted to longer wavelengths compared to their amorphous counterparts (Mastrapa et al., 2008). Conversely, amorphous ice bands are weaker and shifted to shorter wavelengths (Mastrapa et al., 2008).

Specifically, the bands near 1.50 μm , 1.56 μm , 1.65 μm , and 2.0 μm all exhibit minimum shifts of 20 cm^{-1} between phases (Mastrapa et al., 2008). Furthermore, the boxiness of the combined 1.50–1.56 μm feature is reduced in amorphous ice as the 1.56 μm shoulder weakens (Mastrapa et al., 2008).

5.2.2.2 The 1.65 μm discriminator

The most practical diagnostic for phase identification in Cassini VIMS data is the band near 1.65 μm (6050 cm^{-1}).

- Sensitivity: This band is strongly temperature-dependent in crystalline ice, shifting to longer wavelengths and increasing in strength as temperature decreases (Mastrapa et al., 2008).
- Mixtures: The band does not vanish immediately upon the introduction of amorphous material. It remains clearly present in models with up to 80% amorphous content (Mastrapa et al., 2008). Therefore, the mere detection of the 1.65 μm feature does not prove a surface is 100% crystalline (Mastrapa et al., 2008).
- Ratio metric: To robustly distinguish phases, Mastrapa et al. (2008) propose using the ratio of the band area at 1.65 μm to the band area at 1.50 μm (6600 cm^{-1}). This ratio provides a clear separation between amorphous and crystalline trends across the relevant temperature range (20–130 K) and remains valid for various grain sizes (Mastrapa et al., 2008).

5.2.2.3 Amorphous sub-phases

Laboratory data also reveal two distinct regimes within the amorphous phase, likely corresponding to High Density Amorphous (HDA) ice below 70 K and Low Density Amorphous (LDA) ice above 70 K (Mastrapa

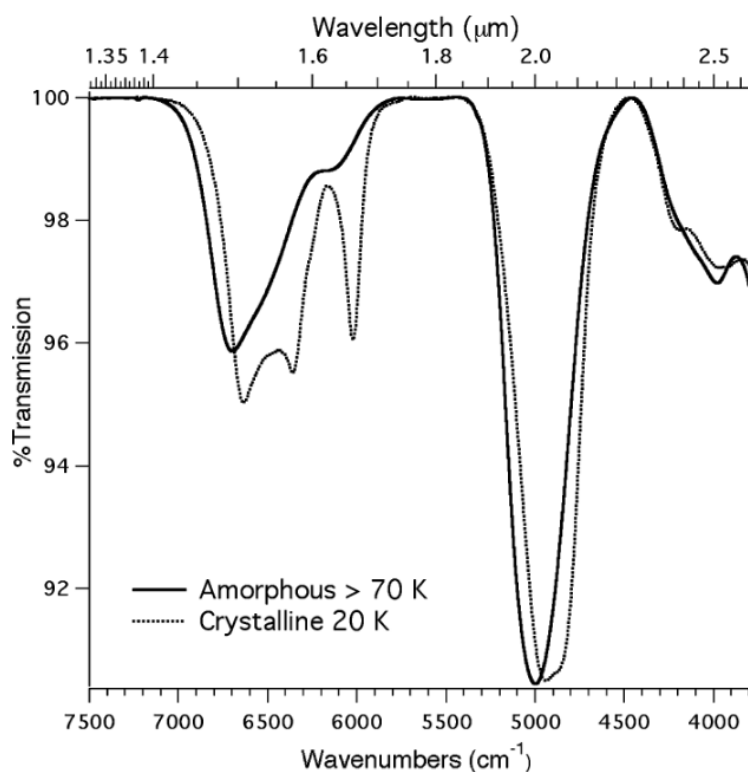


Figure 5.5: Comparison of the infrared transmission spectra of amorphous (> 70 K, solid line) and crystalline (20 K, dotted line) H_2O ice. The crystalline phase exhibits sharper, stronger absorption bands, most notably in the $1.5\text{--}1.65\ \mu\text{m}$ region, where the distinct boxiness of the combined bands is evident, contrasting with the smoother, weaker profile of amorphous ice.

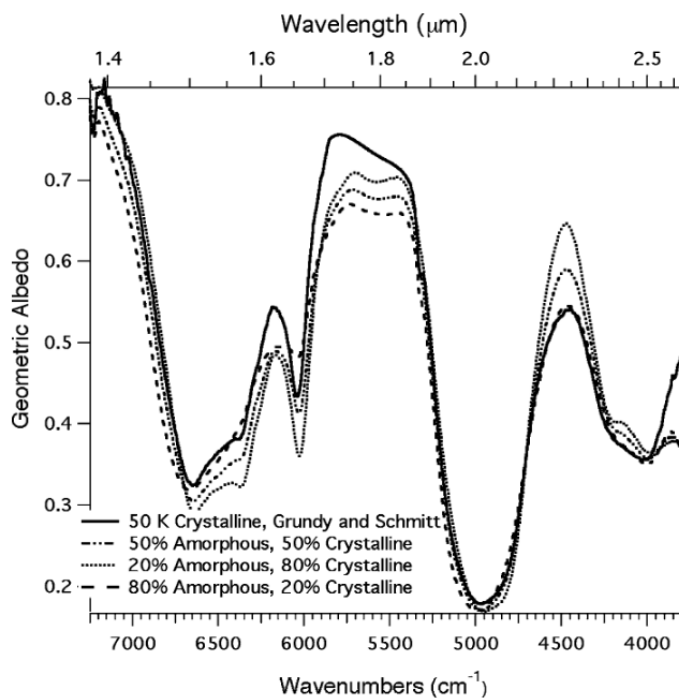


Figure 5.6: Geometric albedo model spectra of H_2O ice mixtures at 50 K. The band near $1.65\ \mu\text{m}$ remains visible even in the spectrum dominated by amorphous ice (80% amorphous, long dashed line), demonstrating that the detection of this feature is not sufficient to claim a surface is fully crystalline.




Figure 5.7: The band area ratio of the $1.65 \mu\text{m}$ band divided by the $1.50 \mu\text{m}$ band as a function of temperature. This metric clearly distinguishes between crystalline ice (filled symbols, exhibiting a strong temperature dependence) and amorphous ice (open circles, displaying a nearly constant low value), providing a robust diagnostic tool for phase identification (Mastrapa et al., 2008).

et al., 2008). Transitions between these regimes are irreversible. Samples deposited below 70 K and then heated will shift their spectral features to match those of the high-temperature phase (Mastrapa et al., 2008).

5.3 Application to E ring modelling

The two preceding sections supply the inputs needed for E ring simulations: the physical formalism (Section 5.1) and the material constants (Section 5.2). For broadband ISS imaging in the visible, the near-zero absorption of hexagonal ice established by Warren and Brandt (2008) confirms that scattering cross sections are dominated by geometry rather than absorption. For VIMS spectral analysis, the Mastrapa et al. (2008) datasets support mixture modelling in which the shifting position of the $1.65 \mu\text{m}$ band and the $1.65/1.50 \mu\text{m}$ area ratio serve as proxies for the crystallinity fraction and temperature of plume-derived grains.

Combining these inputs, the problem of interpreting E ring brightness and polarisation patterns reduces to two coupled models: (i) single-particle scattering computed from the Mueller matrix using the appropriate optical constants, and (ii) the wave-optical imprint of a larger-scale periodic density structure captured by the grating equation. The simulation methodologies available to implement these two components are surveyed below.

5.3.1 Simulation methodologies

The optical behaviour outlined above can be explored with a hierarchy of simulation tools, each valid in a specific range of particle sizes and spatial scales.

5.3.1.1 Geometrical optics and ray tracing

In the limit of very large size parameters, $x = 2\pi a/\lambda \gg 1$, individual grains can be treated within geometrical optics. Rays are traced through refracting and reflecting surfaces, and diffraction is confined to narrow edge effects. Commercial tools such as CODE V and Zemax implement this regime and are frequently used for the design and tolerancing of millimetre-scale and larger optical components.

For the micron-sized grains dominating Saturn’s E ring, this regime is not appropriate for the single-particle

scattering problem itself, but ray tracing remains relevant for modelling the internal propagation of light within the telescope and camera optics.

5.3.1.2 Wave optical propagation

When the structural scales in the medium are comparable to a few wavelengths, diffraction and interference must be treated explicitly. Fourier-based propagation methods, such as Fresnel and Fraunhofer propagation or the angular spectrum approach, describe the evolution of scalar or vector fields between planes by means of Fourier transforms.

In the context of the E ring, these methods are well suited to model the far-field pattern produced by a periodic modulation of optical depth along the line of sight. The periodic structure acts as an effective diffraction grating, and the wave propagator maps the field from the grating plane to the Cassini image plane, allowing direct comparison with the observed stripe spacing and contrast.

5.3.1.3 Full field electromagnetic solvers

For sub-micron to micron-sized grains with complex shapes, porosity, or internal structure, it can be necessary to compute the full vector electromagnetic field. Finite difference time domain (FDTD) and related Maxwell solvers discretise space and time and evolve the electric and magnetic fields directly. They provide detailed information on near fields, resonances, and depolarisation effects, at the cost of substantial computational resources.

In a multi-scale modelling framework these field solvers can be combined with Fourier propagation. Local field maps at the level of individual grains or grain aggregates are first computed with FDTD or related solvers, and the resulting fields are then propagated to larger distances using wave optical methods.

5.3.2 Applicability to the present study

For the work in this thesis, the choice of simulation tool is guided by the relevant physical scales. Single-particle scattering by micron-sized icy grains is most efficiently handled by established light-scattering codes that operate in terms of the Mueller matrix and the optical constants from Warren and Brandt (2008) and Mastrapa et al. (2008). The large-scale imprint of a quasi-periodic density structure in the E ring is naturally treated with Fourier-based propagation methods. Ray tracing is expected to play a limited supporting role for the telescope and camera optics, while full-field electromagnetic solvers are reserved for targeted studies of grain morphology where their additional complexity is warranted.

A compact overview of the different simulation regimes, including their physical applicability and their relevance to this thesis, is provided in Table 5.1.

Table 5.1: Summary of simulation regimes relevant to the scattering and diffraction modelling of the E ring. The choice of tool depends on the size parameter $x = 2\pi a/\lambda$ and on the spatial scale of the periodic structure.

Regime	Physical condition	Representative tools	Relevance to this thesis
Geometrical optics (ray-based)	$x \gg 1$; millimetre-scale grains or larger	CODE V, Zemax	Not applicable to E ring grain scattering. Used only to model the propagation of light inside the ISS telescope and camera optics.
Wave optical propagation	Structure size comparable to a few wavelengths. Micron-scale grains, but periodic structures much larger than λ .	Fourier propagation, Fresnel and Fraunhofer diffraction, angular spectrum methods, RCWA	Primary tool for modelling the diffraction-like banding reported by Rubbrecht et al. (2025). Used to relate the observed stripe spacing to an underlying periodic density modulation.
Full electromagnetic field solvers	$x \sim 1$; sub-micron to micron grains with complex morphology or porosity.	FDTD, finite element Maxwell solvers	Supporting tool for targeted studies of grain morphology. Provides near-field solutions that can be propagated with wave optical methods. Not used for bulk E ring modelling.
Single particle scattering (Mueller matrix)	Applicable for isolated grains with known refractive index and size distribution.	Mie theory, T-matrix, distribution of hollow spheres (DHS)	Baseline computation for S_{11} and S_{12} across wavelength. Supplies the scattering inputs used by the higher-level periodic structure model.

Chapter 6

Thesis outlook

The preceding chapters introduced the physical and observational context required for analysing Saturn's E ring environment. They summarised the properties of plume-derived water ice, the scattering regimes relevant for micron sized grains, the radiative transfer behaviour of icy particles, the modelling approach used in earlier work, the potential of polarimetry for detecting microphysical structure, and the Cassini ISS imaging pathway used to observe the system. These elements together form the conceptual toolbox for interpreting any structures that emerge during the analysis, ensuring that unexpected phenomena can be placed within a consistent physical framework.

A key outcome of the literature review is the identification of a clear research gap. Recent imaging analyses have revealed stripe like, forward-scattered luminous bands in a small set of Cassini orbits near Enceladus. These bands were first reported only months ago by Rubbrecht et al. (2025), who proposed that they may trace an inclined periodic structure embedded within the E ring. Their study established the first indication of such a phenomenon, but was necessarily limited by time, scope and the available dataset. As a result, the physical origin, geometric configuration and broader prevalence of the bands remain unverified. No published study has yet combined imaging, orbit geometry and polarimetry to evaluate the plausibility of this hypothesis.

Moreover, Cassini's polarisation measurements have never been applied to the E ring or to plume-fed grains. This leaves a substantial interpretative gap, since polarisation is directly sensitive to particle size, shape and refractive index. The potential of polarimetric interpretation in this environment therefore remains unexplored.

Together, these gaps motivate a combined imaging, geometric and polarimetric investigation aimed at determining whether the E ring contains an inclined periodic structure and, if so, what physical mechanisms could give rise to it.

6.1 Scientific objectives and research questions

6.1.1 Main research question

How can Cassini ISS imaging and polarimetry be used to reveal the origin of the peculiar scattering stripes in Saturn's E ring, and which physical processes explain this phenomenon?

6.1.2 Sub questions

RQ1. How do the luminous bands behave across Cassini ISS observations in terms of geometry, phase angle, wavelength dependence and repeatability across trajectories?

RQ2. Which polarimetric quantities can be reliably retrieved from ISS polariser sequences, and what are the states of polarisation and their degrees across locations where luminous bands are detected?

RQ3. Which combinations of particle size distributions, particle morphologies and periodic or inclined geometries can reproduce the observed intensity and polarisation signatures in scattering simulations?

RQ4. What is the three dimensional location and geometric configuration of the hypothesised periodic structure

within the E ring, and which structural parameters are required to reproduce the observed inclination, spacing and vertical behaviour of the luminous bands across different observation geometries?

RQ5. What do the imaging, geometric and polarimetric constraints reveal about the physical origin and formation mechanism of the hypothesised periodic structure?

RQ6. How would such a structure evolve within the E ring environment, and how does it interact with the transport and redistribution of plume-derived particles?

6.1.3 Scientific objectives

The objectives of the research follow directly from the questions above.

Scientific objective:

To determine whether an inclined periodic structure is present in Saturn's E ring and to identify the physical mechanism that produces the observed forward-scattering stripes, by combining Cassini ISS imaging with calibrated polarimetric diagnostics.

Methodological objective:

To develop and apply a complete processing chain for Cassini ISS polarimetry, retrieving reliable Stokes parameters and integrating these with geometric mapping tools suited for analysing periodic structures in a three dimensional ring environment.

Modelling objective:

To construct and evaluate forward-scattering simulations linking plausible particle populations and periodic geometries to synthetic observables, allowing direct comparison with the measured intensity and polarisation signatures.

Interpretation objective:

To synthesise the observational and modelling outcomes into a coherent physical explanation for the stripe-like features, assessing their origin, formation pathways and dynamical persistence.

Transport objective:

To evaluate how the existence of a periodic structure interacts with the transport, redistribution and evolution of plume-derived material throughout the E ring, and to identify which transport scenarios remain consistent with the inferred geometry and scattering behaviour.

6.2 Work packages and breakdown structure

This section summarises the complete work breakdown structure for the thesis. The structure is organised into six major work packages labelled A through F. Each work package is decomposed into blue level tasks and further into yellow level atomic tasks. These tasks represent the smallest actionable research items used for planning, Gantt chart construction, and progress monitoring.

The work packages do not represent strict chronological sequences. Instead they group conceptually related research activities. Chronology will be established through the Gantt chart, where yellow level tasks from different packages may interleave depending on dependencies or scientific needs.

6.2.1 Work package A: data discovery, acquisition, and preprocessing

A01. Comprehensive Dataset Discovery

This sub-package extends the work of Rubbrecht et al. by compiling a complete list of Cassini orbits and imaging sequences that may show forward-scattering stripes. It includes screening for relevant observation geometries and documenting the associated metadata in a consistent way.

A02. ISS & VIMS Data Retrieval

This sub-package focuses on retrieving the selected ISS and, where relevant, VIMS datasets from the Planetary

Data System and any mission archives. The data are organised into a structured repository that supports reproducible analysis and cross-referencing between instruments.

A03. Calibration & Preprocessing

This sub-package applies and refines the calibration pipeline for the selected ISS data, including CISSCAL processing, geometric calibration, and co-registration across filters and exposures. The outcome is a set of science-ready images and ancillary products that can be used consistently in later work packages.

6.2.2 Work package B: polarimetric data products

B01. Select and Prepare Polarimetric Sequences

This sub-package identifies all ISS polariser triplets relevant to the E ring and plume environment, focusing on well-exposed scenes and stable pointing. It performs initial quality control, including checks on alignment, exposure level and suitability for Stokes retrieval.

B02. Stokes Retrieval Pipeline

This sub-package implements the mathematical inversion from measured intensities in the polariser triplets to Stokes parameters. It includes calibration of transmission factors, alignment corrections and validation tests to ensure that the retrieved Stokes vectors are physically consistent.

B03. Scene-Level Polarimetric Mapping

This sub-package constructs spatial maps of polarisation quantities over the scenes of interest, such as Degree and Angle of Polarisation. It registers these maps to the geometric reference frames developed in Work Package A, enabling comparison with the locations and morphology of the luminous bands.

B04. Optional DPG Pre-Processing

This sub-package explores differential and geometry-aware polarimetric diagnostics, including elements inspired by the Differential Polarimetric Geometry framework where appropriate. It is designed as an innovative extension that can be activated if time and results permit, providing a pathway to enhanced contrast or curvature-based interpretation.

6.2.3 Work package C: luminous band detection & geometric characterisation

C01. Photometric Band Detection

This sub-package develops and applies methods to detect stripe-like forward-scattering features in the calibrated ISS images, both in previously known orbits and in newly identified sequences. It quantifies band visibility, contrast and basic morphological properties as a function of viewing geometry and wavelength.

C02. Geometric Reconstruction of Band Structure

This sub-package reconstructs the three dimensional geometry of the bands by mapping their projected locations into Saturn-centred coordinate systems. It investigates inclination, spacing and vertical behaviour of the putative structure and assesses how these quantities behave across different orbits.

C03. Polarimetric Cross-Checks Triggered by Band Findings

This sub-package links the band detections from C01–C02 with the polarimetric products from Work Package B. It tests whether band locations correlate with specific polarisation signatures and uses these correlations to refine hypotheses about the underlying structure.

6.2.4 Work package D: optical scattering & band formation simulations

D01. Define Simulation Requirements & Choose Modelling Approach

This sub-package translates the observational findings into requirements for optical models, including relevant particle sizes, morphologies, structure geometries (, and scattering refractive and diffractive process and eventually photonic). It selects appropriate optical and radiative transfer methods and defines the parameter space to be explored.

D02. Build the Forward Simulation Framework

This sub-package implements the chosen simulation framework, linking particle and structure models to synthetic Cassini-like observables. It includes the generation of intensity and polarisation predictions for a range

of viewing geometries comparable to the actual ISS observations.

D03. Simulation ↔ Observation Comparison

This sub-package systematically compares the simulated observables to the measured band characteristics and polarimetric maps. It identifies which combinations of particle and structure parameters remain consistent with the data and documents rejected scenarios.

6.2.5 Work package E: integrated scientific interpretation

E01. Integrate Geometric & Simulated Results

This sub-package combines the geometric reconstruction of the bands with the constraints from the simulation framework. It evaluates which structural configurations are most plausible in light of both the recovered geometry and the optical modelling.

E02. Integrate Polarimetric Findings

This sub-package incorporates the polarimetric diagnostics into the interpretation, using sensitivity to particle size, shape and refractive index to further discriminate between competing scenarios. It assesses how robustly the polarimetric evidence supports or rejects candidate structures.

E03. Integrate Optional DPG Findings (if used)

If DPG-inspired diagnostics are applied, this sub-package integrates their results with the more conventional analyses. It examines whether differential or curvature-like quantities add independent constraints on the structure or particle properties.

E04. Scientific Hypothesis Testing and Final Interpretation

This sub-package synthesises all available evidence into a coherent physical picture of the luminous bands and the hypothesised periodic structure. It formulates and evaluates specific hypotheses about formation mechanisms and assesses their plausibility in the broader context of Enceladus plume physics and E ring dynamics.

6.2.6 Work package F: thesis writing, milestones, and defence

F01. Writing & Synthesis

This sub-package covers the drafting, revision and integration of the thesis chapters, running in parallel with the technical work. It ensures that methods, results and interpretations are documented clearly and consistently with the faculty requirements.

F02. Midterm Review Preparation

This sub-package prepares the deliverables, documentation and presentation material required for the Midterm Review. It aligns the internal WBS progress with the formal TU Delft milestone and incorporates feedback into the subsequent planning.

F03. Green Light Preparation

This sub-package prepares the draft thesis and supporting material needed for the Green Light Review. It focuses on ensuring that the scientific narrative, structure and preliminary conclusions are sufficiently mature for approval to proceed to finalisation.

F04. Final Thesis & Defence

This sub-package covers the submission of the final thesis, repository upload and plagiarism checks, followed by the public defence and closed examination. It concludes the project and delivers the formal academic output of the research.

6.3 Execution strategy and timeline

6.3.1 Adaptive research strategy

While the work packages defined above provide a structured breakdown of tasks, the scientific execution need not be strictly linear. Given the exploratory nature of applying polarimetry to the E ring for the first time, an adaptive strategy is adopted to maximise scientific return within the available time constraints.

This strategy, illustrated in Figure 6.1, prioritises an initial application of the polarimetric pipeline (WP B) to the previously identified dataset defined by Rubbrecht et al. (2025). If this initial analysis reveals strong, novel polarimetric signatures sufficient to address the main research questions, the project focus will shift immediately to simulation and interpretation (WPs D and E). In this scenario, the extensive data expansion and geometric re-characterisation (WPs A and C) become lower priority extensions.

Conversely, if the initial polarimetric results are ambiguous on the limited existing dataset, the full execution of WPs A and C will be triggered to provide the necessary statistical weight and geometric coverage. Crucially, the polarimetric pipeline (WP B) will then be re-applied to these newly discovered features to ensure a consistent dataset before moving to simulation. This conditional approach ensures that effort is focused on the most promising investigative avenues.

6.3.2 Project timeline

Figure 6.2 translates the work packages defined in Section 6.2 into a chronological schedule. The timeline spans a total of 32 active working weeks, commencing with the literature study kick-off in Week 42 and concluding with the final defence in Week 22. The schedule accounts for the winter break (Weeks 52 and 01) by allocating a reduced load equivalent to one working week, while the spring break is fully utilised for research continuity.

Notable in this planning is the parallel execution of tasks around the Midterm Review (Week 8). This milestone is deliberately scheduled to overlap with the finalisation of Band Detection (WP C). Furthermore, the timeline is backwards-planned from the Green Light deadline (Week 18), necessitating a halt of new analysis by Week 16 to meet the Faculty's requirement for a complete thesis draft twenty working days prior to defence.

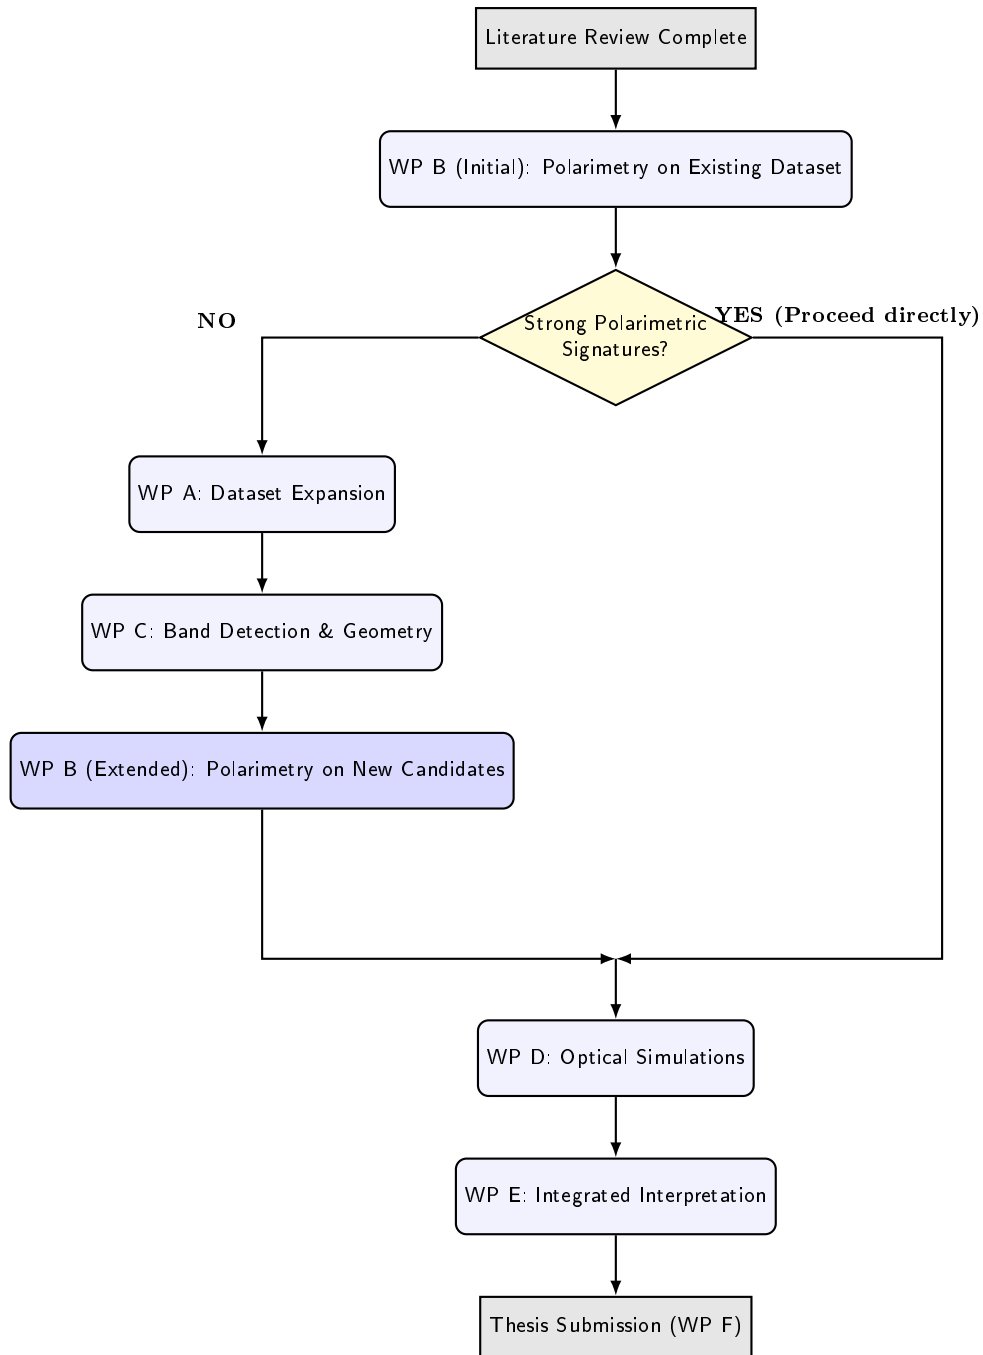


Figure 6.1: Adaptive research strategy flow diagram. The execution includes a critical decision point following initial polarimetric analysis. If the existing dataset yields insufficient signal, the dataset is expanded (WP A) and new bands are detected (WP C) and characterised polarimetrically (WP B extended) before proceeding to simulation.




Figure 6.2: Gantt chart illustrating the research timeline. The schedule maps Work Packages A through F over the 32-week thesis duration, highlighting the critical path from data acquisition to the final defence, including the mandatory Midterm and Green Light milestones.

Bibliography

- Asphaug, E., & Reufer, A. (2013). Late origin of the saturn system. *Icarus*, 223(1), 544–565. <https://doi.org/10.1016/j.icarus.2012.12.009>
- Batygin, K., & Morbidelli, A. (2020). Formation of giant planet satellites. *The Astrophysical Journal*, 894(2), 143. <https://doi.org/10.3847/1538-4357/ab8937>
- Blanc, M., Crida, A., Shibaïke, Y., Charnoz, S., El Moutamid, M., Estrada, P., Mousis, O., Salmon, J., Schneeberger, A., & Vernazza, P. (2025, March). *Understanding the formation of saturn's regular moons in the context of giant planet moons formation scenarios* [Preprint]. arXiv: 2503.01351 [astro-ph.EP]. Retrieved November 12, 2025, from <https://arxiv.org/abs/2503.01351>
- Brown, R. H., Baines, K. H., Bellucci, G., Bibring, J.-P., Buratti, B. J., Capaccioni, F., Cerroni, P., Clark, R. N., Coradini, A., Cruikshank, D. P., Drossart, P., Formisano, V., Jaumann, R., Langevin, Y., Matson, D. L., McCord, T. B., Mennella, V., Miller, E., Nelson, R. M., ... Sotin, C. (2004). The Cassini visual and infrared mapping spectrometer (VIMS) investigation. *Space Science Reviews*, 115, 111–168.
- Canup, R. M. (2010). Origin of saturn's rings and inner moons by mass removal from a lost titan-sized satellite. *Nature*, 468(7326), 943–946. <https://doi.org/10.1038/nature09661>
- Canup, R. M., Richter, K., Dauphas, N., & et al. (2023). Origin of the moon. *Reviews in Mineralogy and Geochemistry*, 89(1), 53–102. <https://doi.org/10.2138/rmg.2023.89.02>
- Canup, R. M., & Ward, W. R. (2002). Formation of the galilean satellites: Conditions of accretion. *The Astronomical Journal*, 124(6), 3404–3423. <https://doi.org/10.1086/344684>
- Charnoz, S., Salmon, J., & Crida, A. (2010). The recent formation of saturn's moonlets from viscous spreading of the main rings. *Nature*, 465(7299), 752–754. <https://doi.org/10.1038/nature09096>
- Cilibrasi, M., Szulágyi, J., Mayer, L., & et al. (2018). Satellites form fast & late: A population synthesis for the galilean moons. *Monthly Notices of the Royal Astronomical Society*, 480(4), 4355–4368. <https://doi.org/10.1093/mnras/sty2163>
- Colombo, G., & Franklin, F. (1971). On the formation of the outer satellite groups of jupiter. *Icarus*, 15, 186–189. [https://doi.org/10.1016/0019-1035\(71\)90073-X](https://doi.org/10.1016/0019-1035(71)90073-X)
- Crida, A., & Charnoz, S. (2012). Formation of regular satellites from ancient massive rings in the solar system. *Science*, 338(6111), 1196. <https://doi.org/10.1126/science.1226477>
- Čuk, M., Dones, L., & Nesvorný, D. (2016). Dynamical evidence for a late formation of saturn's moons. *The Astrophysical Journal*, 820(2), 97. <https://doi.org/10.3847/0004-637X/820/2/97>
- de Pater, I., Martin, S. C., & Showalter, M. R. (2004). Near-infrared observations of saturn's e and g rings during earth's ring-plane crossing in august 1995. *Icarus*, 172, 446–454. <https://doi.org/10.1016/j.icarus.2004.04.018>
- Dougherty, M. K., Buratti, B. J., Seidelmann, P. K., & Spencer, J. R. (2018). Enceladus as an Active World: History and Discovery. In P. M. Schenk, R. N. Clark, C. J. A. Howett, A. J. Verbiscer, & J. H. Waite (Eds.), *Enceladus and the icy moons of saturn* (p. 3). https://doi.org/10.2458/azu_uapress_9780816537075-ch001
- Dougherty, M. K., Khurana, K. K., Neubauer, F. M., Russell, C. T., Saur, J., Leisner, J. S., & Burton, M. E. (2006). Identification of a dynamic atmosphere at enceladus with the cassini magnetometer. *Science*, 311(5766), 1406–1409. <https://doi.org/10.1126/science.1120985>
- Dražkowska, J., & Szulágyi, J. (2018). Dust evolution and satellitesimal formation in circumplanetary disks. *The Astrophysical Journal*, 866(2), 142. <https://doi.org/10.3847/1538-4357/aae0fd>
- Fuller, J., Luan, J., & Quataert, E. (2016). Resonance locking as the source of rapid tidal migration in the jupiter and saturn moon systems. *Monthly Notices of the Royal Astronomical Society*, 458(4), 3867–3879. <https://doi.org/10.1093/mnras/stw609>
- Gao, P., Kopparla, P., Zhang, X., & Ingersoll, A. P. (2016). Aggregate particles in the plumes of enceladus. *Icarus*, 264, 227–238. <https://doi.org/10.1016/j.icarus.2015.09.030>

- Hadid, L. Z., Chust, T., Wahlund, J.-E., Morooka, M. W., Roussos, E., Witasse, O., Rabia, J., Pisa, D., Kim, K., Edberg, N. J. T., Rymer, A. M., Lamy, L., Kotsiaros, S., Aizawa, S., Jeandet, A., Modolo, R., André, N., Canu, P., Bowers, C. F., ... Dougherty, M. K. (2026). Evidence of an extended alfvén wing system at enceladus: Cassini’s multi-instrument observations [e2025JA034657 2025JA034657]. *Journal of Geophysical Research: Space Physics*, *131*(2), e2025JA034657. <https://doi.org/https://doi.org/10.1029/2025JA034657>
- Haffert, S. Y., Bohn, A. J., de Boer, J., & et al. (2019). Two accreting protoplanets around the young star PDS 70. *Nature Astronomy*, *3*, 749–754. <https://doi.org/10.1038/s41550-019-0780-5>
- Hamilton, D. P. (1993). Motion of dust in a planetary magnetosphere—orbit-averaged equations for oblateness, electromagnetic, and radiation forces, with application to saturn’s e ring. *Icarus*, *101*, 244–264. <https://doi.org/10.1006/icar.1993.1017>
- Hamilton, D. P., & Burns, J. A. (1994). Origin of saturn’s e ring: Self-sustained, naturally. *Science*, *264*(5158), 550–553. <https://doi.org/10.1126/science.264.5158.550>
- Hansen, C. J., Shemansky, D. E., Esposito, L. W., Stewart, A. I. F., Lewis, B. R., Colwell, J. E., & et al. (2011). The composition and structure of the enceladus plume. *Geophysical Research Letters*, *38*, L11202. <https://doi.org/10.1029/2011GL047415>
- Hedman, M. M., Gosmeyer, C. M., Nicholson, P. D., Sotin, C., Brown, R. H., Clark, R. N., Baines, K. H., Buratti, B. J., & Showalter, M. R. (2012). An observed correlation between the vertical structure and the optical depth of saturn’s e ring. *Icarus*, *217*, 322–338. <https://doi.org/10.1016/j.icarus.2011.11.014>
- Herschel, J. F. W. (1847). *Results of astronomical observations made during the years 1834, 5, 6, 7, 8, at the cape of good hope* [Accessed October 2025. Page 415 contains the proposed nomenclature for Saturn’s satellites, including Enceladus.]. Smith, Elder; Co. <https://archive.org/details/Resultsastronom00Hers/page/415/mode/2up>
- Herschel, S. W. (1790). *On the satellites of the georgian planet* (Vol. 80) [Accessed October 2025]. Philosophical Transactions of the Royal Society of London. https://archive.org/details/bim_eighteenth-century_on-the-satellites-of-the_herschel-sir-william-1_1790/page/n1/mode/2up
- Hsu, H.-W., Postberg, F., Sekine, Y., Shibuya, T., Kempf, S., Horányi, M., Juhász, A., Altobelli, N., Suzuki, K., Masaki, Y., Kuwatani, T., Tachibana, S., Sirono, S.-i., Moragas-Klostermeyer, G., & Srama, R. (2015). Ongoing hydrothermal activities within enceladus. *Nature*, *519*(7542), 207–210. <https://doi.org/10.1038/nature14262>
- Iess, L., Stevenson, D., Parisi, M., Hemingway, D., Jacobson, R., Lunine, J., Nimmo, F., Asmar, S., Ducci, M., & Tortora, P. (2014). The gravity field and interior structure of enceladus. *Science (New York, N.Y.)*, *344*, 78–80. <https://doi.org/10.1126/science.1250551>
- Ingersoll, A. P., & Ewald, S. K., Shawn P. and Trumbo. (2020). Time variability of the enceladus plumes. *Icarus*, *344*, 113345. <https://doi.org/doi.org/10.1016/j.icarus.2019.06.006>
- Ingersoll, A. P., & Ewald, S. P. (2017). Decadal timescale variability of the enceladus plumes: Analysis of cassini iss images. *Icarus*, *282*, 260–275. <https://doi.org/10.1016/j.icarus.2016.09.018>
- Jia, Y.-D., Russell, C. T., Khurana, K. K., Wei, H. Y., Ma, Y. J., Leisner, J. S., Persoon, A. M., & Dougherty, M. K. (2011). Cassini magnetometer observations over the enceladus poles. *Geophysical Research Letters*, *38*(19). <https://doi.org/https://doi.org/10.1029/2011GL049013>
- Juhász, A., & Horányi, M. (2007). Seasonal variations in saturn’s e ring. *Geophysical Research Letters*, *34*, L09104. <https://doi.org/10.1029/2007GL029285>
- Kempf, S., Beckmann, U., Schmidt, J., & et al. (2008). The e ring in the vicinity of enceladus: I. icy grain properties derived from cassini cda measurements. *Icarus*, *193*(2), 420–437. <https://doi.org/10.1016/j.icarus.2007.06.027>
- Kempf, S., Beckmann, U., & Schmidt, J. (2010). How the enceladus dust plume feeds saturn’s e ring. *Icarus*, *206*, 446–457. <https://doi.org/10.1016/j.icarus.2009.09.016>
- Lainey, V., Casajus, L. G., Fuller, J., & et al. (2020). Resonance locking in giant planets indicated by the rapid orbital expansion of titan. *Nature Astronomy*, *4*, 1053–1058. <https://doi.org/10.1038/s41550-020-1120-5>
- Lainey, V., Karatekin, Ö., Desmars, J., & et al. (2012). Strong tidal dissipation in saturn and constraints on enceladus’ thermal state from astrometry. *The Astrophysical Journal*, *752*(1), 14. <https://doi.org/10.1088/0004-637X/752/1/14>
- Lambrechts, M., & Johansen, A. (2012). Rapid growth of gas-giant cores by pebble accretion. *Astronomy & Astrophysics*, *544*, A32. <https://doi.org/10.1051/0004-6361/201219127>
- Lyra, W., Johansen, A., Klahr, H., & Piskunov, N. (2009). Standing on the shoulders of giants: Trojan earths and vortex trapping in low-mass self-gravitating protoplanetary disks of gas and solids. *Astronomy & Astrophysics*, *493*(3), 1125–1139. <https://doi.org/10.1051/0004-6361:200810797>

- Mastrapa, R., Bernstein, M., Sandford, S., Roush, T., Cruikshank, D., & Ore, C. D. (2008). Optical constants of amorphous and crystalline h₂o-ice in the near infrared from 1.1 to 2.6 μ m. *Icarus*, *197*(1), 307–320. <https://doi.org/https://doi.org/10.1016/j.icarus.2008.04.008>
- Morello, C. D., West, R. A., & Berg, M. J. (2025). Polarization analysis of enceladus’ surface. *Journal of Quantitative Spectroscopy and Radiative Transfer*, *333*, Article 109331. <https://doi.org/10.1016/j.jqsrt.2024.109331>
- Mosqueira, I., & Estrada, P. R. (2003a). Formation of the regular satellites of giant planets in an extended gaseous nebula i: Subnebula model and accretion of satellites. *Icarus*, *163*(1), 198–231. [https://doi.org/10.1016/S0019-1035\(03\)00076-9](https://doi.org/10.1016/S0019-1035(03)00076-9)
- Mosqueira, I., & Estrada, P. R. (2003b). Formation of the regular satellites of giant planets in an extended gaseous nebula ii: Satellite migration and survival. *Icarus*, *163*(1), 232–255. [https://doi.org/10.1016/S0019-1035\(03\)00077-0](https://doi.org/10.1016/S0019-1035(03)00077-0)
- Nesvorný, D., Vokrouhlický, D., & Morbidelli, A. (2007). Capture of irregular satellites during planetary encounters. *The Astronomical Journal*, *133*(5), 1962–1976. <https://doi.org/10.1086/512850>
- Neubauer, F. M. (1980). Nonlinear standing alfvén wave current system at io: Theory. *J. Geophys. Res.; (United States)*, *85*:A3. <https://doi.org/10.1029/JA085iA03p01171>
- Nicholson, P. D., French, R. G., Tollestrup, E., & et al. (1996). Observations of saturn’s ring-plane crossings in august and november 1995. *Science*, *272*, 509–515. <https://doi.org/10.1126/science.272.5261.509>
- Ormel, C. W., & Klahr, H. H. (2010). The effect of gas drag on the growth of protoplanets. analytical expressions for the accretion of small bodies in laminar disks. *Astronomy & Astrophysics*, *520*, A43. <https://doi.org/10.1051/0004-6361/201014903>
- Pollack, J. B., Burns, J. A., & Tauber, M. E. (1979). Gas drag in primordial circumplanetary envelopes: A mechanism for satellite capture. *Icarus*, *37*, 587–611. [https://doi.org/10.1016/0019-1035\(79\)90016-2](https://doi.org/10.1016/0019-1035(79)90016-2)
- Porco, C. C., Helfenstein, P., Thomas, P. C., Ingersoll, A. P., Wisdom, J., West, R., Neukum, G., Denk, T., Wagner, R., Roatsch, T., Kieffer, S., Turtle, E., McEwen, A., Johnson, T. V., Rathbun, J., Veverka, J., Wilson, D., Perry, J., Spitale, J., ... Squyres, S. (2006). Cassini observes the active south pole of enceladus. *Science*, *311*(5766), 1393–1401. <https://doi.org/10.1126/science.1123013>
- Porco, C. C., Helfenstein, P., Thomas, P. C., Ingersoll, A. P., Wisdom, J., West, R., & et al. (2018). Cassini observes the active south pole of enceladus. In P. M. Schenk, R. N. Clark, C. J. A. Howett, A. J. Verbiscer, & J. H. Waite (Eds.), *Enceladus and the icy moons of saturn* (pp. 1393–1401). University of Arizona Press. <https://doi.org/10.1126/science.1123013>
- Porco, C. C., West, R. A., Squyres, S. W., McEwen, A., Thomas, P., Murray, C. D., DelGenio, A. D., Ingersoll, A. P., Johnson, T. V., Neukum, G., et al. (2004). Cassini imaging science: Instrument characteristics and anticipated scientific investigations at saturn. *Space Science Reviews*, *115*(1–4), 363–497. <https://doi.org/10.1007/s11214-004-1456-7>
- Postberg, F., Khawaja, N., Abel, B., Choblet, G., Glein, C. R., Gudipati, M. S., Henderson, B. L., Hsu, H.-W., Kempf, S., Klenner, F., Moragas-Klostermeyer, G., Magee, B., Nölle, L., Perry, M., Reviol, R., Schmidt, J., Srama, R., Stolz, F., Tobie, G., ... Waite, J. H. (2018). Macromolecular organic compounds from the depths of enceladus. *Nature*, *558*(7711), 564–568. <https://doi.org/10.1038/s41586-018-0246-4>
- Pryor, W. R., Rymer, A. M., Mitchell, D. G., Hill, T. W., Young, D. T., Saur, J., Jones, G. H., Jacobsen, S., Cowley, S. W. H., Mauk, B. H., Coates, A. J., Gustin, J., Grodent, D., Gérard, J.-C., Lamy, L., Nichols, J. D., Krimigis, S. M., Esposito, L. W., Dougherty, M. K., ... Zhou, X. (2011). The auroral footprint of enceladus on saturn. *Nature*, *472*(7343), 331–333. <https://doi.org/10.1038/nature09928>
- Ronnet, T., & Johansen, A. (2020a). Formation of moon systems around giant planets. I. constraints from pebble accretion with inner cavities. *Astronomy & Astrophysics*, *633*, A93. <https://doi.org/10.1051/0004-6361/201936804>
- Ronnet, T., & Johansen, A. (2020b). Formation of moon systems around giant planets. II. capture and ablation of planetesimals as foundation for a pebble accretion scenario [Use this key only if you cite the part explicitly labelled 2020b in your notes]. *Astronomy & Astrophysics*, *633*, A93. <https://doi.org/10.1051/0004-6361/201936804>
- Rubrecht, N., Cazaux, S., Seignovert, B., Kenworthy, M., Kutsop, N., Le Mouélic, S., & Loicq, J. (2025). Peculiar rainbows in saturn’s E ring: Uncovering luminous bands near enceladus. *Icarus*, *441*, 116650. <https://doi.org/10.1016/j.icarus.2025.116650>
- Sasaki, T., Stewart, G. R., & Ida, S. (2010). Origin of the different architectures of the Jovian and Saturnian satellite systems. *The Astrophysical Journal*, *714*(2), 1052. <https://doi.org/10.1088/0004-637X/714/2/1052>

- Saur, J., Neubauer, F. M., & Schilling, N. (2007). Hemisphere coupling in enceladus' asymmetric plasma interaction. *Journal of Geophysical Research: Space Physics*, *112*(A11). <https://doi.org/10.1029/2007JA012479>
- Shibaike, Y., & Alibert, Y. (2020). Planetesimal formation at the gas pressure bump following a migrating planet. I. basic characteristics. *Astronomy & Astrophysics*, *644*, A81. <https://doi.org/10.1051/0004-6361/202039086>
- Shibaike, Y., Okuzumi, S., Sasaki, T., & Ida, S. (2017). Satellitesimal formation via collisional dust growth in steady circumplanetary disks. *The Astrophysical Journal*, *846*(1), 81. <https://doi.org/10.3847/1538-4357/aa8454>
- Showalter, M. R., Cuzzi, J. N., & Larson, S. M. (1991). Structure and particle properties of saturn's e ring. *Icarus*, *94*, 451–473. [https://doi.org/10.1016/0019-1035\(91\)90236-Y](https://doi.org/10.1016/0019-1035(91)90236-Y)
- Slipher, V. M. (1914). Mimas and Enceladus turn always the same face to Saturn. *Lowell Observatory Bulletin*, *2*, 70–72.
- Smith, B. A., Soderblom, L., Batson, R., Bridges, P., Inge, J., Masursky, H., Shoemaker, E., Beebe, R., Boyce, J., Briggs, G., Bunker, A., Collins, S. A., Hansen, C. J., Johnson, T. V., Mitchell, J. L., Terrile, R. J., Cook, A. F., Cuzzi, J., Pollack, J. B., . . . Suomi, V. E. (1982). A new look at the saturn system: The voyager 2 images. *Science (New York, N.Y.)*, *215*(4532), 504–537. <https://doi.org/10.1126/science.215.4532.504>
- Spahn, F., Schmidt, J., Albers, N., Hörning, M., Makuch, M., Seiss, M., Kempf, S., Srama, R., Dikarev, V., Helfert, S., Moragas-Klostermeyer, G., Krivov, A. V., Sremcevic, M., Tuzzolino, A. J., Economou, T., & Grün, E. (2006). Cassini dust measurements at enceladus and implications for the origin of the e ring. *Science*, *311*(5766), 1416–1418. <https://doi.org/10.1126/science.1121375>
- Spencer, J. R., et al. (2009). Enceladus: An active cryovolcanic satellite. In M. K. Dougherty, L. W. Esposito, & S. M. Krimigis (Eds.), *Saturn from cassini-huygens* (pp. 683–724). Springer.
- Spencer, J. R., Pearl, J. C., Segura, M., Flasar, F. M., Mamoutkine, A., Romani, P., & et al. (2006). Cassini encounters enceladus: Background and the discovery of a south polar hot spot. In *Science* (pp. 1401–1405, Vol. 311). <https://doi.org/10.1126/science.1121661>
- Stryk, T. (2017). Enceladus' plumes, seen by voyager 1 in 1980 [Accessed: 2026-04-25]. <https://www.planetary.org/articles/20170221-enceladus-plumes-voyager1>
- Szulágyi, J., Binkert, F., & Surville, C. (2022). Meridional circulation of dust and gas in the circumstellar disk: Delivery of solids onto the circumplanetary region. *The Astrophysical Journal*, *924*(1), 1. <https://doi.org/10.3847/1538-4357/ac32d1>
- Takasao, S., Tomida, K., Iwasaki, K., & et al. (2022). Three-dimensional simulations of magnetospheric accretion in a T tauri star: Accretion and wind structures just around the star. *The Astrophysical Journal*, *941*(1), 73. <https://doi.org/10.3847/1538-4357/ac9eb1>
- Thomas, P., Tajeddine, R., Tiscareno, M., Burns, J., Joseph, J., Lored, T., Helfenstein, P., & Porco, C. (2016). Enceladus's measured physical libration requires a global subsurface ocean. *Icarus*, *264*, 37–47. <https://doi.org/10.1016/j.icarus.2015.08.037>
- Waite, J., Combi, M., Ip, W.-H., Cravens, T., McNutt, R., Kasprzak, W., Yelle, R., Luhmann, J., Niemann, H., Gell, D., Magee, B., Fletcher, G., Lunine, J., & Tseng, W. (2006). Cassini ion and neutral mass spectrometer: Enceladus plume composition and structure. *Science (New York, N.Y.)*, *311*, 1419–22. <https://doi.org/10.1126/science.1121290>
- Waite, J. H., Glein, C. R., Perryman, R. S., Teolis, B. D., Magee, B. A., Miller, G., Grimes, J., Perry, M. E., Miller, K. E., Bouquet, A., Lunine, J. I., Brockwell, T., & Bolton, S. J. (2017). Cassini finds molecular hydrogen in the enceladus plume: Evidence for hydrothermal processes. *Science*, *356*(6334), 155–159. <https://doi.org/10.1126/science.aai8703>
- Warren, S. G., & Brandt, R. E. (2008). Optical constants of ice from the ultraviolet to the microwave: A revised compilation. *Journal of Geophysical Research*, *113*(D14), D14220. <https://doi.org/10.1029/2007JD009744>
- Yeoh, S. K., Li, Z., Goldstein, D. B., Varghese, P. L., Levin, D. A., & Trafton, L. M. (2017). Constraining enceladus' plume using numerical simulation based on cassini inms data. *Icarus*, *281*(1), 357–378. <https://doi.org/10.1016/j.icarus.2016.08.028>
- Zhu, Z., Andrews, S. M., & Isella, A. (2018). On the radio detectability of circumplanetary discs. *Monthly Notices of the Royal Astronomical Society*, *479*(2), 1850–1865. <https://doi.org/10.1093/mnras/sty1503>

Part II
Scientific Article

Automated Detection and Characterisation of Luminous Bands in Saturn’s E Ring from Cassini ISS Data

Arnaud Mathieu^{a,1}, Stéphanie Cazeaux^a and Jérôme Loicq^a

^a *TU Delft, Delft, Netherlands*

ARTICLE INFO

Keywords:

Saturn
E ring
Enceladus
Cassini ISS
Fourier transform
VIMS
Ring structure
Image processing

ABSTRACT

Saturn’s E ring harbours faint, quasi-periodic, inclined brightness structures called luminous bands, first characterised from Cassini VIMS and ISS observations during Enceladus flybys. We present the first automated, catalogue-scale survey of luminous bands, drawing on 23 Cassini flybys. Detection exploits the chromatic character of the bands, whose apparent position shifts with wavelength. Because these bands present a quasi-periodic inclined brightness structure, a frequency-domain search using a zero-padded two-dimensional Fourier angular spectrum is applied, with candidates assessed against a phase-randomised null ensemble. Of 573 processed images, 62 yield positive detections across eight flybys. Six of the eight carry no prior published detection, including the first ISS detection for E13, and five are non-targeted flybys, demonstrating bands are present in more ISS images than initially found. The observed band directions are compared against a diffraction model: 39 of the 62 detections agree to within 2 degrees, confirming the grating interpretation; the remaining 23 show a systematic offset whose origin is not yet explained. Band contrasts decrease with Cassini–Enceladus distance, suggesting the strongest structure near the plume source. Notably, no detections occur in the morning ansa despite comparable coverage. A complementary VIMS spectral analysis reveals selective suppression of sub-micron grains on the morning-side ring, consistent with size-selective electromagnetic clearing by Lorentz forces, explaining the asymmetric detection rate.

1. Introduction

Saturn’s E ring is a diffuse, extended torus of micron-sized ice grains, sustained by the cryovolcanic plumes of Enceladus (Porco et al., 2006) and shaped by a continuous interplay between plume injection, plasma drag, and electromagnetic forcing (Spahn et al., 2006; Postberg et al., 2011; Hamilton, 1993). Its optical depth is orders of magnitude below that of the main rings; any brightness structure within it reflects the local balance between particle injection near Enceladus and dispersed loss processes acting across the full torus (Hedman et al., 2012). Within this environment, Rubbrecht et al. (2025) discovered a previously unknown category of fine structure: faint, quasi-periodic, inclined brightness features, termed luminous bands, detected in Cassini Visual and Infrared Mapping Spectrometer (VIMS) data from the E13, E17, and E19 Enceladus flybys, and in Imaging Science Subsystem (ISS) data from E17. Their simultaneous detection by both instruments during E17, using two completely independent optical systems, confirms that the bands are a real external physical phenomenon rather than an instrumental artefact. Their quasi-periodic character and wavelength-dependent positional shift are consistent with diffraction by a millimetre-scale periodic structure in the ring plane, though the formation mechanism of such a structure has not yet been established.

The central question left open by that discovery is whether luminous bands are confined to the specific flybys in which they were first detected, or are a recurring feature of high-phase E ring imaging. Rubbrecht et al. (2025) identified the bands under high phase angles and near-ring-plane viewing geometry, the conditions under which forward-scattering by E ring grains dominates the scene. This observational regime motivates the search strategy of the present survey; whether bands are a common feature of high-phase E ring imaging, or are instead linked to only a few specific flybys, remains an open question. Importantly, the bands are chromatic: their position shifts with wavelength, which is central to the diffraction interpretation and means detection can best be done by comparing images at different wavelengths. The ISS NAC filter set, which spans the full visible range and includes polarised-light filters not used by Rubbrecht et al. (2025), is therefore well suited for extending the survey.

The survey yields 62 positive detections across eight flybys, extending ISS coverage well beyond the single flyby (E17) reported by Rubbrecht et al. (2025). Band contrasts decrease with Cassini–Enceladus distance, and no detections fall in the morning ansa despite comparable coverage there. A complementary VIMS spectral analysis points to a selective depletion of sub-micron grains on the morning side, consistent with shadow-induced Lorentz forcing (Hedman et al., 2012), and provides a physical basis for the asymmetric detection rate.

ORCID(s): <https://orcid.org/0000-0003-0320-3578> (S. Cazeaux); <https://orcid.org/0000-0003-1863-6085> (J. Loicq)
¹amathieu@tudelft.nl (A. Mathieu)

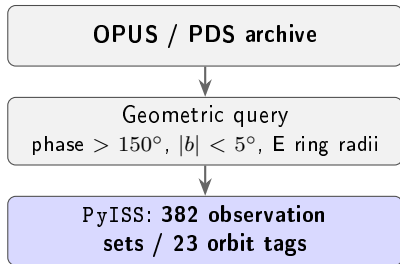


Figure 1: Archive query and observation-set reconstruction chain (b : observer’s ring-plane elevation angle). OPUS: Online Planetary Unified Search, the PDS (Planetary Data System) Ring-Moon Systems Node archive query service. PyISS: the Python package developed for this work (Appendix A). The highlighted box marks the 382 observation sets that form the input to the image reduction stage (Section 3).

This paper presents the first automated, catalogue-scale ISS survey of luminous bands. The paper is organised as follows. Section 2 presents the observational dataset and the steps by which it is constructed. Section 3 details the image reduction and residual construction. Section 4 details the frequency-domain detection pipeline and its significance assessment. Section 5 presents the detection catalogue, a diffraction-model orientation analysis for all 62 detections, and photometric contrast characterisation. Section 6 situates the findings with respect to Rubbrecht et al. (2025) and discusses the physical implications of the asymmetric detection rate.

2. Survey Observations

The observational strategy of this survey is grounded in the viewing geometry identified by Rubbrecht et al. (2025) as the detection regime for luminous bands, which serves as the working hypothesis and starting point of the present study. In their characterisation, bands are consistently detected under two conditions: high phase angles, where E ring grains scatter light strongly toward the observer, making the ring bright; and near-ring-plane viewing, where a long line of sight through the ring intercepts more scattering particles, maximising sensitivity. Adopting these conditions as search criteria, the full Cassini ISS NAC archive is queried through the PDS Ring-Moon Systems Node’s OPUS search service, retaining only flybys where bands could plausibly be detected given this hypothesis. Whether bands also appear outside this regime is an open question; in principle the same pipeline can be rerun with broader search constraints (Table 1) to probe other geometries systematically. The query and observation-set reconstruction chain is summarised in Figure 1.

Conducting such a search by eye is not practical at archive scale: the candidate set is large, manual inspection introduces selection bias, and band contrasts

are often near the noise level. An automated approach is therefore adopted. While no suitable polarised NAC observations were found in the archive, the automated pipeline yields a detection catalogue extending well beyond the single flyby in which Rubbrecht et al. (2025) reported ISS detections.

The Cassini Imaging Science Subsystem Narrow Angle Camera (ISS NAC; Porco et al. 2004) has an angular resolution of $\sim 6 \mu\text{rad px}^{-1}$ and a spectral range of 200–1100 nm. All frames are calibrated using the CISSCAL 4.0 pipeline (Knowles et al., 2020), which converts raw instrument counts to I/F , the ratio of reflected intensity to incident solar flux and a standardised measure of reflectance. Within each orbit tag, sequential NAC frames acquired at the same pointing through different filters constitute an observation set: one complete filter sequence, typically a broadband CLEAR-filter exposure followed by one or more narrow-band science-filter exposures. Each filter selects a different wavelength band, so the images in a set show the same scene at different wavelengths, taken within a few minutes of each other. Figure 2 illustrates this for the running example of this paper. Because the archive metadata lacks explicit sequence identifiers, recovering and grouping these sets is handled automatically by PyISS, a Python package developed for this work and described in Appendix A.

The filtered archive spans 23 Cassini orbit tags from February 2007 to June 2012. Each orbit tag corresponds to one flyby; within it, PyISS identifies all observation sets that contain at least one CLEAR exposure and at least one science-filter exposure, since these are the sets from which CLEAR-subtracted residuals can be produced. A total of 382 such sets are recovered across the 23 orbit tags. These orbit tags fall into two categories. The first are targeted Enceladus flybys: encounters in which Cassini’s trajectory was specifically manoeuvred for a close approach to Enceladus, designated by sequential E-numbers across the mission. This category includes E13, E17, and E19, the encounters studied by Rubbrecht et al. (2025), as well as additional targeted flybys not previously searched for luminous bands. The second are non-targeted flybys: orbits in which Enceladus fell within the NAC field at greater distance due to Cassini’s regular orbital geometry, without a dedicated trajectory correction manoeuvre, and therefore carrying no E-number designation. Both categories are represented in the survey. In all cases the OPUS query ensures Enceladus was within the NAC field of view at the time of observation.

3. Image Reduction and Residual Construction

From the 382 observation sets identified in Section 2, usable residual images are constructed through

Table 1

Geometric constraints applied to the OPUS query. The phase-angle and ring-elevation cuts together isolate the forward-scattering, near-ring-plane regime in which the E ring continuum dominates the scene. Values are informed by the detection conditions reported in Rubbrecht et al. (2025) and are not intended as proven exclusive constraints.

Parameter	Min	Max	Physical rationale
Target		Enceladus	Centres NAC field on E ring
Instrument		ISS NAC	High-resolution imaging
Phase angle (α)	150°	180°	Forward-scattering regime
Ring radius	1.8×10^5 km	4.8×10^5 km	Radial extent of E ring
Observer ring elev. (b)	-5°	$+5^\circ$	Long line of sight through ring
Solar opening angle	-15°	$+15^\circ$	Low solar incidence on ring plane




Figure 2: Observation set from orbit 163EN (E17, 27 March 2012). Five sequential NAC exposures acquired at the same pointing within 328 s, each through a different spectral filter: (a) broadband CLEAR (reference), (b) IR3 (~ 930 nm), (c) IR1 (~ 752 nm), (d) RED (~ 617 nm), and (e) BL1 (~ 451 nm). All five frames show the same scene – Enceladus and the surrounding E ring – at different wavelengths. The time offset Δt from the CLEAR exposure is indicated for each panel.




Figure 3: Image reduction chain from observation sets to the formal survey input. The highlighted box marks the 573 residual images that enter the detection pipeline.

three stages summarised in Figure 3: CLEAR-filter subtraction, Enceladus disc masking, and a final quality-control review. The motivation for this procedure was introduced in the introduction: luminous bands are chromatic features, established as such by Rubbrecht et al. (2025), whose brightness and apparent spacing vary with wavelength, while the dominant scene background is spectrally neutral forward-scattered E ring light. Subtracting a broadband CLEAR image from a science-filter frame removes this achromatic continuum

and leaves the wavelength-dependent brightness residuals in which luminous bands appear.

3.1. CLEAR-filter subtraction

For each observation set, a residual is constructed by pairing each available science-filter frame with the nearest CLEAR exposure in time. The CLEAR filter is the ISS broadband channel, spanning approximately 350–1050 nm; it records the full forward-scattered E ring continuum. A science-filter frame records the same continuum modulated by any wavelength-dependent brightness structure in the scene. The residual

$$I_{\text{res}} = I_X - k I_{\text{CLEAR}} \quad (1)$$

subtracts the scaled CLEAR frame from the science-filter frame, where I_X is the I/F image in a given science filter (the ISS NAC offers a range of narrow-band spectral filters; Figure 2 shows an example set of four) and k is a passband scaling factor correcting for the throughput difference between the two filter paths. The factor k is computed from the ratio of passband integrals weighted by the solar spectral irradiance (accounting for the Sun's colour in the normalisation) for each (CLEAR, science-filter) combination, evaluated from the tabulated ISS transmission curves (Porco et al., 2004). These curves are fixed by the instrument calibration, so k has a unique predetermined value for each filter pair and carries no free parameters. Selecting the nearest CLEAR exposure minimises residuals from




Figure 4: False-colour CLEAR-subtracted residual for the pair (N1711553312 / N1711553432; IR3 filter, orbit 163EN, 27 March 2012), rendered with a diverging colour map centred at zero. Red denotes positive residuals (science brighter than scaled CLEAR) and blue denotes negative deviations. The residual is well balanced; faint inclined ring structures are visible across the field, and the bright red region at the lower limb traces emission from the Enceladus south-polar plume.

pointing drift and short-timescale flux variations between the two frames. A per-image three-sigma outlier filter then suppresses single-pixel artefacts, replacing flagged pixels with the residual median.

The resulting image has values close to zero in spectrally neutral ring regions and deviates wherever the scene has a chromatic brightness structure. Figure 4 shows a representative residual in which the subtraction is well-behaved: the field is balanced around zero and faint inclined ring structures are visible against the flat background.

3.2. Enceladus disc masking

The principal contaminating signal in the residual is the partially subtracted Enceladus disc. Even with closely time-matched exposures, small differences in pointing, timing, and scattered-light level leave a disc-shaped halo in the residual that would otherwise overwhelm the faint ring structure and must therefore be masked.

The disc centre and apparent radius in the NAC focal plane are computed for each exposure using SPICE kernels (Acton, 1996; Acton et al., 2018), a standard toolkit that reconstructs the geometric state of the Cassini spacecraft and solar system bodies at any point in time from navigation and telemetry data. A circular mask of radius $1.5 R_{\text{Enc}}$ is placed around the computed disc centre, with its boundary softened




Figure 5: Enceladus disc mask for image N1711553432 (IR3, orbit 163EN). The circular mask boundary at $1.5 R_{\text{Enc}}$ (white dashed) surrounds the SPICE-computed disc centre; the cosine taper region (yellow shading) smoothly attenuates the transition to avoid ringing in the Fourier transform.

by a cosine taper over 60 pixels to avoid sharp edges that would otherwise introduce ringing artefacts in the subsequent analysis. A cosine taper is preferred over a linear one because its continuous first and second derivatives at both endpoints produce a smoother roll-off, further suppressing spectral leakage in the Fourier transform. For ten images in which Saturn and its rings enter the NAC field, an additional fixed rectangular mask is applied manually. Figure 5 shows the mask applied to the running example of this paper.

3.3. Quality control and formal survey definition

Applying the subtraction to the 382 observation sets yields 624 residual images. Before any detection is attempted, all 624 are reviewed visually to exclude frames unsuitable for analysis regardless of band presence. Three categories prompt exclusion: frames in which the Enceladus disc fills or dominates the unmasked field; frames with large time separations between the CLEAR and science exposures, producing incoherent residuals; and frames with sensor or readout artefacts surviving the outlier filter. Fifty-one frames are excluded on these grounds, leaving the formal survey of 573 residual images summarised in Table 2. These constitute the input to the detection pipeline described in Section 4.

4. Fourier Angular Spectrum Detection

The 573 CLEAR-subtracted residual images produced in Section 3 make the chromatic brightness

Table 2

Survey catalogue summary. The reduction chain proceeds from observation sets through CLEAR subtraction to the 573 residual images forming the formal survey input.

Quantity	Value
Orbit tags	23 (2007–2012)
Observation sets	382
Residual images (pre-QC)	624
Excluded by quality control	51
Formal survey	573
Science filters	IR3, IR1, RED, BL1 (most frequent); additional filters present
Phase-angle range	150°–180°

structure of the E ring directly accessible. Luminous bands, when present, appear as faint inclined periodic stripes across the residual field. Their contrast is often near the noise level, and selecting detections by visual inspection alone would introduce observer bias: images with obvious bands are noticed while faint ones are overlooked, undermining the statistical completeness the survey requires. The key to systematic detection lies in their periodic character. A periodic brightness stripe at inclination angle θ_{band} and spatial frequency f produces a concentrated pair of peaks in the two-dimensional Fourier power spectrum at polar coordinates $(f, \theta_{\text{band}} + 90^\circ)$, regardless of the band's absolute brightness or the smooth background in which it is embedded. This transformation turns a low-contrast stripe detection problem into a peak-localisation problem: the periodic signal accumulates into a single spectral feature while incoherent background noise distributes uniformly across all frequencies. Direct image-space methods such as profile fitting or template matching do not offer this accumulation property and are sensitive to the choice of initial parameters. ISS NAC images are square arrays, directly compatible with this two-dimensional transform; the spatial data cube format of VIMS does not share this property.

Figure 6 illustrates this principle with a synthetic image constructed for this purpose. A sinusoidal brightness stripe at 135° , the simplest possible model of a luminous band in a CLEAR-subtracted residual, is the sole content of the artificially generated image shown in panel (a). Its angular power spectrum $E(\theta)$, computed from the mean-subtracted image without any additional processing, concentrates entirely into a single sharp spike at $\theta^* = 45^\circ$ (panel (b)). Since the 2D Fourier power concentrates perpendicular to the band, $\theta^* = 45^\circ$ recovers the band orientation as $\theta_{\text{band}}^* = \theta^* + 90^\circ = 135^\circ$. The rest of the spectrum is at zero. This is the signature the detection pipeline is designed to recover: a narrow peak standing above a flat baseline. The angle of 135° is representative; any other band orientation produces the same concentration of power at the perpendicular direction.

In practice, each residual image is first conditioned by a preprocessing chain that suppresses the three dominant contamination sources (Enceladus disc signal, large-scale brightness gradient, and detector readout striping) which would otherwise dominate the Fourier power (Section 4.1). The conditioned image is then transformed to the Fourier domain and its power is projected onto a one-dimensional angular spectrum $E(\theta)$; zero-padding the image before the transform (described in Section 4.2) sharpens the angular resolution enough to localise band orientations to sub-degree precision. The significance of each candidate peak is assessed against a phase-randomised null ensemble (a set of surrogate images sharing the same statistical properties but containing no coherent spatial structure), yielding a null-corrected excess spectrum $E_{\text{diff}}(\theta)$ from which the band orientation is extracted (Section 4.2). The procedure concludes with a hybrid detection verdict that combines the automated spectral result with a per-image visual inspection against the spectral diagnostic (Section 4.3).

4.1. Noise management and preprocessing

Unlike the idealised case of Figure 6, real CLEAR-subtracted residuals contain signal sources in addition to any luminous band. Figure 4 shows a representative example: the Enceladus disc dominates the lower portion of the image, a smooth brightness gradient runs across the ring field, and faint row-aligned banding is visible from detector readout. Each of these distributes Fourier power in a way that would either bury a genuine band peak or generate a false one at a specific orientation. Figure 5 already showed how the disc is handled geometrically; the remaining contamination sources are suppressed by a fixed four-step preprocessing chain applied to every residual image before the Fourier transform.

(i) *High-pass filtering.* A Gaussian-smoothed copy of the residual ($\sigma = 80$ px, an empirically determined value) is subtracted, cutting the very lowest spatial frequencies that span the full image width. Luminous bands repeat over tens to hundreds of pixels and are not affected by this cut; the smooth large-scale brightness




Figure 6: Synthetic demonstration of the Fourier detection principle. *Panel (a):* a pure sinusoidal brightness stripe at 135° , the geometric idealisation of a luminous band in a CLEAR-subtracted residual. *Panel (b):* the corresponding one-dimensional angular power spectrum $E(\theta)$ computed from the mean-subtracted image. The angular spectrum shows a single sharp spike at $\theta^* = 45^\circ$ (red dashed), perpendicular to the band, with no power at any other angle; the band orientation is recovered as $\theta_{\text{band}}^* = \theta^* + 90^\circ = 135^\circ$. This concentration of a periodic spatial structure into a single spectral feature is the core property that makes the Fourier angular spectrum the natural space for systematic luminous band detection.

trends that would otherwise dominate the angular spectrum are removed.

(ii) *Disc masking.* The Enceladus disc mask (Section 3.2) is applied: masked pixels are set to zero and the cosine taper smoothly attenuates the transition region. The partially subtracted disc halo would otherwise produce a broad isotropic pedestal in the angular spectrum, elevating the noise floor at all orientations equally.

(iii) *Row and column detrending.* The CLEAR subtraction described in Section 3 already suppresses most of the detector read-out striping, since the readout pattern is common to both frames in a pair. When the two

exposures differ slightly in readout state, however, a residual stripe along rows or columns remains. Per-row and per-column medians, computed over the unmasked region, are subtracted to remove it. As a conservative safeguard, the two detector-axis orientations (0° and 90°) are also excluded from the angular spectrum in the analysis steps that follow.

(iv) *Arcsinh compression.* Occasional bright outlier pixels that survive the outlier filter of Section 3.1 would otherwise dominate the Fourier amplitude. The arcsinh transformation $\text{arcsinh}(x/s)$, with image-adaptive scale $s = 0.05 \sigma_{\text{img}}$ (where σ_{img} is the standard deviation of the preprocessed pixel values, so s is set to 5% of the image noise level), compresses these bright pixels while preserving the relative brightness of faint band structure. The CLEAR subtraction already limits outlier severity by removing the Enceladus disc signal; the arcsinh step handles what remains.

4.2. Angular spectrum construction and calibration

The two-dimensional Fast Fourier Transform (FFT) converts the preprocessed residual into the frequency domain, where each point encodes a spatial frequency and orientation. Summing the 2D FFT magnitude over all spatial frequencies R at each orientation $\theta \in [0^\circ, 180^\circ]$ (where 0° and 180° represent the same undirected orientation) yields the one-dimensional angular power spectrum $E(\theta)$: a peak at angle θ signals periodic brightness structure at that orientation in the image. The summation uses 720 bins of width 0.25° ; the DC centre, regions within $\pm 5^\circ$ of the 0° and 90° detector-axis orientations, and high-frequency bins are excluded. A $1/R$ radial weight is applied to normalise for the increasing number of FFT grid points at larger radii, so that each spatial-frequency increment contributes equally to $E(\theta)$ regardless of its absolute frequency. Figure 7(b) shows the 2D FFT log-magnitude for the running example, and panel (c) shows the resulting $E(\theta)$.

The spectrum in panel (c) illustrates two problems that make the raw $E(\theta)$ unsuitable for detection. First, the image corners extend further along the 45° and 135° diagonals than along the detector axes, concentrating Fourier power at those directions in any square image regardless of its content. This systematic diagonal bias means the 135° peak in Figure 7(c) recurs across the catalogue even in images with no detectable band. Second, a band whose spatial frequency falls between two discrete grid points has its power spread across neighbouring bins, making the peak broad and its location unreliable by several degrees. Zero-padding resolves the second problem, and the null ensemble resolves the first.

Zero-padding surrounds the image with zeros before the FFT, growing the array from $N \times N$ to $(PN) \times (PN)$. The added zeros contain no new information;




Figure 7: Preprocessing and first-pass Fourier analysis for the running example pair (N1711553312/N1711553432; IR3 filter, orbit 163EN, 27 March 2012). *Panel (a)*: preprocessed residual after high-pass filtering, disc masking, row/column detrending, and arcsinh compression. The disc boundary (white dashed) and mask edge at $1.5 R_{\text{Enc}}$ (yellow solid) are visible at the lower edge; a faint inclined band runs from lower-left to upper-right. *Panel (b)*: log-power of the 2D FFT magnitude; the bright cross along the detector axes is excluded by the $\pm 5^\circ$ wedge mask, and a faint diagonal concentration reflects the band's directional signal. *Panel (c)*: one-dimensional angular power spectrum $E(\theta)$; the highest peak sits at $\theta^* = 135^\circ$ (red dashed), but at this frequency resolution the estimate is too coarsely localised to be reliable. Zero-padding (Section 4.2 and Figure 8) is required to sharpen this peak.

they force the FFT to evaluate the same continuous frequency spectrum at P times more grid points, interpolating between the original coarse bins. With $P = 8$, the grid spacing is reduced eightfold, turning the broad blob of Figure 7(c) into the well-defined ridge in Figure 8(a). Applying the same integration to the zero-padded magnitude $|G_{\text{zp}}(u, v)|$ yields $E_{\text{zp}}(\theta)$, shown in Figure 8(b).

Even a sharp peak in $E_{\text{zp}}(\theta)$ can be misleading. The diagonal bias survives zero-padding, and any image-specific amplitude distribution can produce angular anisotropy that mimics a band signal; a global significance threshold cannot account for these because they differ from image to image. For each image, a null ensemble is therefore constructed by randomising the phase of each Fourier coefficient while keeping its amplitude unchanged. The N_{null} surrogate images produced this way preserve the power at each individual spatial frequency, the radial profile, and the diagonal bias of the original, but destroy any coherent spatial structure: they are statistically indistinguishable from the original in all non-directional respects. Their median,

$$E_{\text{null,zp}}(\theta) = \text{median}_k E_{\text{zp}}^{(k)}(\theta), \quad (2)$$

defines the image-specific noise baseline, and subtracting it yields the excess spectrum

$$E_{\text{diff}}(\theta) = E_{\text{zp}}(\theta) - E_{\text{null,zp}}(\theta), \quad (3)$$

which isolates power that cannot be explained by the image's own amplitude distribution or geometry. The

Pearson correlation between E_{zp} and $E_{\text{null,zp}}$ has a mean of 0.91 and a median of 0.93 across the catalogue, confirming that the null accurately tracks the spectral envelope. Figure 8(b) shows E_{zp} alongside the null-ensemble median and spread; both lie near zero at this scale, as the dominant spike in E_{zp} accounts for nearly all the spectral power.

The null-corrected excess spectrum $E_{\text{diff}}(\theta)$ is the final detection statistic. Luminous bands have finite spatial extent in the image and excite a range of spatial frequencies, producing a broad feature in E_{diff} rather than an isolated spike. A single-bin maximum would therefore be vulnerable to isolated noise spikes that momentarily exceed the surrounding level. A moving-window average is applied to E_{diff} before peak extraction, smoothing away such isolated spikes while preserving the broader genuine band feature. The peak angle θ^* and full width at half maximum (FWHM) are then obtained by linear interpolation of the half-maximum crossings, yielding sub-bin precision independent of the bin width. The band orientation in the image plane is recovered as $\theta_{\text{band}}^* = (\theta^* + 90^\circ) \bmod 180^\circ$. Figure 8(c) shows the full result for the running example: a single clean peak at $\theta^* = 130^\circ$, corresponding to a band orientation of 40° .

4.3. Detection verdict

The automated spectral analysis returns a candidate peak and its diagnostics for every image, regardless of whether a physical luminous band is present. A second stage converts this into a binary detection




Figure 8: Angular spectrum construction and calibration for the running example pair (N1711553312 / N1711553432; IR3 filter, orbit 163EN, 27 March 2012). *Panel (a):* log-magnitude of the 2D zero-padded FFT $\log |G_{zp}(u, v)|$, cropped to the accepted radial annulus ($R \leq 6$ native cycles image $^{-1}$). The $8\times$ zero-padded transform resolves the diffuse directional concentration of Figure 7(b) into a sharply defined ridge (cyan). *Panel (b):* angular power spectrum $E_{zp}(\theta)$ (blue) with the null-ensemble median (dashed) and 16th/84th percentile spread (grey band) from 100 phase-randomised surrogates overplotted; both lie near zero at this scale because the dominant spike in E_{zp} accounts for nearly all the spectral power. *Panel (c):* null-corrected difference spectrum $E_{diff}(\theta)$ after moving-window smoothing. A single clean peak at $\theta^* = 130^\circ$ (red dashed), corresponding to a band orientation of 40° (orange dotted), emerges above a flat baseline; the FWHM is marked by the green shading.

verdict through a structured visual review. Each image is presented alongside the preprocessed residual, the zero-padded Fourier magnitude, and E_{diff} with the peak overlaid. An image is labelled *positive* when a coherent inclined band is visible in the residual and its orientation is consistent with the spectral peak angle θ^* ; *doubtful* when the spatial and spectral evidence are inconsistent; and unlabelled otherwise.

This two-stage design is deliberate. Fully automated thresholds alone produce too many false positives, since partially subtracted disc halos and cosmic-ray residuals can generate spurious peaks in E_{diff} , while a purely visual workflow removes the reproducibility that automation provides. The spectral step narrows the candidate set; the visual step confirms or rejects each one.

5. Detection Statistics and Band Properties

The pipeline was applied to all 573 image pairs of the formal survey. Following visual inspection of the standardised multi-panel diagnostic, 62 pairs yield positive detections (10.8%) and 26 are classified as doubtful; the remainder are unlabelled. Positive detections are concentrated in the IR3, IR1, RED, and BL1 filters, which together account for 84% of all detections. The 62 positive detections are distributed across seven flybys as listed in Table 3; six of the eight had no prior published luminous-band search; 163EN (E17) was studied by Rubbrecht et al. (2025) in both VIMS and ISS, and 142EN (E13) was studied by Rubbrecht et al. (2025) in VIMS only, with the present survey providing the first ISS detections for that encounter.

Table 3

Positive detections by Cassini flyby. Six of the eight flybys have no prior published luminous-band detection.

Flyby	Detections	Notes
163EN	8	Targeted (E17); Rubbrecht et al. (2025)
142EN	5	Targeted (E13); Rubbrecht et al. (2025) (VIMS only)
131EN	12	Targeted (E10); no prior luminous-band search
106EN	15	Non-targeted; no prior luminous-band search
161EN	8	Non-targeted; no prior luminous-band search
098EN	8	Non-targeted; no prior luminous-band search
144EN	4	Non-targeted; no prior luminous-band search
122EN	2	Non-targeted; no prior luminous-band search
Total	62	

Figures 9 and 10 illustrate the detection diagnostic for a representative positive and null case. In the positive example (image N1616349628, IR1, orbit 106EN), parallel inclined bands are visible in the preprocessed residual, the 2D FFT magnitude shows a clear directional concentration, and $E_{diff}(\theta)$ yields a clean unimodal peak at $\theta^* \approx 61^\circ$, corresponding to a band direction $\theta_{band}^* \approx 151^\circ$. In the null case (image N1643188892, IR1, orbit 122EN), no coherent inclined feature is visible in the residual and $E_{diff}(\theta)$ is flat across all orientations.

Luminous bands in Saturn's E ring




Figure 9: Detection diagnostic for image N1616349628 (IR1, orbit 106EN). (a) Preprocessed CLEAR-subtracted residual; the cyan 7×7 mosaic marks the detected band direction ($\theta_{\text{band}}^* \approx 151^\circ$). (b) Zero-padded 2D FFT log-magnitude $\log |G_{\text{ZP}}(u, v)|$; the cyan line marks the spectral peak axis. (c) Smoothed excess angular power $E_{\text{diff}}(\theta)$; the green shading marks the FWHM of the peak at $\theta^* \approx 61^\circ$.




Figure 10: Detection diagnostic for N1643188892 (IR1, orbit 122EN): a representative non-detection. (a) Preprocessed residual shows no coherent inclined feature; the 7×7 mosaic at the best-guess direction confirms the absence of organised bands. (b) 2D ZP FFT shows no strong directional concentration. (c) $E_{\text{diff}}(\theta)$ is flat, with no peak above the null-ensemble baseline.

5.1. Band orientation angles

The spectral peaks are well-localised, with a mean FWHM of 26° (range: 16° – 49° across all 62 detections), confirming clear directional coherence in each individual detection. However, the raw 2D FFT angle is expressed in the image plane, which is tied to the camera's specific orientation at the moment of each exposure and is therefore not directly comparable across different encounters. To enable a cross-encounter comparison, each detection angle θ^* is placed into a universal, instrument-independent reference frame: it is rotated from the 2D image plane into the J2000 inertial frame using Cassini's attitude history (C-kernels from the SPICE toolkit) and then projected onto Saturn's

equatorial ring plane. The result is the observed band direction \vec{b}_{obs} , expressed as an azimuthal angle in the ring plane and independent of instrument pointing geometry.

The diffraction interpretation proposed by Rubbrecht et al. (2025) provides a physical baseline against which these observed directions can be tested. Under the diffraction hypothesis, the grating vector must lie in the scattering plane (the plane spanned by the directions from the illuminated ring point toward the Sun and toward Cassini), and the bands, running perpendicular to the grating vector, must lie along the intersection of that scattering plane with the ring plane. This yields a predicted band direction \vec{b}_{pred} , computed




Figure 11: Geometric construction for the diffraction model orientation prediction. The E ring plane (grey disc) and the scattering plane (blue, defined by \vec{k}_{Sun} and \vec{k}_{Cassini}) intersect along the grating vector \vec{v}_{grating} (red). The predicted band direction \vec{b}_{pred} (green) is perpendicular to the grating vector within the ring plane. The right-angle marker in the ring plane shows the 90° constraint between \vec{v}_{grating} and \vec{b}_{pred} . *Note:* The observer elevation and phase angles in this schematic are exaggerated for visual clarity. The actual ISS observations occur at near-grazing elevations ($|b| < 5^\circ$) and high forward-scattering phase angles ($> 150^\circ$), where the scattering plane is nearly parallel to the E ring plane.

for each image purely from Sun and Cassini positions retrieved from SPICE ephemeris kernels: the two unit vectors define the scattering plane, whose intersection with Saturn's equatorial plane gives the grating vector, and rotating that vector by 90° within the ring plane gives \vec{b}_{pred} . Because the prediction follows from celestial geometry alone, it carries no free parameters and is independent of grain size, grain shape, or grating formation mechanism.

The deviation angle $\Delta\theta$ between \vec{b}_{obs} and \vec{b}_{pred} then quantifies agreement with the diffraction hypothesis; $\Delta\theta = 0^\circ$ would correspond to an exact match. Figure 11 illustrates the geometric construction underlying \vec{b}_{pred} .

Figure 12 shows the distribution of $\Delta\theta$ for all 62 positive detections. The distribution is clearly split into two groups. For 39 of the 62 detections, spanning the 131EN, 142EN, 144EN, 122EN, and 163EN flybys, the deviation angle is close to zero, with a median of only 1.6° (just six angular-spectrum bins at the 0.25° resolution of the zero-padded transform, well within the expected measurement uncertainty). This close agreement between \vec{b}_{obs} and \vec{b}_{pred} is consistent with the diffraction interpretation of Rubbrecht et al. (2025), and extends the supporting evidence to a sample six times larger than previously available and across a much greater diversity of viewing geometries.

The remaining 23 detections, from the 106EN and 098EN flybys, show a consistent systematic offset of $\Delta\theta \approx 11\text{--}13^\circ$. Because this offset is identical across different filters and independent flybys, it is not a measurement error. Instead, it indicates that while the underlying diffraction physics remain sound, an additional, currently unknown geometric or physical factor is shifting the band alignment under those specific orbital conditions. Identifying this factor is an open question for future work.

Establishing that these structures behave as diffraction gratings physically explains why the automated detection pipeline succeeds in the first place. Because a diffraction grating inherently separates light by wavelength, the apparent position and spacing of the bands shift depending on the colour of the light being observed. This chromatic behaviour is exactly why the broadband CLEAR-filter subtraction technique (Section 3) is able to isolate these faint bands: by subtracting the spectrally flat background of the E ring, only the wavelength-dependent diffraction signature remains. The chromatic effect is not just an observational quirk; it is the fundamental optical property that makes the bands detectable at all.

5.2. Band contrast δ

Even with a clean Fourier detection, individual bands are often so faint that they risk being lost




Figure 12: Distribution of the deviation angle $\Delta\theta$ between the observed band direction and the direction predicted by the diffraction model, for all 62 positive detections. The diffraction model predicts $\Delta\theta = 0^\circ$, i.e. the band orientation is fully determined by the Sun–ring–Cassini scattering geometry with no free parameters. Blue: five flybys in close agreement with this prediction ($n=39$, median $\Delta\theta = 1.6^\circ$); orange: 098EN and 106EN combined ($n=23$), sharing a systematic offset of approximately 12° .

in residual noise. To pull the signal out, a multi-slice stacking technique is applied. For each positive detection, 15 parallel slices are extracted perpendicular to the detected orientation θ^* and averaged together (the number 15 was selected empirically as a practical balance between averaging gain and spatial coverage of the band). Stacking reinforces the brightness of the band, which is spatially coherent across all 15 slices and therefore adds constructively, while reducing uncorrelated background noise by a factor of approximately $\sqrt{15} \approx 4$.

From this stacked profile, the photometric contrast metric δ is defined as

$$\delta = I_{\max} - I_{\min}, \quad (4)$$

where I_{\max} is the peak intensity and I_{\min} is estimated from the two nearest flanking troughs. This approach cancels image-wide radiometric offsets common in archival ISS data, giving a single photometric quantity that can be compared across detections from different flybys, filters, and orbital geometries. Standard Michelson contrast is not used because CLEAR-subtraction produces negative residual values that make the denominator physically uninterpretable. A per-image signal-to-noise ratio $\text{SNR} = \delta/\sigma_{\text{bg}}$ is also computed, where σ_{bg} is the standard deviation of the profile in quiet off-band regions.

Contrast measurements were obtained for 56 of the 62 positive detections; the remaining 6 could not be characterised. Most bands sit near the detection threshold, with individual contrast values spanning

1.2×10^{-6} to $5.8 \times 10^{-5} I/F$ and a median SNR of 2.3.

Figure 13 shows the profile extraction for image N1711553432 (IR3, orbit 163EN), the running example of this article. The stacked profile (right panel) shows a well-defined peak rising above the local background, with $\delta = 2.9 \times 10^{-5} I/F$ and $\text{SNR} = 7.0$.

It should be noted that δ is measured in raw I/F units from the CLEAR-subtracted residual and is not normalised by the local E ring background brightness; observed trends with geometry may therefore partly reflect variations in the overall scene brightness rather than intrinsic band-to-background contrast. This limitation is discussed in Section 6.3 and the trends below should be treated as indicative rather than physically definitive.

To investigate whether band contrast depends on observing geometry, δ is examined as a function of phase angle, Cassini–Enceladus distance, Cassini–Saturn distance, and integrated line-of-sight column through the E ring (total ray-path length within a cylindrical-annulus model bounded radially by 180,000–480,000 km from Saturn and vertically by $|z| \leq 2,000$ km), for the four filters with four or more contrast measurements (BL1, RED, IR1, IR3). Figure 14 shows δ as a function of each geometric variable, separately for each of the four filters.

The most consistent pattern across all four filters is a decrease in δ with increasing Cassini–Enceladus and Cassini–Saturn distance: bands detected at closer distances are systematically more contrasted. A moderate negative trend with phase angle is also apparent in the IR1 and IR3 filters, possibly reflecting a difference in the phase function of the band material relative to the ambient ring. No trend is found with the integrated E ring line-of-sight column, suggesting that the bands are spatially localised features rather than an accumulation effect along the full ring width. Given the small sample per filter (4–15 detections), these trends are indicative rather than definitive and should be revisited with a larger catalogue before physical interpretations are drawn.

5.3. Viewing geometry

Figure 15 shows the viewing geometry of all positive and doubtful detections projected onto Saturn's equatorial plane. In all cases Cassini was positioned well outside the E ring (8–12 R_S), looking back through the ring at a high phase angle. No detections fall in the morning ansa of Saturn's ring system; all lines of sight sample the evening-side or anti-solar hemisphere, consistent with the forward-scattering geometry required for luminous-band visibility and with the coverage imposed by the selection constraints of Table 1. To verify that this absence is not a sampling bias, Figure 16 shows the azimuthal distribution of all 624 processed image pairs as a histogram of Enceladus

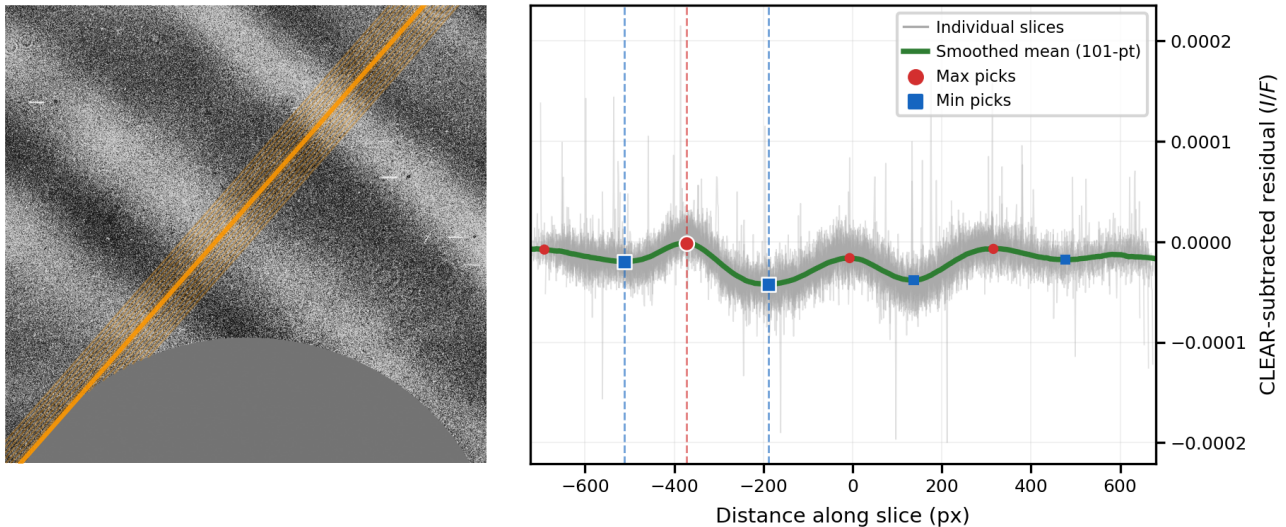


Figure 13: Multi-slice profile extraction for N1711553432 (IR3, orbit 163EN, 2012-03-27), the running example of this article ($\delta = 2.9 \times 10^{-5} I/F$, $\text{SNR} = 7.0$). *Left:* preprocessed display residual with the 15 slice lines overlaid. *Right:* individual slice samples (grey) and the smoothed stacked mean (green, 101-point moving average); red circles mark max picks and blue squares mark min picks; dashed lines indicate the positions of the max and flanking min picks used to compute $\delta = I_{\text{max}} - I_{\text{min}}$.

orbital longitude, with positive detections overlaid. The morning ansa is sampled by a comparable number of pairs as the evening and anti-solar ansae; the absence of positive detections there is not a consequence of incomplete survey coverage.

6. Discussion

Rubbrecht et al. (2025) reported luminous-band detections primarily in VIMS data from the E13, E17, and E19 Enceladus flybys, with additional ISS detections confined to E17 (flyby 163EN in this survey's notation). The present work extends the detection catalogue substantially. The survey spans eight flybys, six of which had no prior published luminous-band search. Luminous bands are therefore not a property exclusive to the encounters studied by Rubbrecht et al. (2025), but recur across a much broader range of observing geometries and orbital epochs. Five of the eight are non-targeted flybys, in the sense that they were not specifically planned for plume or E ring science, demonstrating that luminous bands are detectable in archival data acquired for other purposes and that dedicated targeting is not a prerequisite for detection. The three targeted flybys in the catalogue are 163EN (E17), 142EN (E13), and 131EN (E10); of these, only 163EN was previously searched in ISS, and 142EN was searched in VIMS only. The present pipeline provides the first ISS detection of luminous bands in the E13 encounter.

The 163EN (E17) detections in this survey include the images analysed by Rubbrecht et al. (2025); the present pipeline independently recovers those bands,

confirming that the automated Fourier approach reproduces results obtained by expert visual inspection in the reference study. Beyond extending the catalogue, the deviation-angle analysis of Section 5.1 places all 62 detections in a physically motivated coordinate frame. The close agreement with the diffraction model for 39 of 62 detections ($\Delta\theta < 2^\circ$) is consistent with the grating interpretation of Rubbrecht et al. (2025), extending the supporting evidence to a sample six times larger than previously available and across a much greater diversity of viewing geometries.

A point that the present results reinforce is the chromatic and spatially localised nature of luminous bands. Rubbrecht et al. (2025) demonstrated that the bands have a spectral signature distinct from the ambient E ring continuum, being more prominent in certain narrow-band filters than in the broadband CLEAR channel. The filter-dependent detection rates in this survey (IR3, IR1, RED, and BL1 together account for 84% of all detections) are therefore physically meaningful rather than a consequence of image counts; they reflect the spectral energy distribution of the band material relative to the E ring background, and future work calibrating these rates for exposure time and filter bandwidth could constrain the band spectral slope.

6.1. Morning-ansa absence and size-selective grain clearing

The complete absence of positive detections in the morning ansa (Section 5.3), despite a survey that samples all four ansae comparably (Figure 16), is a key observational result that demands physical interpretation. Hedman et al. (2012) documented systematic

Luminous bands in Saturn's E ring

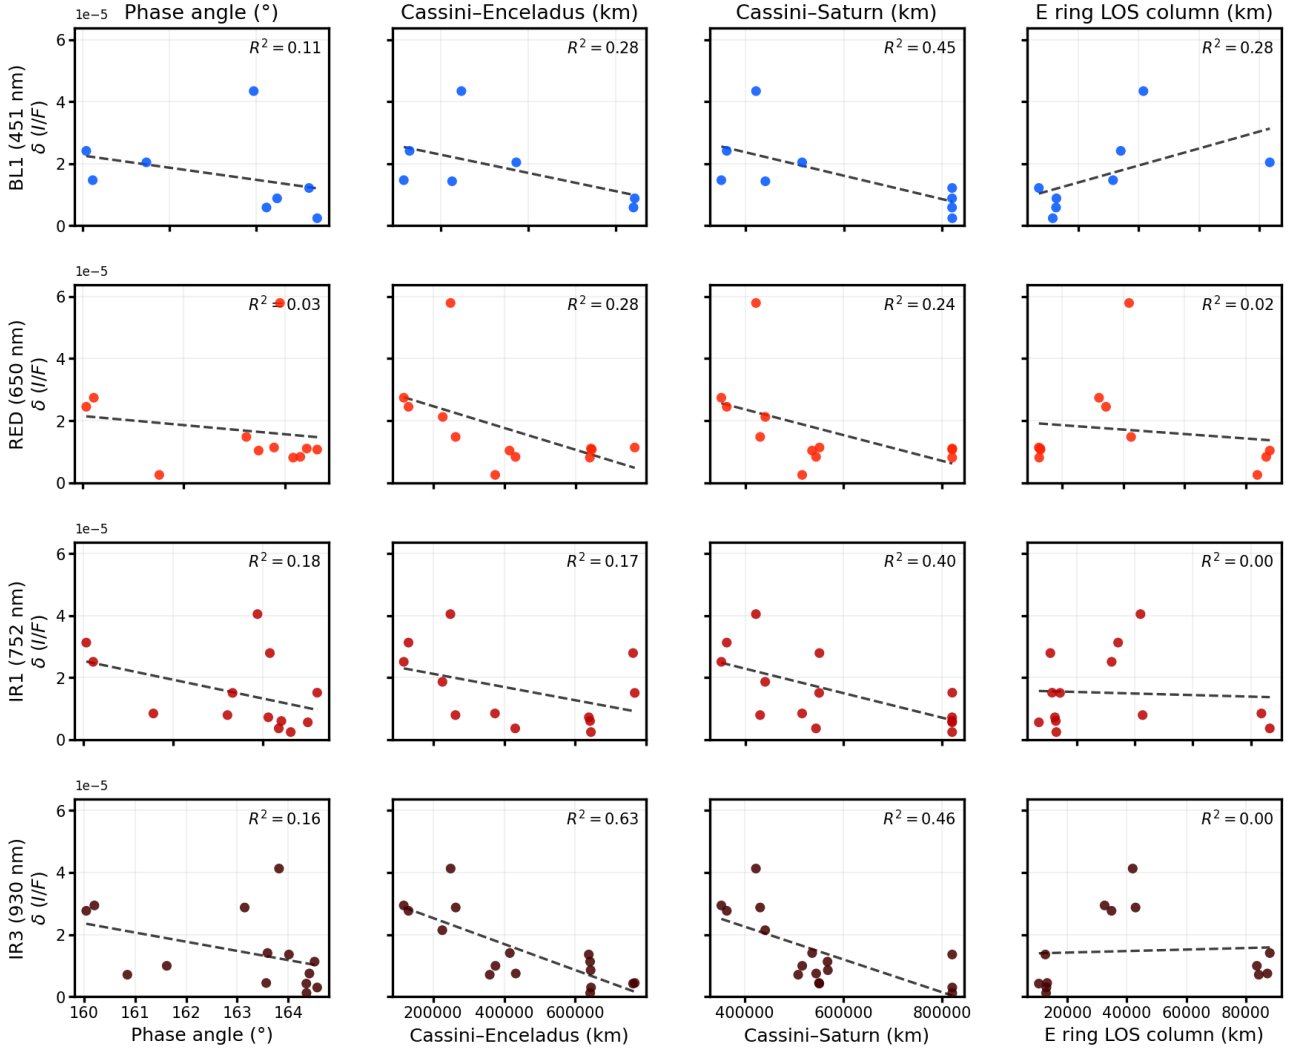


Figure 14: Band contrast δ as a function of observing geometry for the four filters with four or more contrast measurements (rows: BL1, RED, IR1, IR3; columns: phase angle, Cassini–Enceladus distance, Cassini–Saturn distance, integrated E ring line-of-sight column). Dashed lines show linear fits to guide the eye. A systematic decrease in contrast with Cassini–Enceladus and Cassini–Saturn distance is visible in all four filters, while no trend is apparent in the LOS column.

longitudinal asymmetries in the inner E ring attributed to electromagnetic perturbations at Saturn’s shadow boundary: fine charged grains crossing into shadow experience an abrupt halt in photoelectric emission, shifting the equilibrium grain potential and delivering a size-dependent Lorentz impulse that alters the grain’s orbital parameters. If this mechanism selectively depletes the sub-micron grain population on the morning side interior to Enceladus’s orbit (Hedman et al., 2012), a luminous band with the same intrinsic periodic structure could fall below the ISS detection threshold there. At the high phase angles of this survey, band visibility depends on forward scattering by sub-micron grains; their depletion would reduce the contrast δ below the detection threshold even if the underlying periodic structure persists. The coverage analysis (Figure 16) confirms this absence is not a

sampling artefact: the morning ansa is observed comparably to the other three. The proposed explanation is therefore that shadow-induced electromagnetic forcing selectively removes sub-micron grains from the morning ansa, depleting precisely the grain population responsible for band visibility at high phase angles.

To probe whether the morning ansa shows a measurable spectroscopic difference from the other three ansae, ambient E ring background spectra were extracted from Cassini VIMS data, using flybys drawn from the same 23-flyby dataset as the ISS survey. For each of 3–4 flybys per ansa, a median background spectrum was computed from five representative VIMS cubes, selecting only pixels located away from the Enceladus disc and the plume column, a critical selection requirement since plume pixels carry a distinct spectral signature dominated by freshly emitted micron-sized




Figure 15: Viewing geometry of positive (blue) and doubtful (orange) detections projected onto Saturn's equatorial plane. Each line connects the Cassini position to the corresponding Enceladus position; reference circles mark the F, G, E, T, D, and R ring and moon-orbit radii. All observations sample the evening-side or anti-solar hemisphere from outside the E ring at high phase angle.

water-ice grains (Hedman et al., 2009) and would bias the ambient ring estimate. Per-cube medians were averaged within each flyby and then across flybys to produce one representative background spectrum per ansa; uncertainties are the standard deviation across individual cube medians. This comparison should be treated as a first-pass diagnostic: the per-ansa sample is small (~ 15 – 20 cubes). All 23 flybys were selected under the same high-phase, near-ring-plane constraints (Section 2), so the viewing geometry is broadly comparable across ansae, though a formal phase-angle and emission-angle matching has not been performed. The steps required to place the comparison on a firmer footing are described in Section 6.3.

The results are shown in Figure 17. The $2.2\ \mu\text{m}$ continuum level is comparable across all four ansae, suggesting the morning ansa has a column density of particles comparable to the other ansae. The key divergence appears at longer wavelengths: from $\approx 2.8\ \mu\text{m}$ onward the morning ansa spectrum falls below those of the sub-solar, anti-solar, and evening ansae and remains suppressed throughout the 3.0 – $4.5\ \mu\text{m}$ water-ice absorption region.

A uniform reduction in particle number would decrease I/F across all wavelengths while leaving the spectral shape unchanged. A decrease concentrated in the 3.0 – $4.5\ \mu\text{m}$ region at a comparable short-wavelength continuum instead signals a change in the grain size distribution (Bohren and Huffman, 1983).

Sub-micron grains, whose size is small relative to the absorption length at these wavelengths, scatter predominantly from their surfaces and maintain a relatively elevated I/F even in this strongly absorbing regime. Micron-scale and larger grains, by contrast, allow sufficient photon penetration for volume absorption to decrease I/F across the full 3.0 – $4.5\ \mu\text{m}$ interval. A population relatively depleted of sub-micron grains therefore appears darker in this region while its short-wavelength continuum is preserved. It should be noted that Hedman et al. (2012) reported a more nuanced picture than a simple morning-side brightness deficit. Their high-phase mosaics of the morning ansa show a brightness enhancement between $210,000$ and $240,000\ \text{km}$ near Enceladus's orbit, but a sharp brightness depletion interior to $\approx 210,000\ \text{km}$. The sub-micron depletion hypothesis is relevant specifically in the inner radial zone where shadow-charging dynamics scatter the smallest grains out of the ring plane at a much higher rate than larger ones. It should be noted that the 3.0 – $4.5\ \mu\text{m}$ suppression is interpreted here as a grain-size signature, but amorphization of morning-ansa ice could contribute an additional spectral modification in the same region: both UV Ly- α photons and magnetospheric ions efficiently amorphize crystalline water ice (Leto and Baratta, 2003), and grains that have experienced more cumulative shadow crossings and radiation exposure would be expected to be more amorphized than freshly delivered material near the sub-solar ansa. Amorphization shifts the $3\ \mu\text{m}$ band minimum position and reshapes its profile, effects distinct in principle from the broad amplitude suppression produced by grain-size depletion. The spectral resolution and signal-to-noise of Figure 17 do not cleanly separate the two contributions; a dedicated analysis of the $3\ \mu\text{m}$ band minimum position per ansa would be required to test the amorphization hypothesis independently.

The physical chain connecting the shadow boundary to a size-selective grain depletion is theoretically well-established. As a charged grain enters Saturn's shadow, incoming solar radiation is cut off and photoelectric emission ceases abruptly. The stopping of this emission shifts the equilibrium grain potential to a more negative value, and the resulting Lorentz force delivers a size-dependent impulsive kick to the grain's orbit (Horanyi et al., 1992; Hamilton, 1993). The dynamical effect scales as Q/m , the charge-to-mass ratio of the grain. Because the equilibrium surface potential is set by the local plasma environment rather than by grain size, the charge Q scales as $Q \propto r$ (isolated-sphere capacitance in a plasma whose Debye length greatly exceeds the grain radius r), while mass m scales as $m \propto r^3$, giving $Q/m \propto r^{-2}$. Sub-micron grains therefore receive Lorentz impulses orders of magnitude stronger per shadow crossing than micron-sized grains; repeated crossings progressively pump their orbital eccentricities

Luminous bands in Saturn's E ring

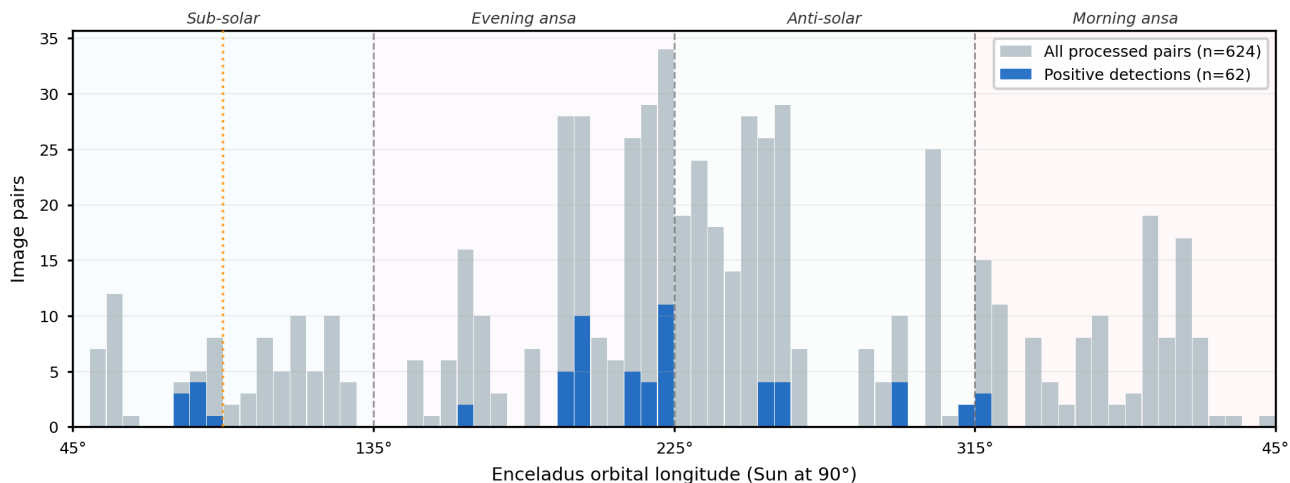


Figure 16: Azimuthal survey coverage: number of (CLEAR, science-filter) image pairs as a function of Enceladus orbital longitude, unfolded over 360°. The four ansae are indicated by shading and italic labels; the Sun direction falls at 90° (orange dotted line) in the middle of the sub-solar quadrant. Gray bars show all 624 processed pairs; blue bars show the 62 positive detections. All observations were acquired under similar viewing geometry (Cassini outside the E ring at high phase angle; see Figure 15). Detections are present in the sub-solar, evening, and anti-solar ansae but are absent from the morning ansa (right quadrant), despite comparable survey coverage in that region.




Figure 17: VIMS ambient E ring background spectra per ansa, averaged over 3–4 flybys and 5 cubes each ($w=11$ moving average). Median I/F from 1.0 to 5.0 μm . The four ansae are nearly identical at short wavelengths but diverge from $\approx 2.8 \mu\text{m}$ onward, where the morning ansa (green) is systematically suppressed through the entire 3.0–4.5 μm region while the 2.2 μm continuum remains comparable across all four ansae.

and inclinations, scattering them out of the ring plane or onto Saturn-impacting orbits (Juhász and Horányi, 2002, 2004). The morning ansa, immediately downstream of the shadow exit, is where this size-selective clearing should be most pronounced.

Independent in-situ evidence that electromagnetic forces sculpt the E ring grain population comes from the Cassini Cosmic Dust Analyzer (CDA; Srama et al., 2004). Kempf et al. (2008) constructed three-dimensional E ring density profiles from CDA impact rates and documented orbit-to-orbit grain number-density fluctuations of order 20% near Enceladus’s

orbit, with some evidence for angular variation in the grain size distribution. Hsu et al. (2011) further demonstrated that the smallest grains in the Saturn system are violently size-selected by the dust–plasma–magnetosphere interaction, with charge-to-mass scaling forcing nanometre-scale particles out of the ring entirely, providing model-independent evidence that the Saturnian dust population is sculpted in a strongly grain-size-dependent fashion.

The connection to the ISS non-detections follows directly. At the high phase angles of this survey (150°–180°), the forward-scattered brightness of the E ring is

dominated by the sub-micron grain population, whose broad forward-scattering peak intercepts Cassini's ISS at these angles; micron-scale grains concentrate scattered light into a much narrower peak confined close to the exact forward direction (Bohren and Huffman, 1983), consistent with constraints from high-phase eclipse imaging (Ingersoll and Ewald, 2011), which derived a median equivalent-volume grain radius of $3.1 \pm 0.5 \mu\text{m}$ from phase angles of 173° – 179° and confirmed that sub-micron grains contribute disproportionately to the visible I/F at the geometries of the present survey. Luminous bands produce a detectable contrast δ only when a sufficient column of such fine grains is present above the local ring background. If shadow-induced electromagnetic clearing has selectively depleted this sub-micron population on the morning side, band structures, even if the underlying periodic geometry persists, would generate too little forward-scattered flux to exceed the photometric detection threshold of the ISS pipeline. The morning-ansa absence is therefore consistent with, and potentially explained by, the selective removal of precisely the grain population that makes luminous bands visible at high phase angle. The observed decrease in band contrast δ with Cassini–Enceladus distance (Section 5.2) is consistent with this picture: band structure is most pronounced close to the plume source, where freshly injected material has not yet experienced repeated shadow crossings.

A key question is whether this picture is consistent with the spectral characterisation of Rubbrecht et al. (2025), whose analysis associates the luminous-band grating with micron-scale crystalline water-ice grains rather than with the sub-micron component. Micron-sized grains, which constitute the mm-scale periodic density clumping identified by Rubbrecht et al. (2025) as the physical grating, experience Lorentz impulses orders of magnitude weaker than sub-micron grains ($Q/m \propto r^{-2}$) and are therefore relatively unaffected by shadow-induced electromagnetic clearing. The sub-micron grains co-located with this structure supply the forward-scattered brightness recorded by ISS at high phase angles; their selective removal reduces the band contrast below the detection threshold without disrupting the underlying periodic geometry. That the grating material is spectroscopically fresh and crystalline (Rubbrecht et al., 2025) further implies continuous plume replenishment: UV and magnetospheric ion irradiation amorphize crystalline ice on timescales set by the local radiation dose accumulation (Leto and Baratta, 2003), so the persistence of a crystalline signature constrains the grating to consist of recently sourced material. A falsifiable prediction follows: VIMS, whose 1 – $5 \mu\text{m}$ sensitivity is matched to the surviving micron-scale scatterers (Hedman et al., 2009; Rubbrecht et al., 2025), could detect luminous bands on the morning ansa in the near-infrared even where ISS cannot. A targeted search for band features in morning-ansa

VIMS cubes from the flybys surveyed here is the most direct observational test of the size-selective clearing hypothesis.

6.2. Systematic band deviation in 106EN and 098EN: ring warp or attitude error

The 23 detections from the 106EN and 098EN flybys show a consistent offset of $\Delta\theta \approx 11$ – 13° that is absent from all other six flybys in the survey. Its stability across multiple filters and two entirely independent encounters rules out random measurement noise: it is a repeatable signal. The same pipeline that recovers band directions to within 1.6° for 39 other detections resolves this offset sharply, making it a reliable feature of those specific observing configurations rather than an artefact of the detection method.

The origin of this offset is not yet established. It could have a physical origin in the local ring geometry at those specific orbital epochs, or it could stem from a systematic error in the SPICE-reconstructed attitude for those Cassini orbit clusters. Isolating the source requires a dedicated geometry audit of the 106EN and 098EN observations and is identified as a priority for future work.

6.3. Limitations, open pitfalls, and future directions

Several limitations of the present analysis are worth making explicit, both as caveats on the reported results and as a roadmap for future improvements.

Simplified E ring geometry. The integrated line-of-sight column used in the contrast trend analysis assumes a homogeneous, flat-torus E ring model ($r \in [180\text{k}, 480\text{k}] \text{ km}$, $|z| \leq 2 \text{ km}$) with no density weighting. In reality the E ring is denser near Enceladus, falls off radially and vertically, and has azimuthal structure tied to Enceladus's orbital phase and plume activity. Longer lines of sight through the outer, lower-density ring are not equivalent to shorter paths through the inner ring in terms of total grain column. This likely partly explains the absence of a trend between δ and the integrated LOS column; a density-weighted column model would be a more informative diagnostic.

CLEAR-subtraction baseline bias. The direct CLEAR subtraction corrects the broadband scene continuum but does not account for spatially variable scattered light or short-timescale flux variations between the two exposures. When these differ between the CLEAR and science frames, the residual I_{res} carries a DC offset that shifts the field as a whole. This does not affect detection, which operates on spatial frequency content rather than the absolute level, but it means the absolute scale of \hat{I}_{res} should not be used to compare filter brightnesses directly across pairs. A per-pair additive correction from an unstructured off-disc region of the residual would mitigate this.

Review-driven labelling subjectivity. The hybrid detection verdict (automated spectral analysis followed by human adjudication of a fixed diagnostic panel) is more reproducible than a fully manual approach, but the labelling of borderline cases still carries some subjectivity. A double-blind cross-check by independent reviewers on a random subset of the catalogue would quantify the inter-rater reliability and flag the most ambiguous cases; this has not been performed here.

Angular projection and universality. The image-plane band orientations reported here are transformed to celestial position angles using SPICE, but the universality of this angular projection across different encounter geometries has not been fully verified. Future work should confirm that the derived position angles are consistent across the two distinct Cassini attitude clusters identified in Section 5.1.

Temporal completeness. The automated archive query produces no results beyond mid-2012, despite the Cassini mission continuing until 2017. The reason for this cut-off has not been verified and may reflect a gap in OPUS metadata, a change in observing strategy, or a shift in the ring-plane geometry that falls outside the search constraints. The later orbital phase includes additional close Enceladus encounters and a wider range of ring-plane geometries, so a verification of the query output against the full mission archive is an important first step before concluding that no usable data exist in that period. Extending the pipeline to the full archive would substantially enlarge the catalogue and test whether the detection rate, azimuthal asymmetry, and contrast trends are stable over the full mission lifetime.

Extended target coverage. The present OPUS query uses Enceladus as the sole target name. A complementary search targeting the E ring directly, or flybys where Enceladus was not the primary target but the E ring was nonetheless in the NAC field at a suitable phase angle, would extend the catalogue without any modification to the existing pipeline.

Phase-angle coverage below 150°. The detection statistics hint that band contrast increases toward lower phase angles within the surveyed range of 150°–180°. Whether bands remain detectable below 150°, and at what phase angle the forward-scattering peak becomes too narrow to intercept the observer, cannot be assessed here. Future ring-plane imaging campaigns at 140°–150° phase would directly constrain the angular scattering width of the band-forming grain population.

Physical origin. The synthetic data experiments described by Rubbrecht et al. (2025), generating artificial periodic stripes superimposed on realistic ring images, confirmed that the Fourier pipeline responds as expected to controlled inputs and helped establish confidence in the detection method. Extending this approach to test for systematic false-positive signatures from specific artefacts (e.g., partially subtracted disc

halos or cosmic-ray streaks aligned with the ring plane) is a useful next step for validating the pipeline's false-alarm rate.

VIMS spectral comparison. The per-ansa spectral comparison in Section 6.1 is a first-pass analysis. The per-ansa sample of ~ 15 – 20 cubes was not selected by a deterministic rule; phase angle and emission angle are not matched between ansae; and the grain-size interpretation is qualitative rather than model-fitted. Placing the comparison on firmer statistical footing, applying a matched-geometry selection, and overlaying Mie forward models for candidate grain-size distributions (Hedman et al., 2009) would be the main steps toward converting the spectral shape into a quantitative constraint. Water-vapour emission near $2.65\ \mu\text{m}$ (Denny et al., 2024) also falls within the suppression band and should be verified as consistently excluded by the plume-pixel mask.

Morning-ansa VIMS band search. The size-selective clearing hypothesis predicts that the periodic grating survives on the morning ansa even where ISS cannot detect it, because the micron-scale grains that constitute the grating are dynamically resilient to Lorentz forcing. A targeted search for band features in morning-ansa VIMS cubes at wavelengths $\gtrsim 3\ \mu\text{m}$, where micron-scale scatterers dominate, is the most direct observational test of this prediction.

7. Conclusions

This paper presents the first automated survey for luminous bands in Cassini ISS data. The main results are as follows.

- The survey covers 573 (CLEAR, science-filter) image pairs from 23 Cassini flybys. The detection pipeline identifies 62 positive detections across eight flybys; six of these had no prior published luminous-band search, and five are non-targeted flybys. The 142EN (E13) detections are the first confirmed in ISS for that flyby, extending the catalogue of Rubbrecht et al. (2025) to a new encounter and instrument combination.
- The observed band orientations are compared against directions predicted by the diffraction model of Rubbrecht et al. (2025). For 39 of the 62 detections, spanning five flybys, the median deviation angle is only 1.6° , confirming the grating interpretation on a sample six times larger than previously available. The remaining 23 detections, from the 106EN and 098EN flybys, share a systematic offset of $\Delta\theta \approx 11$ – 13° that is stable across filters and two independent flybys; its origin, whether physical or related to attitude reconstruction, requires a dedicated geometry audit.

- Band contrast decreases with increasing Cassini–Enceladus and Cassini–Saturn distance: the bands are brightest close to Enceladus and Saturn. No correlation is found with the integrated E ring line-of-sight column, suggesting the bands are spatially localised structures.
- No detections fall in the morning ansa despite comparable survey coverage. A complementary VIMS spectral analysis shows that the morning-side ambient E ring has a decreased I/F across the 3.0–4.5 μm water-ice absorption region at a comparable short-wavelength continuum. This microphysical signature is consistent with shadow-induced Lorentz forcing selectively removing sub-micron grains from the morning ansa (Hedman et al., 2012), depleting precisely the grain population responsible for band visibility at high phase angles.

Together, these results establish luminous bands as a widespread feature of Saturn's E ring whose detectability is regulated by the sub-micron grain population, an azimuthally asymmetric resource shaped by electromagnetic forces. The pipeline provides the foundation for extending this survey to the full Cassini archive (2012–2017).

Data and code availability

The PyISS package used for OPUS query construction and observation-set inference is openly available at <https://github.com/ArnaudMath/pyiss>. The detection and characterisation pipeline, including all post-pipeline scripts, is available at <https://github.com/ArnaudMath/LuminousBandsInERing>. Cassini ISS calibrated images were retrieved from the NASA Planetary Data System (PDS) via the OPUS search interface (<https://opus.pds-rings.seti.org>). SPICE kernels were obtained from the NASA Navigation and Ancillary Information Facility (NAIF).

A. PyISS: automated observation-set reconstruction

PyISS is a Python package developed to automate the identification and reconstruction of co-temporal ISS NAC observation sets from the Cassini archive. It operates in two stages: a geometry-driven query and an observation-set reconstruction.

The query stage communicates with the PDS Ring-Moon Systems Node's OPUS search service through its public metadata API, filtering frames by geometric quantities computed from SPICE kernels (Acton, 1996; Acton et al., 2018). The search constraints applied in this survey are listed in Table 1. The query parameters are fully configurable, so the same package can be

reused for surveys with different geometric requirements without modifying the reconstruction logic.

The reconstruction stage identifies observation sets from the metadata returned by the query. Because the Cassini archive does not expose filter-sequence groupings directly, PyISS recovers them by querying OPUS for all NAC frames within a configurable time window around each seed frame, then retaining those that share the same pointing. Frames already visited during a prior seed query are cached to avoid redundant API calls. The output is two relational tables: a *sets table* with one row per observation set, recording start and end times, constituent filters, and the triggering seed frame; and a *members table* with one row per frame, recording filter, exposure time, and OPUS identifier. Applied to the targeted-Enceladus archive under the constraints of Table 1, PyISS recovers 382 distinct observation sets across 23 Cassini orbit tags spanning 2007–2012.

References

- Acton, C., Bachman, N., Semenov, B., Wright, E., 2018. A look towards the future in the handling of space science mission geometry. *Planetary and Space Science* 150, 9–12. doi:10.1016/j.pss.2017.02.013.
- Acton, C.H., 1996. Ancillary data services of NASA's navigation and ancillary information facility. *Planetary and Space Science* 44, 65–70. doi:10.1016/0032-0633(95)00107-7.
- Bohren, C.F., Huffman, D.R., 1983. *Absorption and Scattering of Light by Small Particles*. John Wiley & Sons, New York.
- Denny, K.E., Hedman, M.M., Bockelée-Morvan, D., Filacchione, G., Capaccioni, F., 2024. Constraining time variations in Enceladus's water-vapor plume with near-infrared spectra from Cassini's visual and infrared mapping spectrometer. *The Planetary Science Journal* 5, 144. URL: <https://doi.org/10.3847/PSJ/ad4c69>, doi:10.3847/PSJ/ad4c69.
- Hamilton, D.P., 1993. Motion of dust in a planetary magnetosphere: orbit-averaged equations for oblateness, electromagnetic, and radiation forces with application to Saturn's E ring. *Icarus* 101, 244–264. doi:10.1006/icar.1993.1017.
- Hedman, M.M., Gosmeyer, C.M., Nicholson, P.D., Sotin, C., Brown, R.H., Clark, R.N., Baines, K.H., Buratti, B.J., Showalter, M.R., 2012. An observed correlation between the vertical structure and the optical depth of Saturn's E ring. *Icarus* 217, 322–338. doi:10.1016/j.icarus.2011.11.014.
- Hedman, M.M., Nicholson, P.D., Showalter, M.R., Brown, R.H., Buratti, B.J., Clark, R.N., 2009. Spectral observations of the Enceladus plume with Cassini-VIMS. *The Astrophysical Journal* 693, 1749–1762. doi:10.1088/0004-637X/693/2/1749.
- Horányi, M., Burns, J.A., Hamilton, D.P., 1992. The dynamics of Saturn's E ring particles. *Icarus* 97, 248–259. doi:10.1016/0019-1035(92)90131-P.
- Hsu, H.W., Kempf, S., Postberg, F., Trieloff, M., Burton, M., Roy, M., Moragas-Klostermeyer, G., Srama, R., 2011. Cassini dust stream particle measurements during the first three orbits at saturn. *Journal of Geophysical Research: Space Physics* 116. doi:10.1029/2010JA015959.
- Ingersoll, A.P., Ewald, S.P., 2011. Total particulate mass in Enceladus plumes and mass of Saturn's E ring inferred from Cassini ISS images. *Icarus* 216, 492–506. doi:10.1016/j.icarus.2011.09.018.
- Juhász, A., Horányi, M., 2002. Saturn's E ring: A dynamical approach. *Journal of Geophysical Research: Space Physics* 107, SMP 1–1–SMP 1–10. doi:10.1029/2001JA000182.

- Juhász, A., Horányi, M., 2004. Seasonal variations in Saturn's E-ring. *Geophysical Research Letters* 31. doi:10.1029/2004GL020999.
- Kempf, S., Beckmann, U., Moragas-Klostermeyer, G., Postberg, F., Srama, R., Economou, T., Schmidt, J., Spahn, F., Grün, E., 2008. The E ring in the vicinity of Enceladus I: Spatial distribution and properties of the ring particles. *Icarus* 193, 420–437. doi:10.1016/j.icarus.2007.06.027.
- Knowles, B., West, R., Helfenstein, P., Verbiscer, A., Wilson, D., Porco, C., 2020. End-of-mission calibration of the Cassini imaging science subsystem. *Planetary and Space Science* 185, 104898. doi:10.1016/j.pss.2020.104898.
- Leto, G., Baratta, G.A., 2003. Ly- α photon induced amorphization of I_c water ice at 16 Kelvin: Effects and quantitative comparison with ion irradiation. *Astronomy & Astrophysics* 397, 7–13. doi:10.1051/0004-6361:20021473.
- Porco, C.C., Helfenstein, P., Thomas, P.C., et al., 2006. Cassini observes the active south pole of Enceladus. *Science* 311, 1393–1401. doi:10.1126/science.1123013.
- Porco, C.C., West, R.A., Squyres, S.W., McEwen, A., Thomas, P., Murray, C.D., DelGenio, A.D., Ingersoll, A.P., Johnson, T.V., Neukum, G., et al., 2004. Cassini imaging science: Instrument characteristics and anticipated scientific investigations at Saturn. *Space Science Reviews* 115, 363–497. doi:10.1007/s11214-004-1456-7.
- Postberg, F., Schmidt, J., Hillier, J.K., Kempf, S., Srama, R., 2011. A salt-water reservoir as the source of a compositionally stratified plume. *Nature* 474, 620–622. doi:10.1038/nature10175.
- Rubbrecht, N., Cazaux, S., Seignovert, B., Kenworthy, M., Kutsop, N., Le Mouélic, S., Loicq, J., 2025. Peculiar rainbows in Saturn's E ring: Uncovering luminous bands near Enceladus. *Icarus* 441, 116650. doi:10.1016/j.icarus.2025.116650.
- Spahn, F., Schmidt, J., Albers, N., Hörning, M., Makuch, M., Seiss, M., Kempf, S., Srama, R., Dikarev, V., Helfert, S., Moragas-Klostermeyer, G., Krivov, A.V., Sremcevic, M., Tuzzolino, A.J., Economou, T., Grün, E., 2006. Cassini dust measurements at enceladus and implications for the origin of the E ring. *Science* 311, 1416–1418. doi:10.1126/science.1121375.
- Srama, R., Ahrens, T.J., Altobelli, N., Auer, S., Bradley, J.G., et al., 2004. The Cassini Cosmic Dust Analyzer. *Space Science Reviews* 114, 465–518. doi:10.1007/s11214-004-1435-z.

Part III

Research Chapters

About Part III

The literature study (Part I) established the scientific context for this thesis and the article (Part II) presented the findings. Part III describes how the detection system was built. Its purpose is technical: to document the software, data preparation steps, and detection algorithm that underlie the results. The scientific motivation for each methodological choice is already established in the preceding parts; these chapters focus on how.

The central idea of the detection strategy follows from the chromatic hypothesis of Rubbrecht et al. (2025): if a broadband CLEAR-filter image of the E ring is subtracted from a simultaneous narrowband science-filter image of the same field, the achromatic ring continuum cancels out and only the wavelength-dependent brightness remains. Under this hypothesis, luminous bands would survive the subtraction while the neutral background does not, making CLEAR subtraction a natural detection tool. Building and searching a catalogue of such subtracted images is the aim of the survey.

Operationalising this idea required solving three problems in sequence, each addressed by one chapter.

The first problem is archival. The Cassini ISS archive stores every individual exposure as an independent record; there is no metadata field linking a CLEAR exposure to the science-filter exposures taken alongside it a minute later. Before any CLEAR subtraction can be performed, the observation structure must be reconstructed from timing and geometric metadata. Chapter 7 describes PyISS, a Python package developed for this thesis to query the archive and recover the paired observations automatically, yielding 573 (CLEAR, science-filter) image pairs spanning 23 Enceladus flyby sequences.

The second problem is instrumental. A raw ISS exposure is not yet a residual image: it must be calibrated, the broadband continuum must be subtracted with the correct passband scaling, the bright Enceladus disc must be masked before it contaminates the analysis, and pairs with severe artefacts must be identified and excluded. Chapter 8 describes the data preparation pipeline that performs these steps.

The third problem is statistical. A CLEAR-subtracted residual image of the E ring is faint and noisy. Detecting an inclined periodic structure in such an image reliably, across 573 frames and without selecting by eye, requires a method that can separate a directional signal from isotropic noise and assess whether a candidate detection is statistically significant. Chapter 9 describes how this is achieved, following the

methodological development from a first attempt with a standard Fourier transform through to the zero-padded angular-spectrum method and phase-randomised null ensemble that form the core of the pipeline.

Chapter 8

PyISS: From Archive to Image Pairs

Detecting luminous bands in Saturn’s E ring requires a large set of calibrated image pairs. Each pair consists of a broadband CLEAR-filter exposure and a narrowband science-filter exposure of the same field, taken within a few minutes of each other. The CLEAR filter is the ISS broadband channel, transmitting all visible wavelengths without spectral selection; it records the full forward-scattered E ring continuum. A science-filter frame (BL1, GRN, RED, IR1, IR3, etc.) records the same continuum modulated by any wavelength-dependent brightness structure in the scene. Subtracting a scaled CLEAR frame from the science-filter frame isolates exactly this residual: the signal in which luminous bands appear.

The challenge is that the Cassini ISS archive stores every exposure as an independent record with no explicit grouping. In practice, Cassini acquired E ring images in *observation sets*: groups of exposures taken at the same pointing within a single scheduling block, typically spanning a few minutes. Most commonly, a CLEAR exposure is taken first, followed by several science-filter exposures. However, this is not always the case: the CLEAR exposure sometimes falls in the middle of a set rather than at the start, and some sets contain no CLEAR exposure at all. This variability means that the observation structure cannot be assumed — it must be inferred from the data. Figure 8.1 shows a representative example of one complete observation set.

PyISS is a Python package developed for this thesis to solve this problem in a reproducible and scriptable way. It provides two pipeline-facing capabilities: a query builder for OPUS (the Outer Planets Unified Search interface, NASA’s public search portal for Cassini science data), and an observation-set inference routine that reconstructs which frames belong together. Figure 8.2 summarises the data flow from the archive to the 573 image pairs that feed the data-preparation pipeline (Chapter 9).

8.1 Querying the archive

OPUS allows frames to be retrieved by geometric metadata pre-computed from the original SPICE kernels (C. H. Acton [1996](#); C. Acton et al. [2018](#)): phase angle,




Figure 8.1: A representative observation set from orbit 163EN (E17, 27 March 2012). Five sequential NAC exposures were acquired at the same pointing within 328 s, each through a different spectral filter: (a) broadband CLEAR, (b) IR3 (~ 930 nm), (c) IR1 (~ 752 nm), (d) RED (~ 617 nm), and (e) BL1 (~ 451 nm). The CLEAR frame serves as the broadband reference; the science-filter frames are the residual inputs. The time offset Δt from the CLEAR exposure is indicated for each frame. The goal of PyISS is to automatically identify and group all such sets from the archive, across 23 Cassini orbit sequences.

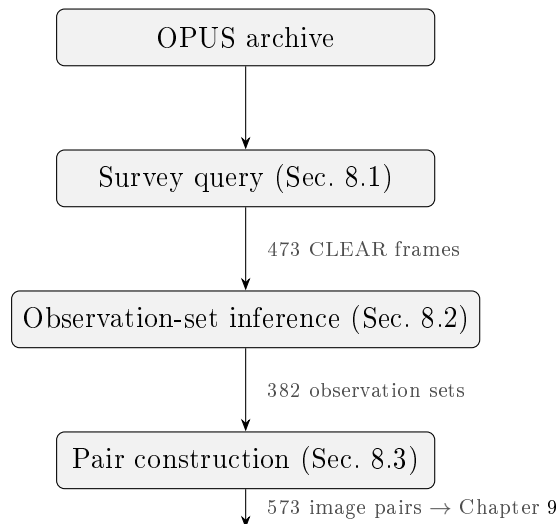


Figure 8.2: Data flow through PyISS, from the raw archive to the image pairs that enter the data-preparation pipeline (Chapter 9).

Table 8.1: Geometric constraints applied to the OPUS survey query. The scientific rationale for each parameter is discussed in Part II, Table 1.

Quantity	Adopted range
Target body	Enceladus
ISS filter	CLEAR
Phase angle	150°–180°
Ring-plane radius	180 000–480 000 km
Observer ring-plane elevation	−5°–+5°
Solar ring-plane opening angle	−15°–+15°

ring-plane radius, spacecraft elevation above the ring plane, and more.

For the luminous-band survey, the query targets CLEAR-filter NAC frames satisfying the six geometric constraints listed in Table 8.1. The phase-angle range (150°–180°) selects the high-phase forward-scattering regime in which the E ring brightens dramatically and luminous bands have been observed (Rubbrecht et al. 2025). The ring-plane-radius range spans the radial extent of the E ring, and the small elevation and solar-opening-angle limits ensure the line of sight passes near the ring mid-plane, maximising column depth. Applied to the full Cassini archive, this query returns 473 CLEAR frames across the Enceladus flyby sequences of the mission.

8.1.1 Query builder

The `query()` function exposes a simple, chainable interface: geometric constraints are added with `.param(name, value)`, the result count is capped with `.limit(n)`, and the query is executed with `.fetch(*columns)`, which issues the HTTP requests to OPUS and returns the results as a tabular dataset. The listing below shows a minimal usage example; Table 8.2 shows the five frames it returns.

Listing 8.1: Querying the archive and inferring an observation set from the first result.

```

1 from pyiss import query, infer_set
2
3 # Fetch five CLEAR frames matching the E-ring survey geometry
4 results = (
5     query()
6     .param("target", "ENCELADUS")
7     .param("COISSfilter", "CLEAR")
8     .param("phase1", 150)
9     .param("phase2", 180)
10    .limit(5)
11    .fetch("opusid", "time1", "RINGGE0phase1")
12 )
13
14 # Infer the observation set for the first returned frame
15 obs_set = infer_set(results.df.iloc[0]["opusid"])

```

Table 8.2: First five CLEAR frames returned by a representative PyISS survey query targeting the Enceladus E-ring geometry (phase 150° – 180° , ring-plane elevation $|b| < 5^\circ$, ring radius 180 000–480 000 km).

OPUS identifier	Start time (UTC)	Phase angle ($^\circ$)
co-iss-n1487335087	2005-02-17 12:11	152.64
co-iss-n1487335135	2005-02-17 12:12	152.62
co-iss-n1487335287	2005-02-17 12:13	152.58
co-iss-n1511803932	2005-11-27 17:03	161.17
co-iss-n1511803965	2005-11-27 17:03	161.17

8.2 Reconstructing observation sets

With 473 CLEAR frames identified, the next step is to find the science-filter frames that were acquired alongside each one. The archive provides no sequence identifier to link them, so membership must be inferred. The key observation is that frames acquired within the same observation set are separated by short time gaps (tens to hundreds of seconds), while frames from different sets are separated by much longer gaps (hours to days). The inference algorithm exploits this timing signature to recover the set structure automatically.

Starting from one of the 473 CLEAR frames, the algorithm retrieves a wide time window of ISS frames from the archive, both before and after that reference frame. No geometric constraints are applied during this retrieval, so science-filter frames that fall outside the survey’s phase-angle or geometry limits are still included — they are needed to reconstruct the full observation set even if they would not themselves appear in the survey query. The retrieved frames are then sorted by time and the algorithm identifies the boundaries between observation sets as points where the gap between consecutive frames jumps sharply.

8.2.1 Temporal grouping

The threshold separating gaps within a set (short) from gaps between sets (long) is found by two-means clustering on the logarithm of the gap values: an adaptive method that learns the characteristic cadence of each orbit from the data itself, rather than relying on a fixed cutoff that would not generalise across the varying scheduling of different flybys. The observation set of the reference frame is the longest unbroken sequence of frames linked by short gaps that includes it.

As a secondary check, when OPUS reports body-surface pointing metadata near a candidate boundary, a large geographic jump in pointing is used to confirm the split. This prevents two physically distinct observation pointings from being accidentally merged by a coincidentally short time gap. Once a set has been inferred, it is stored in a cache keyed on the frame identifier; any subsequent reference frame belonging to the same set is resolved from the cache without an

additional network request.

Table 8.3 shows the full observation set recovered for the example frame in the query above. The set belongs to a polarimetric sequence comprising repeated GRN+P0/P60/P120 triplets alongside UV3, IR1, and IR3 exposures; the algorithm recovers this multi-filter structure correctly without any prior knowledge of the filter pattern.

Table 8.3: Observation set inferred by `infer_set` from seed `co-iss-n1487335087` (2005-02-17). Each row is one ISS NAC frame in the sequence; the CLEAR exposure is the photometric reference used for residual construction.

OPUS identifier	Start time (UTC)	Filter
<code>co-iss-n1487334437</code>	12:00:26	GRN+P0
<code>co-iss-n1487334509</code>	12:01:38	GRN+P60
<code>co-iss-n1487334581</code>	12:02:44	GRN+P120
<code>co-iss-n1487334629</code>	12:03:52	UV3
<code>co-iss-n1487334663</code>	12:04:36	IR1
<code>co-iss-n1487334711</code>	12:05:17	IR3
<code>co-iss-n1487334751</code>	12:06:03	GRN+P0
<code>co-iss-n1487334791</code>	12:06:43	GRN+P60
<code>co-iss-n1487334831</code>	12:07:22	GRN+P120
<code>co-iss-n1487334879</code>	12:08:08	UV3
<code>co-iss-n1487334912</code>	12:08:46	IR1
<code>co-iss-n1487334952</code>	12:09:22	IR3
<code>co-iss-n1487334987</code>	12:10:00	GRN+P0
<code>co-iss-n1487335020</code>	12:10:33	GRN+P60
<code>co-iss-n1487335053</code>	12:11:06	GRN+P120
<code>co-iss-n1487335087</code>	12:11:41	CLEAR
<code>co-iss-n1487335135</code>	12:12:20	CLEAR
<code>co-iss-n1487335287</code>	12:13:22	CLEAR

8.3 Building image pairs

Once the observation sets are known, image pairs are constructed by matching each science-filter frame in a set to the temporally nearest CLEAR frame in the same set. The time separation between the two exposures is recorded; frames with no available CLEAR counterpart, or whose nearest CLEAR is flagged for saturation or pointing offset, are excluded.

Figure 8.1 in the chapter introduction illustrates the result: from one observation set, several (CLEAR, science-filter) pairs are extracted, one per science filter. Applied to all 473 CLEAR frames, the inference and pairing steps recover 382 observation sets spanning 23 Cassini orbit sequences between 2007 and 2012, yielding 573 image pairs that enter the data-preparation pipeline described in Chapter 9.

8.4 Software

PyISS is implemented in Python and depends on `requests` for HTTP communication with OPUS, `pandas` for tabular data handling, and `scikit-learn` for the two-means clustering step. The package is available at <https://github.com/ArnaudMath/pyiss>; the version used for this thesis is v0.4.2. The 573 image pairs it produces are the input to the data-preparation pipeline described in Chapter 9.

Chapter 9

Data Preparation Pipeline

The 573 image pairs assembled by PyISS (Chapter 8) arrive as raw calibrated frames from the Planetary Data System (PDS), the NASA archive that distributes mission science data in standardised formats. Turning them into residual images ready for band detection requires four operations: CLEAR-filter subtraction, Enceladus disc masking, download and format conversion, and a final quality-control review. Figure 9.1 shows the full data flow.

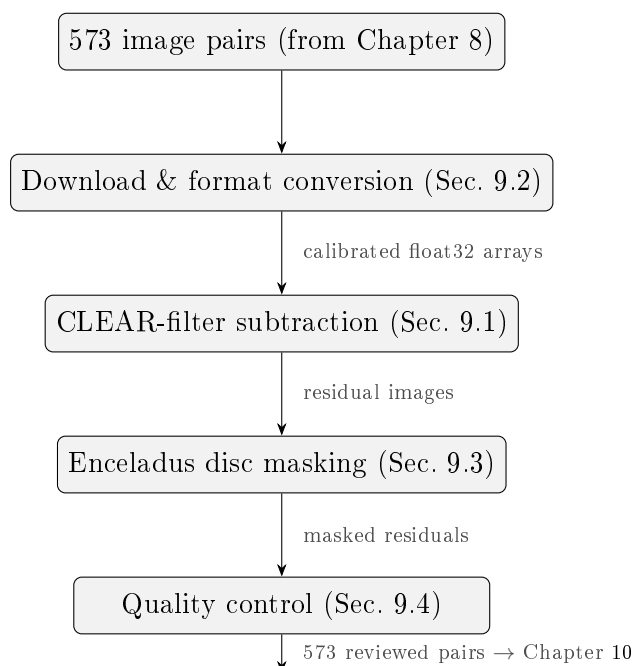


Figure 9.1: Data preparation pipeline, from the raw image pairs to the quality-reviewed residual images that enter the detection pipeline (Chapter 10).

9.1 The residual image and why it reveals bands

Rubbrecht et al. (2025) reported that luminous bands have a chromatic character: under this hypothesis, a band’s apparent position shifts with wavelength, distin-

guishing it from the E ring background. The E ring background is spectrally neutral, consisting of broadband forward-scattered light from micron-sized grains whose scattering cross section varies little across the ISS filter range. Subtracting a broadband CLEAR-filter image from a narrowband science-filter image therefore removes the achromatic continuum and isolates the wavelength-dependent brightness in which, under this hypothesis, the bands would appear.

For each image pair, the residual is constructed as

$$I_{\text{res}} = I_X - k I_{\text{CLEAR}}, \quad (9.1)$$

where I_X is the science-filter I/F image, I_{CLEAR} is the simultaneous CLEAR-filter I/F image, and k is a passband scaling factor that corrects for the throughput difference between the two filter paths. Here I/F is the standard reflectance unit used by the Cassini calibration pipeline: the ratio of observed intensity to the intensity of a Lambertian surface at normal incidence. The factor k is computed from the ratio of solar-spectrum-weighted passband integrals,

$$k = \frac{\int T_X(\lambda) Q(\lambda) F_{\odot}(\lambda) d\lambda}{\int T_{\text{CLEAR}}(\lambda) Q(\lambda) F_{\odot}(\lambda) d\lambda}, \quad (9.2)$$

using the tabulated ISS system transmissions $T(\lambda)$ and detector quantum efficiency $Q(\lambda)$ from Porco et al. (2004), and the solar spectrum $F_{\odot}(\lambda)$ evaluated at Saturn's mean heliocentric distance of 9.54 au. The factor is pre-computed once per filter combination. Selecting the nearest CLEAR exposure in time minimises residuals from pointing drift and short-timescale flux variations between the two frames.

The result is an image whose pixel values are close to zero in spectrally neutral ring regions, and deviate wherever the scene has a chromatic brightness structure. Figure 9.2 shows a representative well-behaved residual: the field is balanced around zero and faint inclined ring structures are visible against the flat background. In practice, a fraction of pairs show a residual baseline shift toward positive or negative values (Figure 9.3). The factor k corrects for the mean throughput ratio of the two filter paths, but small additive offsets can persist, likely from spatially variable scattered light or short-timescale flux variations between exposures. This bias does not affect band detection, which operates on spatial frequency content rather than the absolute level of the residual, but it is worth noting as a limitation of the passband scaling in the paired-observation context.

After subtraction, a per-image 3σ outlier filter suppresses single-pixel artefacts by replacing pixels where $|I_{\text{res}} - \tilde{m}| > 3\sigma$ with the residual median \tilde{m} . The positions of flagged pixels are stored in a boolean mask; downstream steps consume this mask rather than the substituted values.

9.2 Download and format conversion

Calibrated ISS frames are distributed by the PDS in the VICAR (Video Image Communication And Retrieval) format, processed through the CISSCAL 4.0 pipeline






Figure 9.2: Well-behaved CLEAR-subtracted residual (IR3 filter, $\Delta t = 56$ s, orbit 163EN). The diverging colour map is centred at zero; red denotes science brighter than scaled CLEAR, blue the reverse. Faint inclined ring structures are visible across the field, and the bright red region at the lower limb traces emission from the Enceladus south-polar plume.



(a) BL1 filter ($\Delta t = 328$ s). The entire field shifts toward negative values: the scaled CLEAR overestimates the science-filter flux.



(b) RED filter ($\Delta t = 254$ s). The field shifts toward positive values: the scaled CLEAR underestimates the science-filter flux.

Figure 9.3: Examples of residual baseline bias from the same observation sequence. The bias reflects the limits of the passband scaling factor k and is distinct from the ring structures; it does not affect the frequency-domain detection.

(Knowles et al. 2020), which converts raw instrument counts to I/F . A VICAR file stores pixel data as a contiguous array of 32-bit floating-point values preceded by an ASCII metadata header and a block of binary label records. Standard PDS parsing libraries do not correctly resolve the pixel offset for CALIB products because they omit a non-zero binary-record contribution; a custom parser was therefore written to read the pixel array directly. Files already present on disk are skipped, making the download step safe to interrupt and restart.

9.3 Enceladus disc masking

The Enceladus disc is the brightest feature in most NAC frames and must be excluded from the residual image: the partially subtracted disc halo would otherwise dominate the image and mask any faint band structures. The disc centre and apparent radius in the NAC focal plane are computed for each exposure from the SPICE kernels, navigation and telemetry data products that reconstruct the geometric state of the Cassini spacecraft and solar system bodies at any point in time (C. H. Acton 1996; C. Acton et al. 2018). The disc radius in pixels and its centre column and row in the 1024×1024 NAC frame follow from the spacecraft-to-Enceladus range and the NAC plate scale.

A hard binary mask applied at the disc boundary would create a sharp step at the mask edge, producing ringing artefacts that spread radially across the image and could be confused with faint ring structure. Instead, a smooth multiplicative weight $w(x, y) \in [0, 1]$ is applied:

$$w(x, y) = \frac{1}{2} \left[1 - \cos \left(\pi \frac{d(x, y) - R_{\text{mask}}}{\Delta_{\text{taper}}} \right) \right] \quad \text{clipped to } [0, 1], \quad (9.3)$$

where $d(x, y)$ is the pixel distance from the disc centre, $R_{\text{mask}} = 1.5 R_{\text{disc}}$ is the inner mask boundary (set large enough to include the near-limb plume emission and any saturation halo), and $\Delta_{\text{taper}} = 60$ px is the taper width. The cosine shape is chosen because it provides a smooth, differentiable transition between fully suppressed and unaffected pixels, which eliminates the ringing artefact. Other smooth window functions such as the Hann or Tukey windows would achieve a similar result; the half-period cosine is the simplest available implementation. Pixels inside R_{mask} are fully suppressed ($w = 0$), those outside $R_{\text{mask}} + \Delta_{\text{taper}}$ are unaffected ($w = 1$), and the boundary region transitions smoothly between the two. Figure 9.4 shows the mask geometry for a representative frame.

9.4 Quality control

Before any pair enters the band-detection step, the full residual image set is reviewed visually. Automatic outlier removal cannot catch all failure modes: some pairs are numerically valid but scientifically unusable. The review identified five recurring categories of unusable pairs:

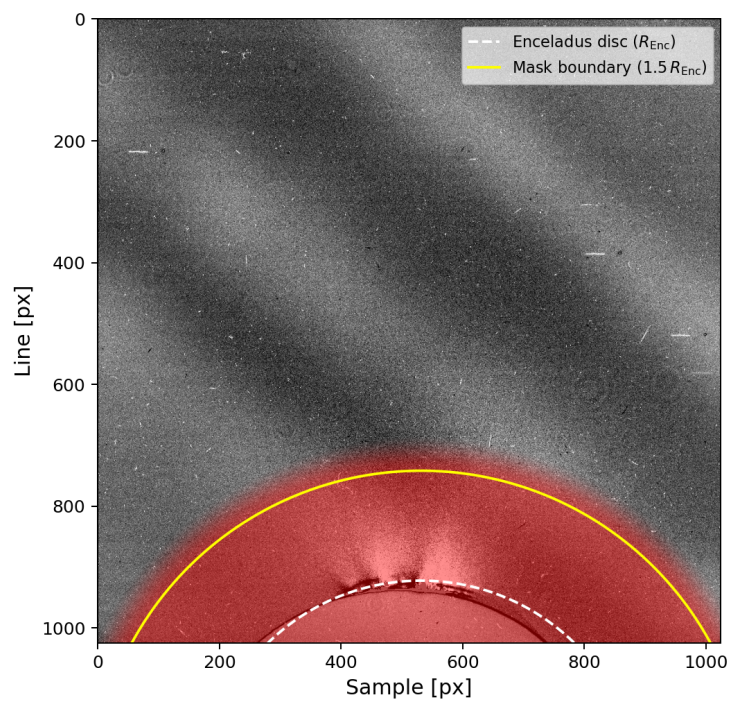


Figure 9.4: Enceladus disc mask for a representative frame (IR3, orbit 163EN). The white dashed circle marks the disc boundary (R_{disc}); the yellow solid circle marks the mask boundary at $1.5 R_{\text{disc}}$. In this frame Enceladus lies below the detector edge and only the lower portion of the mask intersects the image; the E ring band structures in the upper field are unaffected.

- Cosmic-ray residuals: bright trails or clusters that survive the 3σ clean, recognisable by their sharp, linear morphology.
- Saturation halos: large annular residuals around a saturated Enceladus disc, caused by differences in saturation extent between the CLEAR and science exposures.
- Failed reads: pairs returning a non-finite or identically-zero pixel array, indicating a corrupted or incompletely downloaded calibration file.
- Large time separations: pairs whose time gap between the CLEAR and science exposure is large enough for pointing drift to degrade the coherence of the subtraction.
- Disc-centred fields: pairs in which the field of view is centred on the Enceladus disc rather than the surrounding ring plane, making the residual unusable for E ring band analysis regardless of the calibration quality.

Of the 624 pairs for which a residual was produced, 51 were flagged and excluded through this review. The remaining 573 pairs constitute the formal survey input and enter the detection pipeline described in Chapter 10.

Chapter 10

Detecting and Characterising Luminous Bands

The data-preparation pipeline (Chapter 9) delivers 573 gray-level residual images: one per (CLEAR, science-filter) image pair, with the broadband E ring continuum subtracted, the Enceladus disc masked, and obvious artefacts removed. Each image captures a different observation of the ring from a different vantage point, at a different wavelength, at a different moment in the Cassini mission. The task of this chapter is to answer one question: is a luminous band present in this image, and if so, at what orientation and with what contrast?

The challenge is formidable. The bands are faint, with brightness modulations of a few percent of the local ring signal at best. They are inclined at arbitrary angles to the image frame. They are quasi-periodic, meaning they repeat with a characteristic spacing but not with perfect regularity. And the detection must be performed automatically across 573 images, without pre-selection or visual guidance.

This chapter describes how that detection is achieved. It begins not with the final algorithm but with the question that motivated it: given a gray-level residual image, how would one even approach detecting a faint, quasi-periodic stripe? The answer leads through a series of methodological steps: from a first attempt with a standard Fourier transform, through the discovery of its angular precision limitations, to the zero-padded angular-spectrum method and phase-randomised null ensemble that form the core of the pipeline. Figure 10.1 shows the complete detection chain.

10.1 Why frequency-domain detection?

10.1.1 Starting from a gray-level residual image

The input to the detection algorithm is a gray-level image: a 1024×1024 array of floating-point values, where each pixel records the local I/F residual after CLEAR subtraction. Positive values indicate that the scene is brighter in the science filter

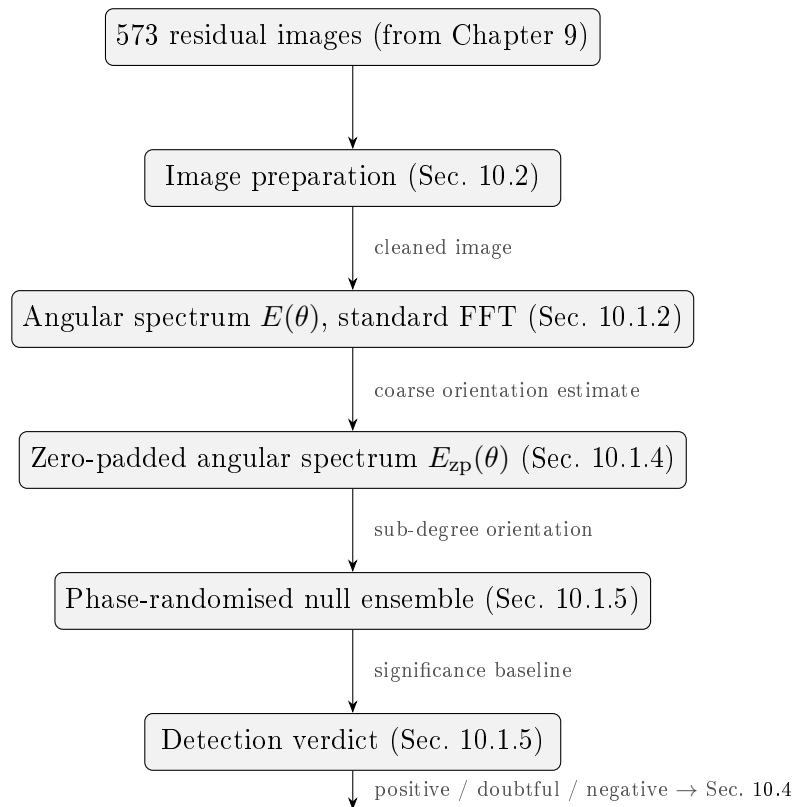


Figure 10.1: The detection chain applied to each of the 573 residual images. The conceptual foundation for each step is developed in Sections 10.1–10.1.5; the implementation parameters are collected in Section 10.3.

than in the scaled CLEAR; negative values indicate the reverse. A luminous band appears as a faint diagonal modulation: alternating light and dark stripes running at some angle across the field.

Finding this modulation is not straightforward. A range of alternatives was considered before settling on the approach described here, including edge detection, the Radon and Hough transforms, Gabor filter banks, wavelet decompositions, matched filtering, and neural network classifiers, among others. Each was set aside once the frequency-domain approach proved well-suited to the structure of the problem.

The property that makes a Fourier approach well-suited is this: *quasi-periodic directional structure concentrates its energy into a specific direction of the 2D frequency spectrum, regardless of where in the image the structure appears.* A band running at 40° to the horizontal leaves a signature at 130° in the Fourier domain (perpendicular to the band), irrespective of whether the band is in the upper-left corner or the lower-right of the image. This concentration property means that even a faint band (whose signal is buried in a single pixel) can be detected by integrating over all the pixels that contribute to the same Fourier direction. This is the key hypothesis that motivates the entire detection strategy.

10.1.2 The 2D Fourier transform and the angular spectrum

For a residual image $f(x, y)$ of size $M \times N$ pixels, the 2D Discrete Fourier Transform (DFT) is

$$F(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x, y) e^{-j2\pi(ux/M+vy/N)}, \quad (10.1)$$

where (x, y) are the spatial pixel coordinates and (u, v) are the spatial frequencies, in cycles per image width and height respectively. The DFT decomposes the image into a sum of 2D sinusoidal waves of all possible frequencies and orientations; the coefficient $F(u, v)$ measures how strongly each wave is present in the image. Because the band orientation is more important than its phase, the analysis uses the magnitude spectrum $|F(u, v)| = \sqrt{\text{Re}(F)^2 + \text{Im}(F)^2}$.

A periodic stripe pattern running at angle ϕ to the horizontal appears in the 2D magnitude spectrum as a pair of symmetric bright peaks oriented at $\phi + 90^\circ$: perpendicular to the spatial bands. The distance of these peaks from the DC component ($u = 0, v = 0$) is proportional to the spatial frequency of the band: closely spaced bands produce peaks far from the centre, while widely spaced bands produce peaks near the centre.

To reduce this 2D information to a 1D detection statistic, the magnitude spectrum is collapsed into an *angular power spectrum* $E(\theta)$: for each angle θ (in 0.25° bins), the power is accumulated by summing $|F(u, v)|/R$ over all frequency-domain pixels (u, v) whose angle falls within the bin and whose radius $R = \sqrt{u^2 + v^2}$ falls within the *accepted radial annulus* $[R_{\text{DC}}, R_{\text{max}}]$. The $1/R$ weighting equalises the contribution of different frequency magnitudes. The accepted radial annulus excludes the DC spike at the origin (below $R_{\text{DC}} = 8$ cycles/image) and high

frequencies dominated by pixel-scale noise (above $R_{\max} = 50$ cycles/image); only the intermediate range, where band signals are physically expected, contributes to $E(\theta)$. A peak in $E(\theta)$ at angle θ^* indicates that the dominant periodic structure in the image runs perpendicular to θ^* , i.e., at orientation $\theta^* + 90^\circ$.

Additionally, a $\pm 5^\circ$ wedge around the horizontal and vertical axes is excluded from the accumulation. This removes a systematic cross-shaped artefact visible in the 2D FFT: rows and columns of the detector each have a small systematic brightness offset that, after CLEAR subtraction and axis detrending, leaves a residual coherent along the detector axes. These systematic contributions would produce spurious peaks at 0° and 90° in $E(\theta)$; the axis exclusion suppresses them.

Figure 10.2 illustrates the principle with a pure sinusoidal stripe pattern: a single-frequency band at 135° produces a single sharp spike in $E(\theta)$ at 45° , perpendicular to the band, with no power at any other angle. This idealised case shows why the Fourier angular spectrum is well suited to detecting inclined periodic structure even in a noisy image: the signal from a band, however faint, is not spread across the spectrum but concentrated in a single direction. Figure 10.3 then shows the same pipeline applied to a real residual image from orbit 163EN (the E17 Enceladus flyby, 27 March 2012, IR3 filter). Panel (a) shows the residual image after the cleaning steps described in Section 10.2: a faint diagonal modulation runs from lower-left to upper-right, and the lower portion of the image is masked where Enceladus lies just below the detector edge. The speckle texture visible throughout the image is the granular E ring background, arising from individual fluctuations in photon noise and residual calibration errors. This background is isotropic (equally distributed in all directions) and therefore contributes uniformly to all bins of $E(\theta)$, rather than concentrating in one direction. Panel (b) shows the central region of the 2D FFT log-power: the bright cross at the horizontal and vertical axes (excluded by the wedge mask) is the detector-axis artefact, and the faint elongated concentration at approximately 135° is the directional signal from the band. Panel (c) shows the resulting $E(\theta)$: the highest peak sits near 135° , consistent with the directional concentration visible in panel (b). At this frequency resolution, however, the estimate is too coarsely localised to be reliable; zero-padding is required to sharpen it into a usable orientation measurement.

10.1.3 The precision problem: scalloping loss

For real noisy images, the standard FFT reveals a directional concentration in the 2D spectrum, but the resulting angular peak is dominated by noise and cannot be claimed as a band orientation. However, it has a fundamental angular precision limitation rooted in sampling: the frequency grid of a 1024×1024 FFT is fixed by the image dimensions, with adjacent grid points separated by exactly 1 cycle/image. When the magnitude spectrum is accumulated into 0.25° angular bins, the angular resolution is similarly limited: the reported peak angle can only take values at bin centres, and can therefore differ from the true band orientation by up to the half-bin width.




Figure 10.2: Synthetic demonstration of the Fourier detection principle. Panel (a): a pure sinusoidal brightness stripe at 135° , the geometric idealisation of a luminous band in a CLEAR-subtracted residual image. Panel (b): the corresponding angular power spectrum $E(\theta)$ computed from the image. A single sharp spike appears at $\theta^* = 45^\circ$ (red dashed), perpendicular to the band; the band orientation is recovered as $\theta^* + 90^\circ = 135^\circ$. This concentration of a periodic spatial structure into a single spectral direction is the core property that makes the Fourier angular spectrum the natural detection space for luminous bands.

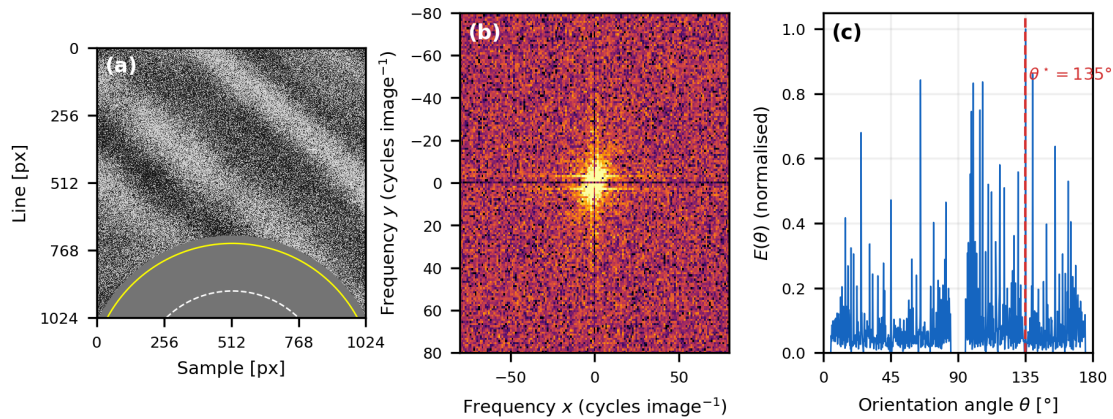


Figure 10.3: Standard Fourier detection for the running example (IR3 filter, orbit 163EN, 27 March 2012). (a) Cleaned residual image: a faint diagonal band runs from lower-left to upper-right; the flat grey region at the bottom is the Enceladus disc mask; the speckle texture is the isotropic E ring background. (b) Central region of the 2D FFT log-power; the bright cross at 0° and 90° is the detector-axis artefact, excluded from the analysis; the faint elongated concentration near 135° is the band signal. (c) Angular power spectrum $E(\theta)$; the highest peak sits near 135° (red dashed), but at this frequency resolution the estimate is too coarsely localised to be reliable. Zero-padding is required to sharpen this result.

More subtly, if the true Fourier peak of the band falls *between* two grid points, its energy is smeared across the neighbouring bins, a phenomenon known as *scallop loss*. The peak appears blunted and its height underestimated, making both the reported orientation and the reported signal strength unreliable. This is a sampling problem in the Fourier domain: the discrete grid undersamples the continuous spectrum, and the true peak position is hidden between grid points.

Figure 10.4 illustrates this with a synthetic 1D signal at a non-integer frequency of 18.7 cycles/image. The coarse FFT (red circles) samples the spectrum only at integer-spaced bins and misses the true peak entirely, reporting a lower amplitude and an offset location. This level of angular uncertainty, potentially several degrees, is acceptable for a coarse first detection but not for characterising band orientations or comparing them across observations.

10.1.4 Zero-padding: achieving sub-degree precision

The solution to the sampling problem is zero-padding: before computing the FFT, the residual image is surrounded by a border of zeros to produce a larger array of size $P \times 1024 \times P \times 1024$, where P is the padding factor. The DFT of this larger array evaluates the same underlying continuous Fourier spectrum but at a P -times finer frequency grid. This is mathematically equivalent to sinc-interpolation between the original coarse grid points: zero-padding does not add new spatial information, but it reveals where between the original bins the true spectral peak lies. A padding factor of $P = 16$ reduces the angular grid spacing from 0.25° to $0.25^\circ/16 \approx 0.016^\circ$,




Figure 10.4: Illustration of scalloping loss and its correction by zero-padding, using a synthetic 1D signal at a non-integer frequency (18.7 cycles/image). Red circles: coarse FFT sampled at integer-spaced bins, which misses the true peak and underestimates its amplitude. Blue curve: zero-padded FFT ($P = 16$), which interpolates the continuous spectrum at $16\times$ finer resolution and recovers the true peak location and amplitude. The gap between the two indicates the scalloping loss.

giving sub-degree angular precision on the band orientation. In practical terms, the difference between reporting a band at 40° and at 45° would place the same band on entirely different ring structures when projected back onto the sky plane; zero-padding is what makes the orientation measurement physically meaningful rather than merely directionally approximate.

The zero-padded angular spectrum $E_{\text{zp}}(\theta)$ is constructed by the same annular accumulation procedure as the standard FFT, applied to the finer-grid ZP magnitude spectrum over the same radial annulus where band signals are expected.

After accumulation, $E_{\text{zp}}(\theta)$ contains two contributions: the genuine directional signal from any band present in the image, and an *isotropic baseline* from the image's own background structure. The isotropic baseline is the angle-averaged power level: the amount of spectral power that would appear at each angular bin even if the image contained no directional structure at all. It arises from the granular texture of the E ring background and from residual systematic effects, and it must be measured and subtracted before any peak can be claimed as a detection. This is the purpose of the null ensemble.

10.1.5 Statistical significance: the null ensemble

To measure the isotropic baseline, 100 *phase-randomised surrogates* of the image are generated. Each surrogate is constructed by taking the complex Fourier transform of the image, replacing all Fourier phases with independent uniform random values in $[0, 2\pi)$, and transforming back to the spatial domain. This operation destroys

any directional signal in the image: the phase coherence that makes a band appear as a stripe is scrambled, while leaving the power spectrum $|F(u, v)|^2$ unchanged at every frequency. Each surrogate therefore has the same total power distribution as the original image, but no directional content.

The ZP angular spectrum is computed for each of the 100 surrogates, and the median, 16th percentile, and 84th percentile are recorded at each angular bin. These form the *null baseline*: the expected level of $E_{\text{zp}}(\theta)$ for an image with the same power spectrum but no directional structure. The null-corrected difference spectrum is

$$E_{\text{diff}}(\theta) = E_{\text{zp}}(\theta) - E_{\text{null,med}}(\theta), \quad (10.2)$$

where $E_{\text{null,med}}(\theta)$ is the null median. A peak in $E_{\text{diff}}(\theta)$ represents genuine directional power above the isotropic baseline.

To make peak detection robust, $E_{\text{diff}}(\theta)$ is smoothed at two scales: first at a narrow window of $w = 9$ bins (2.25°) to identify candidate peak locations, and then at a wide window of $w = 64$ bins (16°) to produce the final smooth spectrum from which peak heights are measured. The wide smoothing suppresses narrow spurious peaks while preserving the broad signature of a luminous band.

Figure 10.5 shows all three steps for the running example. Panel (a) shows the 2D ZP FFT log-magnitude: where the standard FFT (Figure 10.3b) showed only a diffuse diagonal hint, the zero-padded transform resolves a clearly elongated ridge, confirming that the directional signal is real and concentrated. The elongated ridge direction is taken as evidence that the signal corresponds to the band seen in the spatial image: it is oriented perpendicular to the spatial stripe and its position in the frequency domain is consistent with the band spacing observed in the image. Panel (b) shows $E_{\text{zp}}(\theta)$ against the null-ensemble median and 16th/84th percentile band: the peak near $\theta^* = 130^\circ$ stands well above the null baseline. Panel (c) shows $E_{\text{diff}}(\theta)$ after smoothing at $w = 9$ and $w = 64$ bins: the peak at 130° (band orientation 40°) is clearly resolved. The 5° offset from the standard-FFT result (45°) is a direct consequence of scalloping: the standard FFT reported the nearest bin centre, while the ZP result interpolates to the true peak location.

A positive detection requires two independent lines of evidence to agree:

1. A peak in the smoothed $E_{\text{diff}}(\theta)$ that clears the null 84th percentile.
2. A spatially visible band modulation in the cleaned residual image, oriented consistently with the spectral peak angle.

If both criteria are satisfied, the image is classified as a *positive detection* and the peak angle is recorded as the band orientation. If only one criterion is satisfied (a spectral peak without a visible spatial band, or a visible band without a clear spectral peak), the image is classified as *doubtful*. If neither criterion is met, the image is a *non-detection*. This two-criterion approach guards against both false positives from noise fluctuations and false positives from systematic image artefacts that produce spectral peaks without real spatial structure.




Figure 10.5: Zero-padded detection chain for the running example (IR3, orbit 163EN). (a) Log-magnitude of the 2D ZP FFT, cropped to the accepted radial annulus (axes in cycles image⁻¹); the elongated ridge near 130° is clearly resolved and its direction is overplotted (cyan line). Compare with the diffuse diagonal in the standard FFT (Figure 10.3b). (b) ZP angular spectrum $E_{zp}(\theta)$ (blue) against the null-ensemble median (dashed) and 16th/84th percentile band (grey shading). (c) Null-corrected difference spectrum $E_{diff}(\theta)$ smoothed at $w = 9$ bins (light) and $w = 64$ bins (dark); the peak at 130° (red dashed) gives a band orientation of 40° (orange dotted).

10.2 Image preparation

Before the angular spectrum is computed, each residual image is cleaned of signals that would contaminate the Fourier analysis. Four steps are applied in sequence, each targeting a specific known interference source.

The Enceladus disc mask, computed from SPICE-derived disc geometry as described in Chapter 9 (Section 9.3), is applied first. The SPICE computation accounts for light-time delay and stellar aberration: without these corrections, the finite travel time of light and the motion of the spacecraft would shift the computed disc centre by several pixels, leaving a residual halo at the wrong position. Without the mask itself, the partially subtracted disc halo, a broad circularly symmetric residual that survives CLEAR subtraction, would introduce strong power at all angles of $E(\theta)$, obscuring the band signal. The mask weight $w(x, y)$ (Equation 9.3 of Chapter 9) tapers smoothly from fully suppressed at the disc centre to fully unaffected at the outer edge of the transition zone. This smooth taper is critical: a hard step at the mask boundary would generate a spurious circular ridge in the 2D FFT that contaminates the angular spectrum at all orientations equally. Figure 10.6 demonstrates this effect: the hard boundary produces a clear circular structure in the 2D FFT (panel c), while the cosine taper removes it entirely (panel d).

A high-pass Gaussian filter ($\sigma = 80$ px) then removes the smooth, slowly-varying E ring brightness gradient that survives CLEAR subtraction. This gradient contributes power concentrated near the DC component of the Fourier spectrum; removing it prevents this low-frequency power from leaking into the angular bins at small radii.

Row-wise and column-wise median detrending subtracts the median value of each row and each column in turn. Individual detector rows and columns each carry a




Figure 10.6: Effect of the disc mask boundary on the 2D FFT, shown for the running example (N1711553432, IR3, orbit 163EN). Panels (a) and (b): the cleaned residual image after applying a hard binary mask and a cosine taper mask respectively; the yellow dashed circle marks the inner mask boundary. Panels (c) and (d): the corresponding 2D FFT log-magnitude, cropped to the region of interest around the DC component. The hard boundary introduces a circular ringing artefact in (c) that is absent in the cosine-taper result (d). This artefact adds power at all orientations equally, raising the noise floor of $E(\theta)$ and reducing sensitivity to band signals.

Table 10.1: Detection pipeline parameters. The rationale column explains the physical or algorithmic motivation for each value.

Step	Parameter	Value	Rationale
Disc mask	Radius factor	$1.5 \times R_{\text{disc}}$	Includes plume emission and saturation halo
	Taper width	60 px	Suppresses ringing; small vs ring scale
Image cleaning	HP Gaussian σ	80 px	Removes large-scale background; wider than typical band spacing
	Axis detrending	row and column median	Suppresses detector-stripe artefacts
	Arcsinh scale α	$0.05 \sigma_{\text{img}}$	Compresses bright outliers without distorting band signal
Angular spectrum	DC exclusion R_{DC}	8.0 cycles/image	Excludes DC spike and near-DC noise
	Axis exclusion	$\pm 5^\circ$ around $0^\circ, 90^\circ$	Removes detector-axis artefacts
	Outer cutoff R_{max}	50.0 cycles/image	Excludes pixel-scale noise
	Angular bins N_θ	720 ($0.25^\circ/\text{bin}$)	Matches ZP precision after smoothing
ZP FFT	Pad factor P	up to 16 (memory-limited)	Reduces angular grid to $\sim 0.016^\circ/\text{bin}$
	ZP inner radius	1.0 ZP-bin	Excludes the DC component in padded space
	ZP outer radius	6.0 cycles/image	Accepted annulus matching band spatial frequencies
Null ensemble	Surrogates N_{null}	100	Sufficient to estimate median and 84th percentile stably
Smoothing	Narrow window w_1	9 bins (2.25°)	Candidate peak localisation
	Wide window w_2	64 bins (16°)	Final peak detection; width tuned to band angular FWHM

small systematic offset from residual flat-field errors; after CLEAR subtraction these offsets appear as faint horizontal and vertical stripes in the residual image. In the Fourier domain they produce the cross artefact along the 0° and 90° axes seen in Figure 10.3b. Detrending suppresses these offsets before the FFT is computed, and the additional $\pm 5^\circ$ axis exclusion in the angular accumulation removes any residual contribution.

Finally, an arcsinh compression $f \rightarrow \sinh^{-1}(f/\alpha)$ with $\alpha = 0.05 \sigma_{\text{img}}$ is applied, where σ_{img} is the standard deviation of the valid pixels. This nonlinear rescaling compresses the dynamic range: bright pixels from surviving cosmic-ray residuals or plume emission would otherwise dominate the Fourier power and elevate the baseline across all angular bins. Arcsinh compression reduces their weight while preserving the relative contrast of the faint band modulation.

Saturn-intrusion variant. In ten images from one observation sequence, Saturn and its rings enter the NAC field of view. The Saturn disc contributes a strong directional signal coherent along the ring-plane direction that would dominate $E_{\text{diff}}(\theta)$ and mask any luminous-band signal. These images are processed with an additional rectangular mask covering the Saturn disc and ring edge, tapered with the same cosine window. All other steps are identical.

10.3 Pipeline parameters

Table 10.1 collects the full parameter set used for the survey, together with the rationale for each choice.

10.4 Survey results

The detection method described above was applied to all 573 residual images from the survey.

10.4.1 Detection geometry

Of the 573 images, 62 were classified as positive detections and 26 as doubtful. Figure 10.7 shows the viewing geometry of both groups projected onto Saturn’s equatorial plane. Each marker indicates the Enceladus position at the time of observation; the Sun direction is at the top. The positive detections (blue) and doubtful cases (orange) occupy the same broad geometric regime, at high phase angles and near-ring-plane viewing, and are concentrated on the evening-ansa side of the ring (the left quadrant). The absence of detections on the morning ansa despite comparable geometric coverage is one of the main findings of the survey and is discussed in the companion article (Chapter 7).

10.4.2 Example detections and the doubtful category

Figures 10.8–10.10 show the detection diagnostic for three representative cases: a positive detection, a doubtful case, and a non-detection. Each figure shows three panels: (a) the cleaned residual image with small line segments marking the detected band orientation; (b) the log-magnitude of the 2D ZP FFT, cropped to the accepted radial annulus; and (c) the null-corrected difference spectrum $E_{\text{diff}}(\theta)$ with the peak orientation and its full-width at half-maximum indicated. In all three cases $E_{\text{diff}}(\theta)$ is positive; what distinguishes them is whether one angle stands out clearly above the rest or whether the spectrum is flat with no clear winner.

The 26 doubtful cases arise from two distinct situations. In the first, a faint diffuse modulation is visible in the cleaned image but no corresponding peak appears clearly above the null in the ZP spectrum: the band contrast is too low for a confident spectral detection even if the spatial evidence is suggestive. In the second, the ZP spectrum shows a well-defined peak above the null but no band structure can be identified in the image, a spectral anomaly without spatial confirmation. In both situations the evidence from the two diagnostic views is inconsistent. A positive detection requires both to agree; when they do not, the image is set aside as doubtful rather than forced into either category. Resolving these marginal cases systematically, for example by developing a quantitative spatial coherence metric to complement the spectral significance criterion, is deferred to future work.

The pipeline is sensitive to bands whose contrast exceeds the systematic noise floor of the CLEAR subtraction, which is set by pointing drift and short-timescale flux variations between the two exposures. Bands with peak-to-trough contrast δ below approximately $5 \times 10^{-5} I/F$ are generally below this floor and would not produce a significant ZP spectral peak regardless of their spatial extent. Bands with a very large angular half-width (diffuse, slowly-varying modulations) spread their Fourier power across many angular bins rather than concentrating it in a peak, reducing the detection sensitivity for such structures. The pipeline is therefore most sensitive to bands that are moderately narrow in orientation angle and have contrast above the photometric noise floor.




Figure 10.7: Viewing geometry of positive (blue) and doubtful (orange) detections, projected onto Saturn’s equatorial plane. The Sun is at the top; the dashed lines separate the four ansa quadrants. Positive and doubtful detections occupy the same geometric regime, concentrated on the evening-ansa side. The 485 non-detections (not shown) have comparable geometric coverage, demonstrating that the detection asymmetry is physical rather than geometric.

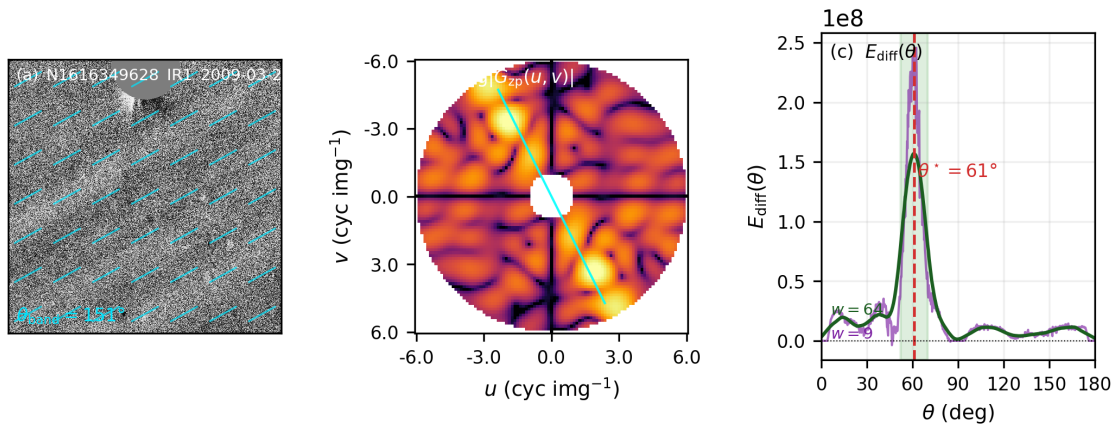


Figure 10.8: Detection diagnostic for a positive detection. Panel (a): a faint diagonal band is visible in the cleaned image; the overlaid segments confirm the recovered orientation. Panel (b): the 2D ZP FFT shows a clearly elongated ridge, concentrated at the band direction. Panel (c): $E_{\text{diff}}(\theta)$ has one angle that dominates clearly over all others, with a sharp, well-defined peak. Both detection criteria are satisfied.




Figure 10.9: Detection diagnostic for a doubtful case (BL1 filter, 20 November 2009). Panel (c): $E_{\text{diff}}(\theta)$ shows a somewhat prominent peak, but it is not sharply dominant over nearby angles. Panel (a): no coherent band structure can be identified in the cleaned image. Panel (b): no elongated ridge is visible in the 2D ZP FFT. The spectral and spatial evidence are inconsistent, placing this observation in the doubtful category.




Figure 10.10: Detection diagnostic for a non-detection. Panel (a): no band structure is visible in the cleaned image. Panel (b): the 2D ZP FFT shows no preferred orientation. Panel (c): $E_{\text{diff}}(\theta)$ is roughly flat across all angles, with no single direction standing out as a clear winner. Neither detection criterion is met.

10.4.3 Band contrast as a function of observing geometry

The band contrast $\delta = I_{\text{max}} - I_{\text{min}}$ is measured from a spatial brightness profile extracted perpendicular to the detected band orientation in the cleaned residual image. The pixels are projected onto the band-perpendicular axis, binned along that axis, and smoothed with a Savitzky–Golay filter to suppress pixel-scale noise. The contrast δ is then the peak-to-trough range of the smoothed profile, measured in I/F units. This procedure is applied to each of the 62 positive detections, for all available filters in each observation set.

As a consistency check, δ was also measured for ten randomly selected non-detections. In all ten cases the smoothed profile showed no statistically significant peak-to-trough variation, with values consistent with zero within the noise level. This confirms that the contrast measurement procedure does not produce spurious non-zero values in blank fields.

Figure 10.11 shows δ as a function of phase angle, Cassini–Enceladus distance, Cassini–Saturn distance, and E ring line-of-sight column for the four filters with four or more measurements (BL1, RED, IR1, IR3). A systematic decrease in δ with increasing Cassini–Enceladus and Cassini–Saturn distance is visible across all four filters. A moderate negative trend with phase angle is visible in the infrared filters. No clear trend appears with the line-of-sight column.

The line-of-sight (LOS) column through the E ring is computed under a simplified geometric model: the ring is treated as a flat torus with radial extent $r \in [180\,000, 480\,000]$ km and half-thickness $|z| \leq 2\,000$ km. For each image, the Cassini position and camera boresight direction from SPICE are used to trace the line of sight through the model torus, and the length of the segment inside the torus is recorded. Figure 10.12 shows the computed LOS paths for all survey images as a geometric verification.

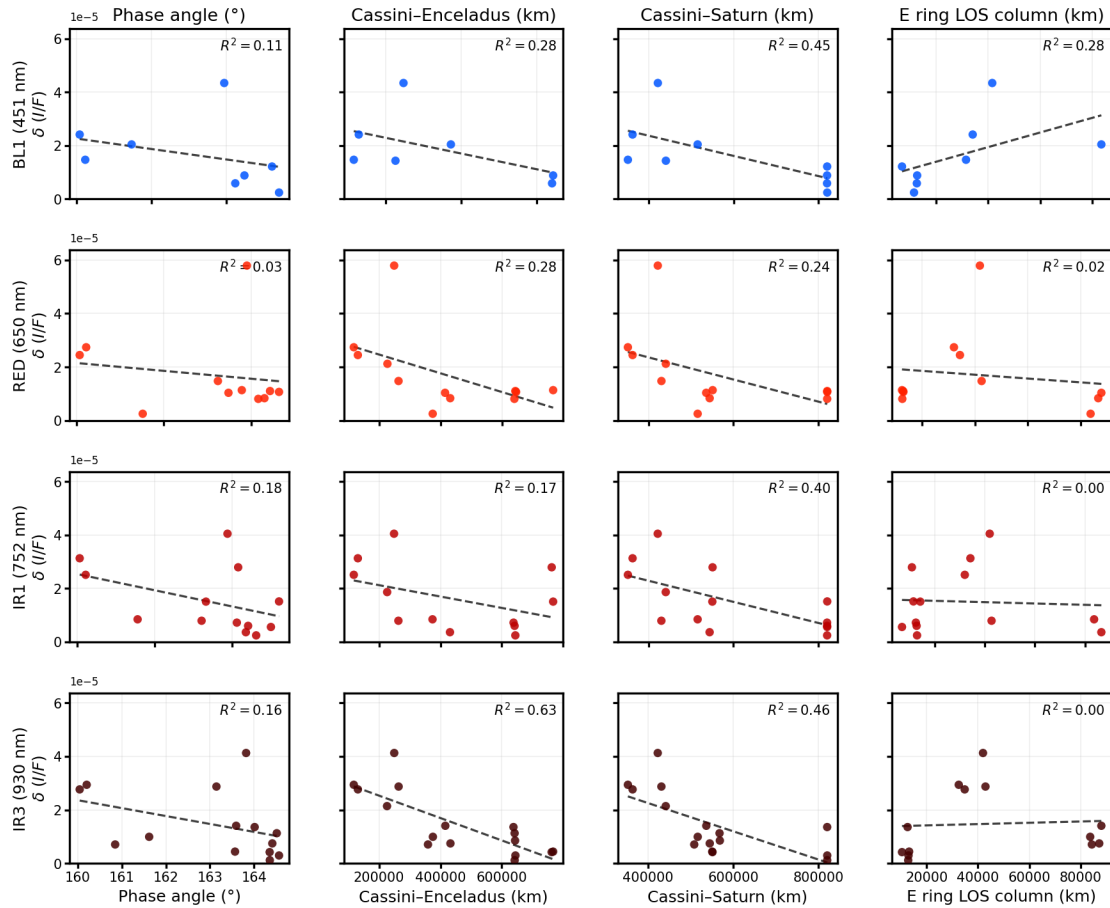


Figure 10.11: Band contrast δ (peak-to-trough in I/F) as a function of observing geometry for the four filters with four or more measurements. Rows: BL1, RED, IR1, IR3. Columns: phase angle; Cassini–Enceladus distance; Cassini–Saturn distance; integrated E ring line-of-sight column. Dashed lines are linear fits to guide the eye. The decrease of δ with distance is visible across all filters.




Figure 10.12: Geometric verification of the E ring line-of-sight column computation. Left: top-down view in the Saturn equatorial plane, showing Cassini positions (purple), Enceladus positions (blue), and LOS segments inside (orange) and outside (grey) the E ring model torus. Right: meridional cross-section confirming that the segments clip correctly to the model boundaries.

The contrast metric δ has two important limitations. First, it is measured in raw I/F units and is not corrected for the residual baseline bias discussed in Chapter 9: a filter-dependent DC offset can shift the profile up or down without affecting δ , but a sufficiently large offset can cause the profile minimum to clip near zero, artificially reducing δ . Second, the LOS model assumes a homogeneous, axisymmetric torus; in reality the E ring is denser near Enceladus and thinner at larger radii. A longer column through the torus does not necessarily correspond to more ring material if the additional path samples the low-density outer regions. Both effects mean that the trends in Figure 10.11 should be read as indicative rather than physically definitive, and that δ values from different filters or observation geometries are not directly comparable. A physically motivated contrast metric that corrects for the passband scaling bias and weights the LOS by a density model is a natural next step for future work.

Taken together, the three chapters of Part III answer the question of how the detection system was built: given the Cassini archive, how does one reliably detect luminous bands across 573 independent observations of the E ring? Chapter 7 recovered the paired observations from the archive; Chapter 8 turned them into calibrated residual images; Chapter 9 searched those images for directional periodic signals and assessed their significance. The result is a catalogue of 62 positive detections and 26 doubtful cases spanning 23 Cassini orbit sequences between 2007 and 2012. The scientific interpretation of this catalogue is presented in Part II.

References (Research Chapters)

- Acton, C., N. Bachman, B. Semenov, and E. Wright (2018). “A look towards the future in the handling of space science mission geometry”. In: *Planetary and Space Science* 150, pp. 9–12. DOI: [10.1016/j.pss.2017.02.013](https://doi.org/10.1016/j.pss.2017.02.013).
- Acton, C. H. (1996). “Ancillary data services of NASA’s Navigation and Ancillary Information Facility”. In: *Planetary and Space Science* 44.1, pp. 65–70. DOI: [10.1016/0032-0633\(95\)00107-7](https://doi.org/10.1016/0032-0633(95)00107-7).
- Knowles, B., R. West, P. Helfenstein, A. Verbiscer, D. Wilson, and C. Porco (2020). “End-of-mission calibration of the Cassini Imaging Science Subsystem”. In: *Planetary and Space Science* 185, p. 104898. DOI: [10.1016/j.pss.2020.104898](https://doi.org/10.1016/j.pss.2020.104898).
- Porco, C. C., R. A. West, S. W. Squyres, A. McEwen, P. Thomas, C. D. Murray, A. D. DelGenio, A. P. Ingersoll, T. V. Johnson, G. Neukum, et al. (2004). “Cassini Imaging Science: Instrument Characteristics and Anticipated Scientific Investigations at Saturn”. In: *Space Science Reviews* 115.1–4, pp. 363–497. DOI: [10.1007/s11214-004-1456-7](https://doi.org/10.1007/s11214-004-1456-7).
- Rubbrecht, N., S. Cazaux, B. Seignovert, M. Kenworthy, N. Kutsop, S. Le Mouélic, and J. Loicq (2025). “Peculiar rainbows in Saturn’s E ring: Uncovering luminous bands near Enceladus”. In: *Icarus* 441, p. 116650. DOI: [10.1016/j.icarus.2025.116650](https://doi.org/10.1016/j.icarus.2025.116650).