Spin-Scanning Cameras for Planetary Exploration: Imager Analysis and Simulation

George F. Brydon

A thesis submitted to **UCL** for the degree of
Doctor of Philosophy

Mullard Space Science Laboratory
Department of Space and Climate Physics
University College London

24th August 2021

I, George Brydon confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

In this thesis, a novel approach to spaceborne imaging is investigated, building upon the scan imaging technique in which camera motion is used to construct an image. This thesis investigates its use with wide-angle (≥90° field of view) optics mounted on spin stabilised probes for large-coverage imaging of planetary environments, and focusses on two instruments.

Firstly, a descent camera concept for a planetary penetrator. The imaging geometry of the instrument is analysed. Image resolution is highest at the penetrator's nadir and lowest at the horizon, whilst any point on the surface is imaged with highest possible resolution when the camera's altitude is equal to that point's radius from nadir. Image simulation is used to demonstrate the camera's images and investigate analysis techniques. A study of stereophotogrammetric measurement of surface topography using pairs of descent images is conducted. Measurement accuracies and optimum stereo geometries are presented.

Secondly, the thesis investigates the EnVisS (Entire Visible Sky) instrument, under development for the Comet Interceptor mission. The camera's imaging geometry, coverage and exposure times are calculated, and used to model the expected signal and noise in EnVisS observations. It is found that the camera's images will suffer from low signal, and four methods for mitigating this – binning, coaddition, time-delay integration and repeat sampling – are investigated and described. Use of these methods will be essential if images of sufficient signal are to be acquired, particularly for conducting polarimetry, the performance of which is modelled using Monte Carlo simulation.

Methods of simulating planetary cameras' images are developed to facilitate the study of both cameras. These methods enable the accurate simulation of planetary surfaces and cometary atmospheres, are based on Python libraries commonly used in planetary science, and are intended to be readily modified and expanded for facilitating the study of a variety of planetary cameras.

Impact Statement

Much of the work of chapter 5 is published in the peer reviewed journal *Planetary and Space Science*. Its content presents new techniques for both planetary observation and instrument study which aim to inform and be utilised in future planetary science missions.

The work of chapter 6 contributed significantly to the development of the Comet Interceptor mission's EnVisS instrument, assisting with its study prior to the mission's selection by the European Space Agency, and with its subsequent development. At the time of writing, the instrument is expected to be part of the mission's payload when it launches in 2029, and to ultimately return images that will enable novel and valuable cometary science to be conducted.

Beyond their academic application, the image simulation techniques developed for this thesis have facilitated the communication of its contents to the wider public. The simulations produced have been used in publications, conference presentations, outreach activities and broadcast media.

Acknowledgements

This work is dedicated to Vera Brydon. It would take more pages than are in this thesis to describe the impact she had throughout my life. Her continual and unconditional encouragement and support made this and many other things possible, where they otherwise wouldn't have been. She is greatly missed.

I am extremely grateful to my parents whose support, encouragement and advice have been invaluable throughout my education. They likely weren't expecting to house me for the last year of my PhD, but under strange circumstances (a putting mat on the hallway floor for 12 months straight) they did it with great patience, and made it a very enjoyable time. Thank you also to my brothers Josh and Jack, and every other family member who has played a part in my getting here.

I must also thank Divya, who every day is (and for the last three years has been) a source of inspiration, support, perspective and humo(u)r, and has opened my eyes to many things. Divya, I would not be who I am without you.

I would, of course, also like to thank my supervisors Geraint Jones and Andrew Coates. Geraint's support, advice and feedback throughout my PhD research have been invaluable and much appreciated. So too has his generosity in providing opportunities for me to get involved with the Comet Interceptor mission, something from which I and this PhD benefited significantly. Both Geraint and Andrew, along with Ollie Price, Roger Stabbins, Will Dunn, Annie Wellbrock, Eleni Bohacek, Yutian Cao, Sam Taylor, Affelia Wibisono and Richard Haythornthwaite made the Planetary group at MSSL an excellent and enriching place to work. Numerous other current and former members of MSSL provided good company and welcome distractions throughout my PhD.

Contents

List of Figures

List of Tables

1 Introduction

Cameras, in one form or another, are a fundamental part of spaceborne planetary science missions. This has been the case practically since the very first planetary camera went to space in 1959. A continuous improvement in the techniques and technologies involved in planetary imaging has led to observations covering an ever-growing number of solar system bodies in ever-greater detail, and has facilitated impressive observations such as large mountain ranges shrouded in a thin atmosphere at Pluto (Figure 1.1 left), and images captured from the surface of a comet (Figure 1.1 right). Some argue that cameras are the most productive of planetary science instruments (Malin, et al., 2007), and indeed they historically often lead the scientific payload of planetary missions (Anderson, et al., 1991), but their value, to both scientists and the wider public, comes also from their unique ability to convey to us what our own eyes would see, were we there ourselves.

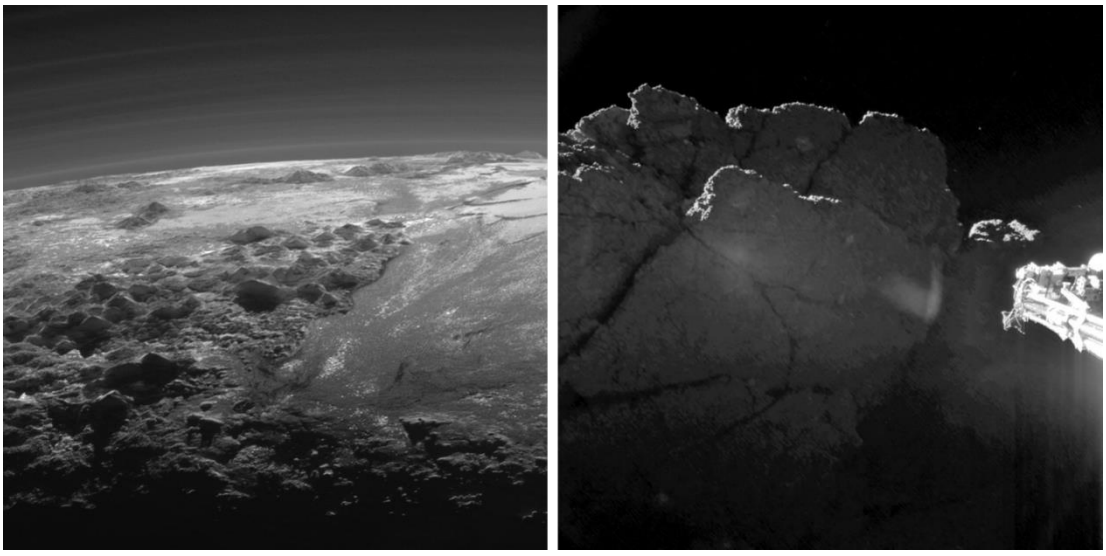

Figure 1.1: (left) Mountain ranges on Pluto's surface, and a thin atmosphere, captured by the New Horizons spacecraft during a 2015 flyby (Stern, 2018). (right) A fractured cliff on the surface of comet 67P/Churyumov-Gerasimenko, captured by the Philae lander of the Rosetta mission, reproduced from Bibring et al. (2015).

If current and recent in-development missions, technology demonstrations and spacecraft concepts are indicative of things to come, the future of planetary exploration is likely to present increasingly varied and diverse opportunities by which to image

planetary bodies. Successful tests of Ingenuity, the first ever Martian helicopter (Balaram, et al., 2021), may pave the way for flying cameras at a number of solar system bodies (see Dragonfly, for example (Lorenz, et al., 2021)). Mission designs such as the L-DART penetrator (Barber, et al., 2018) and the DAVINCI+ Venus probe (Garvin, et al., 2020) are examples of small planetary probes that can facilitate planetary imaging in a variety of scenarios. And the tried and tested orbiter and flyby spacecraft designs that have captured thousands of images over the last six decades will continue to be represented in future missions, transporting cameras throughout the solar system (e.g. JUICE (Grasset, et al., 2013), EnVision (Widemann, et al., 2020)).

Imaging diverse solar system bodies from a wide variety of spaceborne platforms necessitates (and indeed facilitates) an equally wide variety of cameras. Fields of view range from <1° to >90°. Some cameras are precisely pointed at their target, some scan across the scene, and some image whilst they and their spacecraft spin. Planetary cameras can be as large as 1 m long with masses over 50 kg, or as small as 10 cm with mass less than 0.5 kg. Images may be captured from a static platform (e.g. a lander), or whilst travelling hundreds to thousands of metres per second relative to the target. This wide variety of imaging setups is often dictated by the nature of the spacecraft on which the camera is mounted, and the target that it will image. A camera ideal for one imaging scenario can be useless in another. Studying and developing new methods of imaging therefore broadens the range of missions from which spaceborne imaging can be performed. This can facilitate the inclusion of a camera on a spacecraft which would otherwise not have one, helping to strengthen the scientific value and justification of the mission.

1.1 Research Aims

In this thesis, a novel approach to planetary imaging is investigated, with the aim of providing the foundations for the method's further development for, and use on, future missions. It should be noted that this thesis deals exclusively with imaging in the visible and near-visible range of the spectrum. The thesis builds upon the scan imaging technique, used by spaceborne cameras for five decades, in which a scanning motion of the camera is used to construct images. The focus here is a novel implementation of this method, in which rotational scan motion is provided by a spinning spacecraft, and fisheye optics are used to achieve complete coverage of the target. The scanning of the camera's view across the scene allows the construction of large-coverage images, and permits multispectral imaging without filter mechanisms. In this sense, the imaging

approach aims to be suited to small planetary probes, where mass and complexity must be minimised and the probe is more likely to be spin stabilised.

The work of this thesis focusses on two instruments. Firstly, a descent camera for a planetary penetrator. This is an instrument concept, unattached to any specific mission. It images from onboard a spinning penetrator during its descent to a planetary surface. The combination of 360° rotation and wide-angle optics allows the camera to image all of the visible surface. This thesis presents calculations of the camera's imaging geometry, its resolutions and coverage, and their evolution over the course of the probe's descent. Software is developed to simulate the camera's images, which is used for discussion and demonstration of image processing including geometric correction and rectification. The simulated images are used in a comprehensive study of descent stereophotogrammetry, demonstrating and characterising the camera's ability to measure surface topography.

Secondly, the Entire Visible Sky (EnVisS) imager is addressed. This instrument is under development for the European Space Agency's (ESA) Comet Interceptor mission. The work of this thesis began in the early stages of the instrument, prior to the mission's selection by ESA, and contributed to its subsequent development. EnVisS is designed to image during a close comet flyby, from a small spin stabilised probe. Operating similarly to the penetrator concept, but with even wider-angle optics, it is designed to capture images of almost the entire sky in order to study the comet's large-scale structures. The camera's imaging geometry and mechanics are calculated. The image simulation software is expanded to allow exemplar EnVisS data to be produced. A noise model of the instrument is developed and used for comprehensive study of its expected signals and capabilities, and the camera's novel method of imaging polarimetry is characterised.

1.2 Thesis Overview

This introductory chapter is followed by three background chapters. Chapter 2 gives an overview of imaging, presenting the maths necessary for describing cameras, and introducing terminology that will be used throughout this thesis. Techniques and hardware used for spaceborne imaging are subsequently presented.

With this foundation set, chapter 3 gives a history of spaceborne imaging throughout the solar system, demonstrating the diverse ways in which planetary bodies have been observed, and attempting to justify the development of new imaging methods.

Chapter 4 provides background on comets and their observation. An overview of cometary science is given, and the current scientific model of comets is described. Spacecraft encounters with comets are outlined, followed by descriptions of the cameras used by those spacecraft to image their target comets.

The subsequent three chapters present this thesis' results. In chapter 5, the study of the penetrator descent camera is presented, in which its imaging geometry is derived, its data are simulated, image processing steps are developed and the use of its images for descent stereophotogrammetry is demonstrated. The stereophotogrammetry work of chapter 5, along with the background material of chapter 3 section 3.2, is also published as a peer-reviewed article in the journal Planetary and Space Science (Brydon, et al., 2021).

Chapter 6 presents work conducted on the EnVisS instrument. The imaging geometry, coverage, exposure times and exposure frequency are derived. A study of previous cometary observations and cometary dust models is conducted to predict the spectral radiances that EnVisS will observe. These are combined with an instrument model to predict image signals, and a number of techniques for increasing image signal are characterised. Finally, a Monte Carlo simulation of the camera's imaging polarimetry is presented, and its implications on focal plane design, measurement accuracies, and image signals are discussed.

In chapter 7, the methodology that was developed for simulating penetrator descent camera and EnVisS images is presented. The overarching goals of the simulation approach are described. Simulation of planetary surfaces is then discussed and demonstrated. Following this, the physics governing radiative transfer is presented, and a method of physically simulating images of cometary dust structures is thus derived. The thesis ends with chapter 8, where key conclusions are discussed, and opportunities for further research are identified.

2 Scientific Imaging Techniques

2.1 The Camera Model

The function of a camera is to record a two-dimensional (2D) representation (an image) of a three-dimensional (3D) scene (the world). Fundamentally, this requires the mapping of points in 3D space on to a plane, called the focal plane, where their position is described by a 2D coordinate:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} \rightarrow \begin{bmatrix} u \\ v \end{bmatrix} \tag{2.1}$$

The relationship between a point's position in 3D space and its 2D projection on the focal plane is known as the imaging geometry of the camera. The imaging geometry is the foundation of mathematically describing a camera's images, and as is the goal in scientific imaging, using the camera's images to describe the scene it observed. The specific imaging geometry of a camera depends on the nature of its optics (the apertures, lenses and mirrors at the 'front' of the camera which direct light to the focal plane) whose purpose is to perform the mapping of equation 2.1.

2.1.1 The Pinhole Camera

Imaging and images are actually commonplace in nature, and have been exploited by humans for millennia. The surfaces of bodies of water and shiny materials have been used as mirrors for over 8000 years (Enoch, 2006). Images of opaque objects appear as shadows when the Sun shines, and have been used for measuring the time since at least as early as 1200 BC (Pilcher, 1923). And small holes – occurring in the canopies of trees, for example – project images of the Sun on to the ground, a phenomenon about which Aristotle wrote (Lindberg, 1968), and which is now called pinhole imaging.

The first use of this phenomenon to construct a camera, known as a pinhole camera, is generally attributed to Ibn Al-Haytham, who lived circa 1000 and is often referred to as the father of modern optics (Lindberg, 1968; Young, 1989). The pinhole camera is the simplest of camera designs. In its ideal form, rays from points in 3D space are only permitted to reach the focal plane if they happen to pass through a single point known as the pinhole (in more general camera terminology, this point, the nexus of rays travelling through the camera to the focal plane, is known as the optical centre). These

rays travel in a straight line through the optical centre onto the focal plane, thus performing mapping consistent with equation 2.1. The imaging geometry of the pinhole camera is generally referred to as either the pinhole or central perspective geometry. Real pinhole cameras are physically constructed using an opaque, closed chamber with a small hole in one face, which projects an image onto the inside of the opposite face, as illustrated in Figure 2.1. Note that this geometry leads to an inverted image.

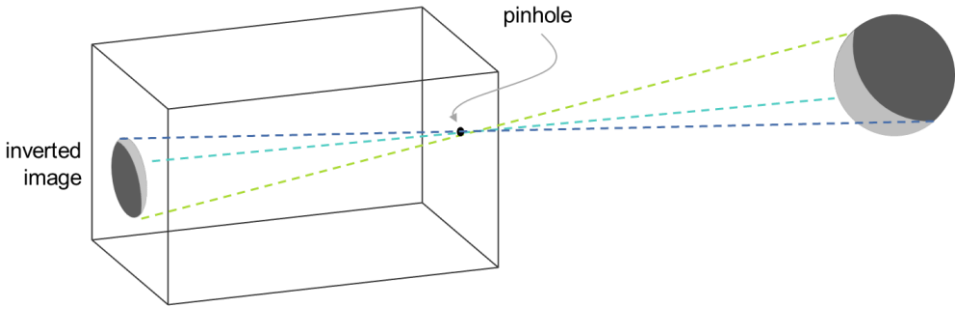

Figure 2.1: A simple pinhole camera. The pin hole is the only way for light to access the inside of the camera. Rays (dashed lines) hit the back face of the camera (focal plane) and form an inverted image.

Because a real pinhole camera produces an inverted image, the convention when mathematically describing the pinhole imaging geometry is to place the focal plane in front of the optical centre, so that light rays strike it *before* passing through the optical centre. It is of course not physically possible to construct a pinhole camera in this way, but it is a geometrically identical model to the real case, and brings the simplification of a non-inverted image. This form of the pinhole camera model is illustrated in Figure 2.2, which shows the relationship between points in 3D space and their projection on the focal plane. The origin of the focal plane's coordinate system is called the principal point. The line connecting the principal point to the optical centre (the origin of the camera's 3D coordinate system) is called the principal axis. The principal axis is perpendicular to the focal plane. The distance from the optical centre to the principal point is defined as the focal length (in a real pinhole camera, as in Figure 2.1, this is the distance from the front face to the back face of the camera).

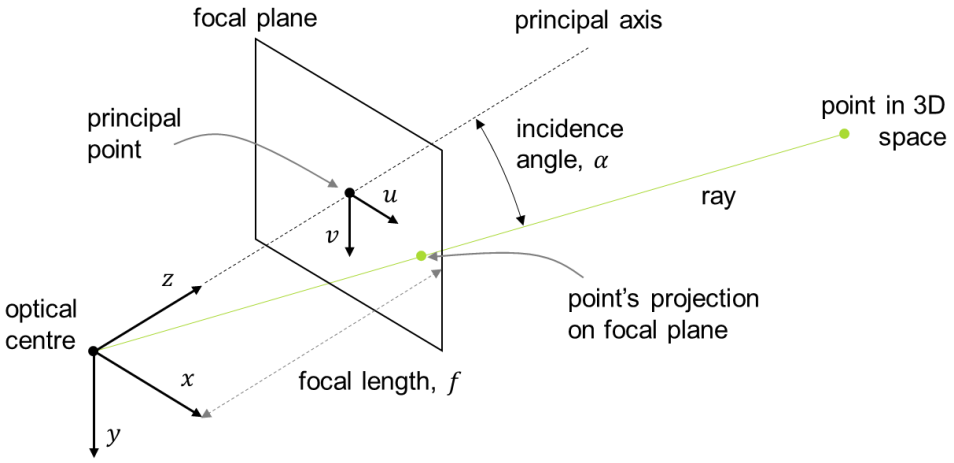

Figure 2.2: Pinhole camera model, with focal plane in front of the optical centre. The camera's 3D coordinate frame $[x \quad y \quad z]^T$ and focal plane's 2D coordinate frame $[u \quad v]^T$ are indicated by the solid black arrows. The green ray indicates the mapping of a 3D point to the focal plane.

Mapping of a 3D point $[x \quad y \quad z]^T$ to its 2D image coordinate $[u \quad v]^T$ is given by

$$\begin{bmatrix} u \\ v \end{bmatrix} = \frac{f}{z} \begin{bmatrix} x \\ y \end{bmatrix} \tag{2.2}$$

where f is the camera's focal length. A single point in 2D image space corresponds to a line in 3D space of constant x/z and y/z. The mapping of equation 2.1 carried out by a camera is not bijective, i.e. a 3D coordinate cannot be recovered from its mapped coordinate on the focal plane. Instead, a point on the focal plane maps to a direction (i.e. a straight line in space emerging from the optical centre), represented by a vector of arbitrary length:

$$\vec{d} = \frac{1}{f} \begin{bmatrix} u \\ v \\ f \end{bmatrix} \tag{2.3}$$

Because, in the pinhole geometry, the ray which maps a real world point to the focal plane is simply a straight line through the optical centre, the angle it makes with the principal axis before reaching the optical centre (called the incidence angle, α) is the same as the angle after passing through the optical centre (called the reflection angle, β). The incidence angle of a ray passing through a 3D point $[x \quad y \quad z]^T$ is described by

$$\tan \alpha = \frac{\sqrt{x^2 + y^2}}{z} \tag{2.4}$$

Let the image radius, r_{im} be the distance from the principal point to a projected point on the focal plane. For a pinhole camera, image radius and incidence angle are related by

$$r_{im} = f \tan \alpha \tag{2.5}$$

The relationship between incidence angle and image radius is called the mapping function (the specific relationship of equation 2.5 is the central perspective mapping function), and it is a defining characteristic of a camera's imaging geometry.

Through equation 2.5, the physical size of the focal plane dictates the extent of the angular range imaged by the camera, known as its field of view (FOV). FOV is generally split into horizontal and vertical components, which are angles measured in the x-z and y-z planes respectively (meaning they align with the u and v axes of the image respectively). A focal plane centred on the principal point with length U and V in the u and v dimensions respectively will capture a horizontal (FOV_u) and vertical (FOV_v) FOV given by

$$FOV_u = 2 \tan^{-1} \frac{U}{2f}$$

$$FOV_v = 2 \tan^{-1} \frac{V}{2f} \tag{2.6}$$

The dimensions U and V of a camera's focal plane are dictated by the physical size of its sensor. The focal length f of the camera's optics then sets its FOV, with longer focal lengths resulting in narrower FOVs. If one required a large FOV (70°, for example), a very small focal-length-to-sensor-size ratio would be required. Realising this in an actual camera requires either a very large sensor, or optics which are positioned very close to the sensor, both of which are physically challenging to achieve. For this reason, the central perspective projection is generally used only for cameras with narrow FOVs, and cameras with large FOVs (called wide-angle cameras) rely on different mapping functions.

2.1.2 Non-perspective Mapping Functions

Most cameras' imaging geometries are that of central perspective (Schneider, et al., 2009), and it is also a good approximation of human vision (Artal, 2015). For that reason, the central perspective model is frequently used to describe cameras' imaging geometries. But, when constructing real cameras, the scope of the central perspective geometry suffers from a limitation: as the incidence angle grows towards 90°, the image radius approaches infinity. As a result, a central perspective camera designed to image

large incidence angles would require a physically large detector that is both expensive to construct and challenging to accommodate.

In practice, large incidence angles are instead imaged with fisheye optics (so named because they mimic the view seen by a fish looking up through the water's surface), whose mapping functions accommodate wide angles on small detectors. Images formed by fisheye optics are notable for their distortion, whereby straight lines appear curved (Schneider, et al., 2009). Unlike for central perspective geometry, incidence angle and reflection angle are not equal in fisheye optics (i.e. $\alpha \neq \beta$) as illustrated in Figure 2.3. Table 2.1 lists the mapping functions of four common fisheye geometries, alongside central perspective.

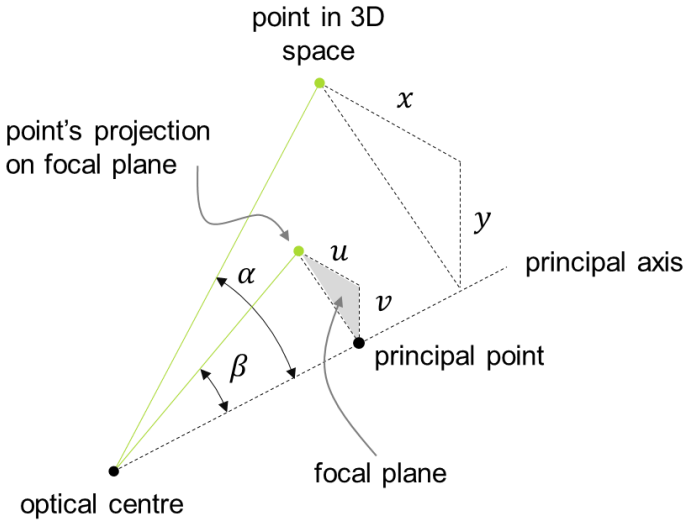

Figure 2.3: Camera model for an imaging geometry with arbitrary mapping function.

Table 2.1: Mapping functions of different imaging geometries (Schneider, et al., 2009).

Projection Type	Mapping Function
Central perspective	$r_{im} \propto \tan \alpha$
Equidistant	$r_{im} \propto \alpha$
Equisolid angle	$r_{im} \propto \sin {}^{\alpha}/_{2}$
Orthographic	$r_{im} \propto \sin \alpha$
Stereographic	$r_{im} \propto \tan {}^{\alpha}/_{2}$

To obtain a general relationship between a 3D point in space and its 2D position on the focal plane, note from Figure 2.3 that

$$\frac{u}{v} = \frac{x}{y} \tag{2.7}$$

because the scene point, image point and principal axis define a single plane. For any imaging geometry, image radius is given by

$$r_{im} = \sqrt{u^2 + v^2} \tag{2.8}$$

and equations 2.7 and 2.8 lead to the result

$$u = \frac{r_{im}}{\sqrt{\left(\frac{y}{x}\right)^2 + 1}}$$

$$v = \frac{r_{im}}{\sqrt{\left(\frac{x}{y}\right)^2 + 1}} \tag{2.9}$$

One of the most commonly used mapping functions in the design and manufacture of fisheye optics is the equidistant projection:

$$r_{im} = f'\alpha \tag{2.10}$$

Indeed, the cameras studied in this thesis (chapters 5 and 6) are both assumed to use optics with an equidistant mapping function. In equation 2.10, the constant f' is the distance from optical centre to principal point, but in a fisheye system this is not technically the focal length (strictly speaking, the focal length of fisheye optics varies with incidence angle). From equations 2.4 and 2.10, the image radius of a point $[x \quad y \quad z]^T$ in equidistant geometry is given by

$$r_{im} = f' \tan^{-1}\left(\frac{\sqrt{x^2 + y^2}}{z}\right) \tag{2.11}$$

By extension, equations 2.9 and 2.11 yield the mapping of 3D coordinates to 2D focal plane coordinates:

$$u = \frac{f' \tan^{-1}\left(\frac{\sqrt{x^2+y^2}}{z}\right)}{\sqrt{\left(\frac{y}{x}\right)^2 + 1}}$$

$$v = \frac{f' \tan^{-1}\left(\frac{\sqrt{x^2+y^2}}{z}\right)}{\sqrt{\left(\frac{x}{y}\right)^2 + 1}} \tag{2.12}$$

For mapping focal plane coordinates to 3D space (recall that this yields a direction vector, not a point), combining equations 2.8 and 2.11 gives the relationship

$$z = \frac{\sqrt{1 + \frac{v^2}{u^2}}}{\tan\left(\frac{\sqrt{u^2+v^2}}{f'}\right)} \tag{2.13}$$

The mapping of $[u \quad v]^T$ to a direction vector $\vec{d}$ of arbitrary length is then found from equations 2.7 and 2.13:

$$\vec{d} = \begin{bmatrix} 1 & \frac{v}{u} & \frac{\sqrt{1 + \frac{v^2}{u^2}}}{\tan\left(\frac{\sqrt{u^2+v^2}}{f'}\right)} \end{bmatrix}^T \tag{2.14}$$

Equations 2.12 and 2.14 describe the inward and outward projection respectively of a camera with equidistant imaging geometry. The FOV of a focal plane with lengths U and V in the u and v dimensions, centred on the principal point, is given by

$$FOV_u = \frac{U}{f'}$$

$$FOV_v = \frac{V}{f'} \tag{2.15}$$

The inward and outward projections and FOVs of any of Table 2.1's mapping functions can be found by this same method.

2.1.3 Coordinate Transforms

In order to increase the usability of the imaging geometry equations derived in sections 2.1.1 and 2.1.2, it is necessary to be able to describe transformations between the camera's intrinsic coordinate frames (both the 3D camera frame and the 2D focal plane frame), and other, extrinsic coordinate frames.

2.1.3.1 *Image Coordinates*

By convention, the origin of an image's coordinate frame is not placed at the principal point of the camera, but instead is offset to one of the image's corners (commonly the top left corner), or its centre (Figure 2.4). Coordinates $[u \quad v]^T$ in the focal plane frame transform to coordinates $[u' \quad v']^T$ in an offset image frame by the equation

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \begin{bmatrix} u - \delta_u \\ v - \delta_v \end{bmatrix} \tag{2.16}$$

where δ_u and δ_v are the horizontal and vertical positions respectively of the offset frame's origin in the focal plane coordinate frame. In an actual camera, the value of this offset will depend on the shape and size of the camera's sensor, and its positioning relative to the optics' principal point.

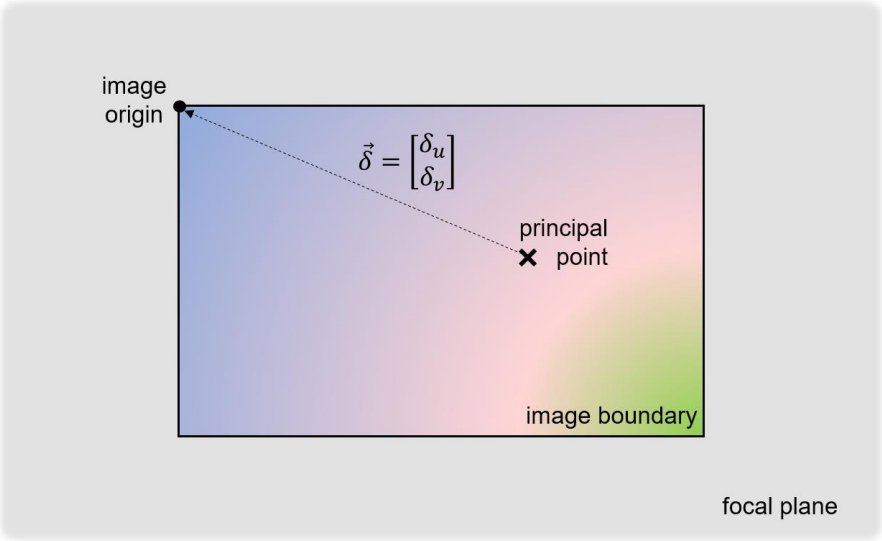

Figure 2.4: Illustration of an image on a camera's focal plane, and the relationship between focal plane coordinates and image coordinates (equation 2.16).

2.1.3.2 *Converting to and from the Camera Frame*

In the imaging geometry equations of sections 2.1.1 and 2.1.2, the camera's intrinsic coordinate frame is used to define 3D positions. It is often much more useful for 3D coordinates to be linked to a separate coordinate frame, such as that of a planet, and it is therefore necessary to be able to describe transformations between 3D Cartesian frames. Consider two such coordinate frames A and B with bases $\{\vec{a}_1, \vec{a}_2, \vec{a}_3\}$ and $\{\vec{b}_1, \vec{b}_2, \vec{b}_3\}$ respectively. Frames A and B are described by matrices M_a and M_b respectively:

$$M_a = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix}$$

$$M_b = \begin{pmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{pmatrix} \tag{2.17}$$

A point with coordinate $\vec{p}_a$ in frame A has a coordinate $\vec{p}_b$ in frame B given by

$$\vec{p}_b = {M_b}^T(\vec{O}_a - \vec{O}_b + M_a\vec{p}_a) \tag{2.18}$$

where $\vec{O}_a$ and $\vec{O}_b$ are the locations of the two frame's origins (see Figure 2.5). Consider a world coordinate frame with base vectors $e_1 = [1 \quad 0 \quad 0]^T$, $e_2 = [0 \quad 1 \quad 0]^T$ and $e_3 = [0 \quad 0 \quad 1]^T$, in which a camera's frame is defined (M_c, $\vec{O}_c$). World points $\vec{p}_w$ are converted to the camera frame by the transformation

$$\vec{p}_c = {M_c}^T(\vec{p}_w - \vec{O}_c) \tag{2.19}$$

and points in the camera frame are transformed to the world frame by

$$\vec{p}_w = M_c\,\vec{p}_c + \vec{O}_c \tag{2.20}$$

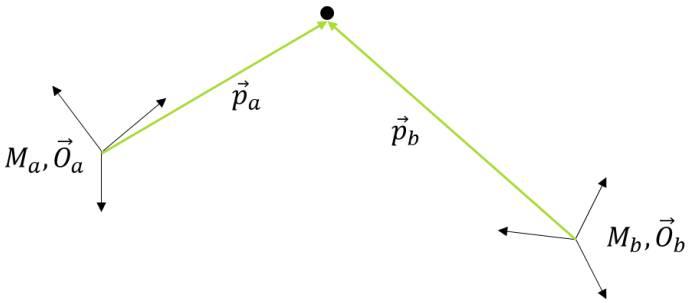

Figure 2.5: Illustration of two coordinate frames, as described in the main text.

2.1.3.3 Spherical Coordinates

Finally, a coordinate transform that will be very useful for the cameras studied in this thesis is the mapping from 3D Cartesian coordinates to spherical coordinates $[r \quad \theta \quad \phi]^T$:

$$\begin{bmatrix} r \\ \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \sqrt{x^2 + y^2 + z^2} \\ \tan^{-1}\left(\frac{y}{x}\right) \\ \cos^{-1}\left(\frac{z}{\sqrt{x^2 + y^2 + z^2}}\right) \end{bmatrix} \tag{2.21}$$

where r is the point's distance from the coordinate frame origin, azimuth θ is the angle measured anticlockwise around the z axis, from the x axis, to the point's projection in the x-y plane, and the polar angle ϕ is measured from the z axis to the point's radius vector (illustrated in Figure 2.6). The inverse transformation is given by

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r\cos\theta\sin\phi \\ r\sin\theta\cos\phi \\ r\cos\phi \end{bmatrix} \tag{2.22}$$

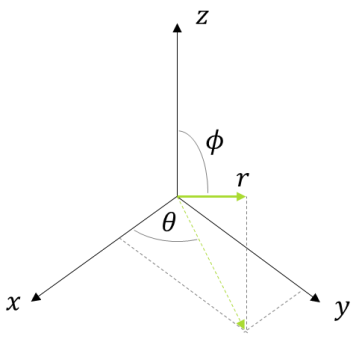

Figure 2.6: Illustration of spherical coordinates, as described in the main text.

2.1.4 Motion

As anyone who has used a camera will know, motion has a serious and important effect on the nature of a captured image. When there is relative motion of a camera and the scene it is imaging, the image plane coordinates u and v corresponding to a given point in the scene are time dependent. Over the finite duration that an image is captured, the changing values of u and v result in a blurring of the image, known as motion blur. Motion blur reduces the angular resolution of an image. In scientific imaging, there is a trade-off between the amount of signal collected in an image (i.e. the duration over which the image is captured) and the resolution of the image. But motion is also often used to facilitate imaging. The cameras studied in this thesis rely on their relative motion with the scene to construct their images. The effect of motion is an important consideration in the work presented in both chapters 5 and 6, and a general mathematical description of it is therefore given here.

For a central perspective camera, equation 2.2 describes the relationship between image plane coordinates and 3D coordinates in the camera frame. Relative motion of the camera and scene results in the time evolution of points' $[x \quad y \quad z]^T$ coordinates. The rate of change of a point's image plane coordinates as a function of the rate of change of its camera frame coordinates is given by:

$$\frac{du}{dt} = f\left(\frac{z\frac{dx}{dt} - x\frac{dz}{dt}}{z^2}\right)$$

$$\frac{dv}{dt} = f\left(\frac{z\frac{dy}{dt} - y\frac{dz}{dt}}{z^2}\right) \tag{2.23}$$

Similar, albeit less succinct relationships can be derived for the equidistant imaging geometry. Differentiating equation 2.9 with respect to time gives

$$\frac{du}{dt} = \frac{u_{denom}\frac{dr_{im}}{dt} - r_{im}\frac{du_{denom}}{dt}}{u_{denom}^2}$$

$$\frac{dv}{dt} = \frac{v_{denom}\frac{dr_{im}}{dt} - r_{im}\frac{dv_{denom}}{dt}}{v_{denom}^2} \tag{2.24}$$

where r_{im} is given by equation 2.10, and u_{denom} and v_{denom} are the denominators of equation 2.12:

$$u_{denom} = \sqrt{\left(\frac{y}{x}\right)^2 + 1}$$

$$v_{denom} = \sqrt{\left(\frac{x}{y}\right)^2 + 1} \tag{2.25}$$

The time derivatives of u_{denom} and v_{denom} are given by

$$\frac{du_{denom}}{dt} = \frac{yx\frac{dy}{dt} - y^2\frac{dx}{dt}}{x^3\sqrt{\left(\frac{y}{x}\right)^2 + 1}}$$

$$\frac{dv_{denom}}{dt} = \frac{xy\frac{dx}{dt} - x^2\frac{dy}{dt}}{y^3\sqrt{\left(\frac{x}{y}\right)^2 + 1}} \tag{2.26}$$

whilst the time derivative of image radius is

$$\frac{dr_{im}}{dt} = \frac{f'\left(zx\frac{dx}{dt} + zy\frac{dy}{dt} - (x^2 + y^2)\frac{dz}{dt}\right)}{(x^2 + y^2 + z^2)\sqrt{x^2 + y^2}} \tag{2.27}$$

These equations are in the camera frame, therefore treating the camera as stationary whilst the scene moves. They can be applied to any combination of camera and scene motion through the appropriate application of above-discussed coordinate transforms.

2.2 Image Acquisition

The nature of a camera's observations is not only a function of its imaging geometry, but also the way in which it is operated. Indeed, a single camera can produce markedly different image products depending on its method of image acquisition. Many modern smartphones exemplify this with their panorama image function: the user slowly rotates the phone whilst it continuously captures images, building up a panoramic photograph which is fundamentally different to the snapshots captured with the standard camera function. Planetary imaging makes use of a wide variety of camera designs and imaging setups, but they can all be categorised as using one of two methods of image acquisition: frame imaging or scan imaging.

2.2.1 Frame Imaging

The frame imaging technique is the method by which almost every consumer camera operates, and will likely be familiar to anyone who has ever taken a photo. Indeed, many of the cameras that have flown or are flying on spacecraft employ this method too (and are referred to as framing cameras). Because the focus of this thesis is scanning cameras, little will be said here on framing cameras, other than to highlight their main differences from scanning cameras.

When a framing camera captures an image, it records a snapshot of what it can see at that moment. The resulting image has a shape, dimensions and FOV tied to the shape and dimensions of the detector and the focal length of the optics (as described in section 2.1). Framing cameras benefit from a constant imaging geometry (ignoring the effects of structural deformation, caused for example by thermal expansion or vibrations experienced during launch). Routine calibration techniques are well established which allow a camera's imaging geometry to be well determined, leading to high levels of geometric fidelity in their images (see section 2.3.2).

Framing cameras are commonly used on three-axis stabilised spacecraft, because of the stable platform they provide, and the flexibility to accurately point the camera in any direction (they are seldom used on spin stabilised spacecraft for the same reasons). Often, spacecraft pointing and vibration will be stringently controlled, allowing accurate pointing of very narrow FOVs (often less than a degree) and long exposures to capture

high-resolution, high-signal images of planetary surfaces. Framing cameras are also common on landed platforms such as rovers.

Figure 2.7 shows four examples of images captured by framing cameras throughout the solar system, to highlight the diversity of observations of which the technique is capable. Note that images (a) and (b), of the Jupiter and Saturn systems respectively, were captured by the same instrument (Voyager 1's Imaging Science Subsystem), demonstrating the flexibility that a framing camera with pointing control provides.

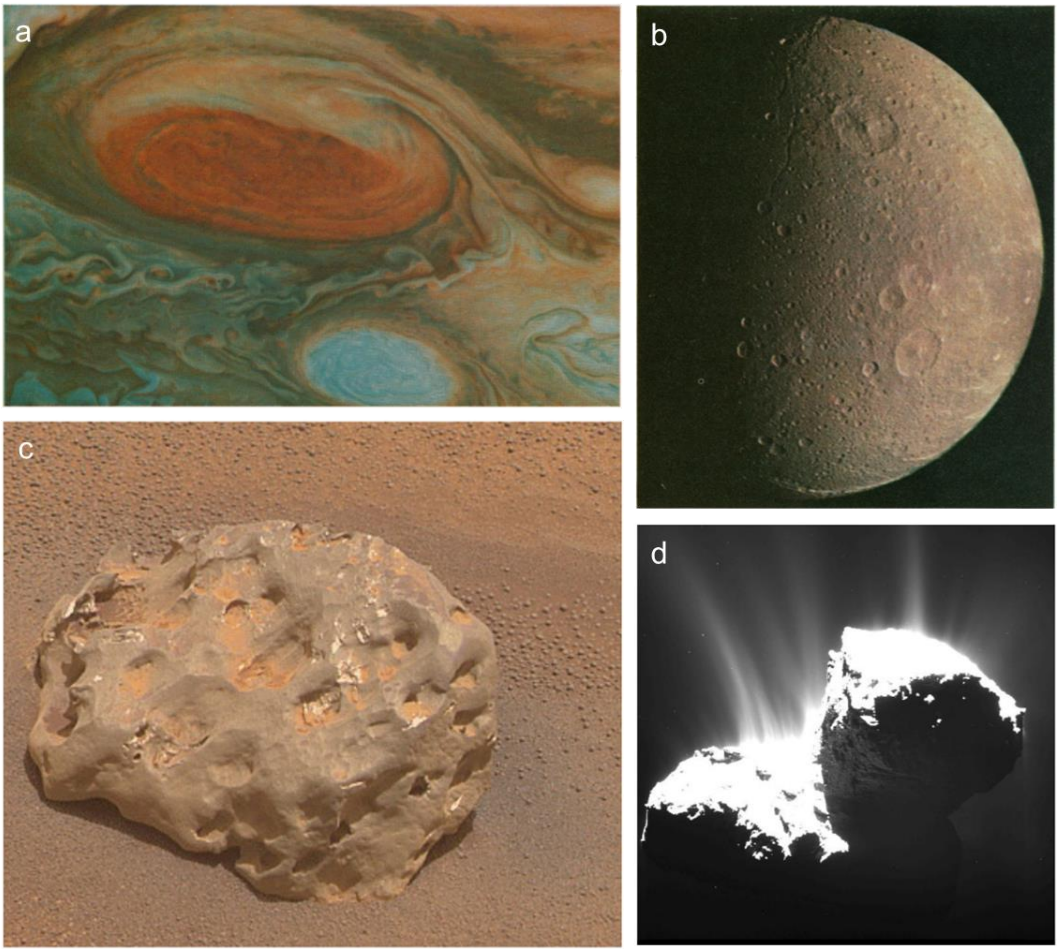

Figure 2.7: Examples of the wide variety of images captured by framing cameras throughout the solar system. a) Jupiter's Great Red Spot imaged by Voyager 1's Imaging Science Subsystem (ISS) (Smith, et al., 1979b). b) Saturn's moon Dione, imaged by Voyager 1's ISS (Smith, et al., 1981). c) Image of an iron meteorite, named 'Heat Shield Rock' on Mars' surface, imaged by the Mars Exploration Rover Opportunity's Pancam (Squyres, et al., 2006). d) The nucleus and dust jets of 67P during the Rosetta encounter, captured by the OSIRIS instrument (Tubiana, et al., 2015).

2.2.2 Scan Imaging

The fundamental difference between frame imaging and scan imaging is the effect that relative motion between the camera and the scene has on the image. In frame imaging, motion is undesirable and impacts both the resolution and signal strength of the obtained image. In contrast, scanning cameras move by design, changing their viewing perspective during image capture so as to extend the scope of their images beyond the instantaneous view of their detectors. Data are recorded whilst the camera moves, and combined to build up a single image, whose size grows with scan duration. This is illustrated in Figure 2.8.

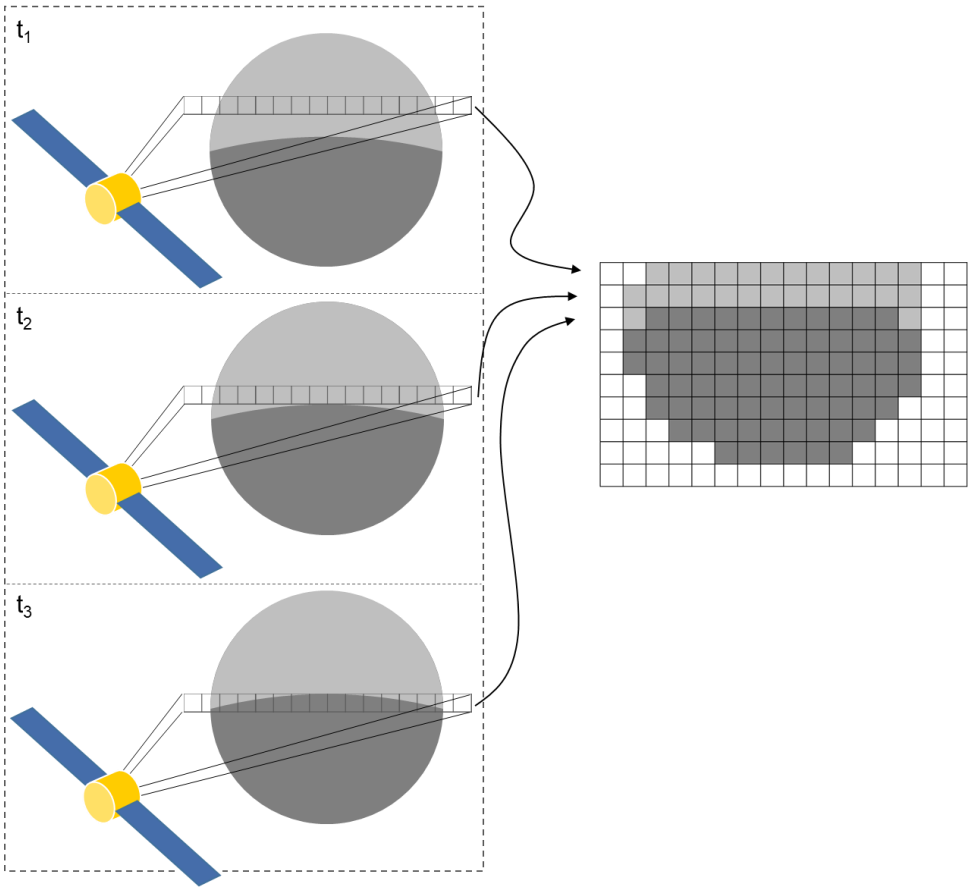

Figure 2.8: Illustration of the scan imaging principle. The detector's footprint moves over the scene (left) and captures narrow images at regular intervals, which are combined to build up a larger 2D image (right).

Effective use of scan imaging requires that the relative camera-scene motion can be constrained, such that the geometry of the final image is as desired. Generally, the instrumental setup has one of the detector's axes aligned with the direction of scene motion (i.e. the scene moves either horizontally or vertically across the detector). The FOV of the camera is then described by its across-track FOV (the extent of the FOV

measured perpendicular to the direction of scene motion) and its along-track FOV (the extent of the FOV measured parallel to the scene motion), as shown in Figure 2.9. Unlike with frame imaging, the final image has a greater coverage than the instrument's FOV. The total region imaged by the final scan is called its field of regard (FOR).

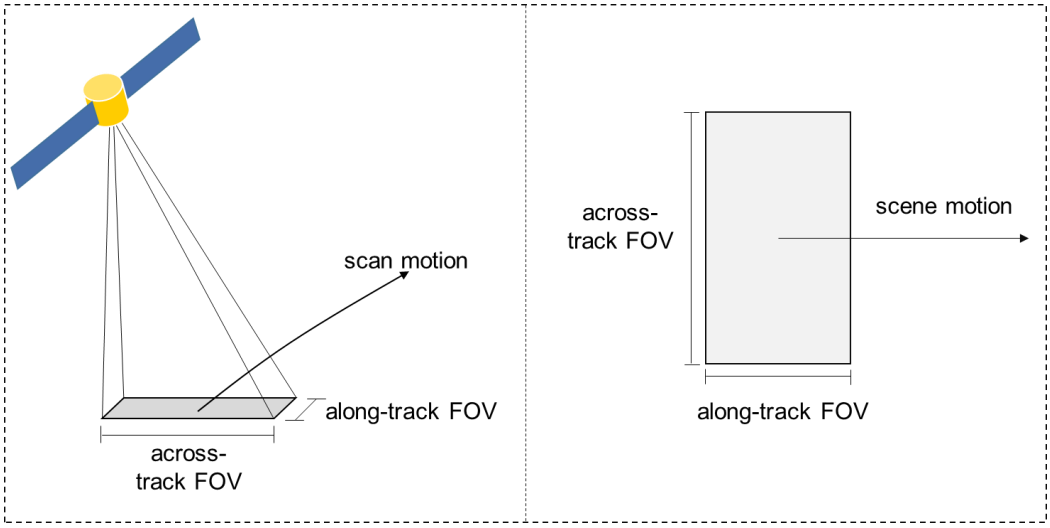

Figure 2.9: Illustration of the across-track and along-track FOVs in the scene frame (left) and the detector frame (right).

Scan imaging is well suited to situations where there is significant relative motion of the spacecraft and imaged object. FORs greater than the instrument's FOV can be achieved, allowing large-coverage images to be captured. However, because an image is captured over a significant period of time, and its imaging geometry is linked to the camera's motion during that time, geometric fidelity of the images is limited by the stability of the camera's platform, and often needs to be corrected for unwanted motion (e.g. Gianinetto and Scaioni (2008)).

To avoid blur and achieve a uniform spatial resolution, the exposure time of a scanning camera is limited by the rate of scene and camera motion. This is characterised by the imager's dwell time, which is the duration for which a single point in the scene remains within a single pixel. It is commonly taken that the dwell time of the system defines the maximum blur-free exposure time (Strojnik & Paez, 1997).

Two distinct methods of scan imaging, namely push-broom and push-frame imaging, are commonly used on spaceborne planetary imagers, and each will be described separately here.

2.2.2.1 Push-broom

Push-broom cameras capture images with a one-dimensional sensor, often called a line sensor. The sensor is oriented such that its long axis aligns with the instrument's across-track FOV, and the along-track FOV is parallel to the sensor's short axis. A 2D image is constructed by concatenating 1D image strips, which are repeatedly captured by the sensor at a frequency linked to the system's dwell time (Gupta & Hartley, 1997). By producing a 2D image from 1D lines, all of which are captured by the same sensor, it is ensured that every line of the final image has the same across-track geometry. Push-broom cameras often require no moving parts, minimising design complexity and protecting against points of failure. An example push-broom image, captured by the Mars Reconnaissance Orbiter's (MRO) Context Camera (CTX), is shown in Figure 2.10.

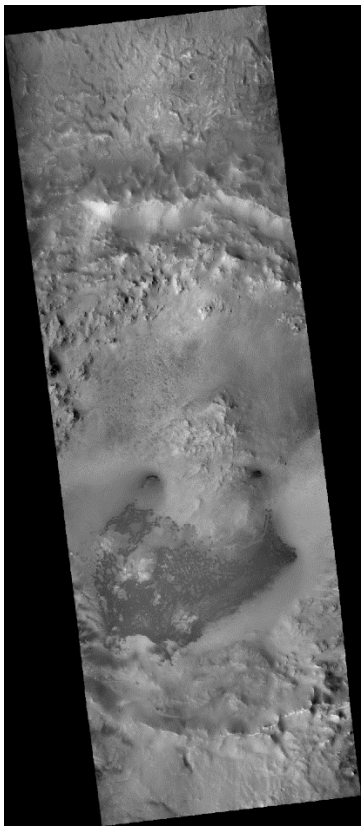

Figure 2.10: A push-broom image of Mars' surface, captured by CTX (Malin, et al., 2007).

A 2D array detector such as a CCD can be used in a push-broom camera to provide multiple parallel scan lines (each column of a detector can act as its own scan line), or several separate line sensors can be incorporated into a single instrument. This allows the simultaneous capture of multiple images. Different spectral filters can be placed over each line sensor, facilitating multispectral imaging without the mass, complexity or risk

of failure associated with a filter wheel. The High Resolution Stereo Camera (HRSC) aboard Mars Express (shown in Figure 2.11), as well as the High Resolution Imaging Science Experiment (HiRISE) aboard MRO (McEwen, et al., 2007) employ this method of multispectral imaging.

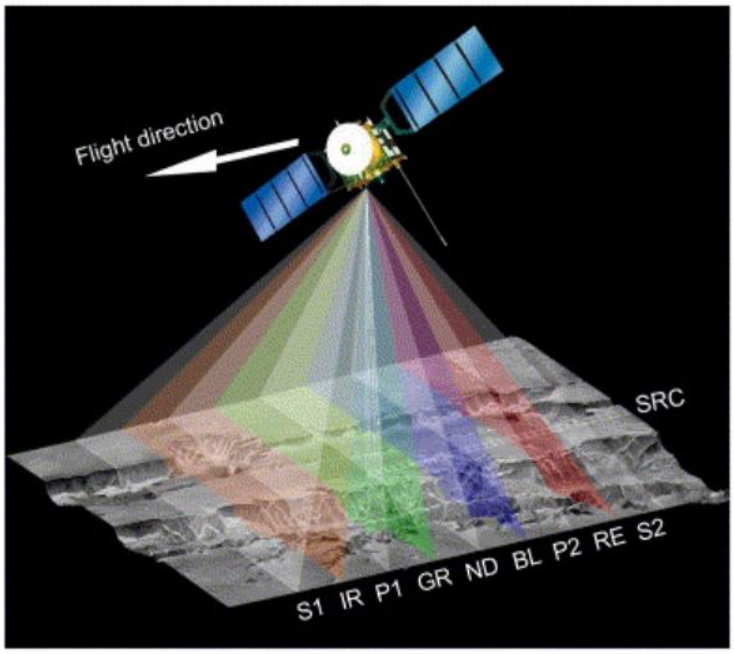

Figure 2.11: The HRSC push-broom operating principle, showing its multiple individual scan lines. Reproduced from Jaumann, et al. (2007).

Push-broom cameras are commonly used on orbital spacecraft at Earth (e.g. OLCI (Nieke, et al., 2012)), the Moon (e.g. LROC NACs (Robinson, et al., 2010)) and Mars (e.g. HiRISE (McEwen, et al., 2007)). The spacecraft's orbital motion and the spin of the target body allows large surface coverage to be obtained in short timescales. Regular, repeat passes over the same region of surface allows systematic change detection (e.g. Di et al. (2014)). In contrast, the Halley Multicolour Camera (HMC) aboard Giotto performed push-broom imaging with motion provided by the spacecraft's spin stabilisation (Keller, et al., 1987). Only a small portion of the instrument's FOR, which coincided with the Comet Halley, was captured and saved.

Scan imagers in general can suffer from distortions to their imaging geometry when unwanted motion occurs during scanning. Gupta and Hartley (1997) describe some of the challenges involved with modelling an orbital push-broom camera's motion (e.g. the varying speed and altitude of the camera according to Kepler's laws, and the spacecraft's slow rotation during orbit), and proceed to simplify the motion to a straight line for the sake of analysis. For any system, there will unavoidably be some level of disparity

between the modelled/measured motion and the true motion of the camera, leading to some presence of geometric errors. Because push-broom imagers only capture a 1D image strip at any one time, inaccurate scan motion can lead to holes in the camera's coverage (Sutton, et al., 2017).

2.2.2.2 Push-frame

Push-frame imaging could be described as a mixture of frame imaging and push-broom imaging. As with the push-broom technique, the scene is scanned across the camera's FOV using their relative motion, and a 2D image is built up over time. Instead of using a 1D line sensor, a push-frame camera uses a 2D detector, which repeatedly captures 2D images, called framelets. Framelets will often take the form of an image strip with a longer across-track axis than along-track axis. Commonly, framelets are captured frequently enough to result in overlap between adjacent framelets. As with push-broom cameras, a large 2D detector can be used to capture multiple scans simultaneously. For example, the Juno spacecraft's JunoCam captures red, green and blue framelets simultaneously, producing three images in a single scan. JunoCam's focal plane and example data are shown in Figure 2.12.

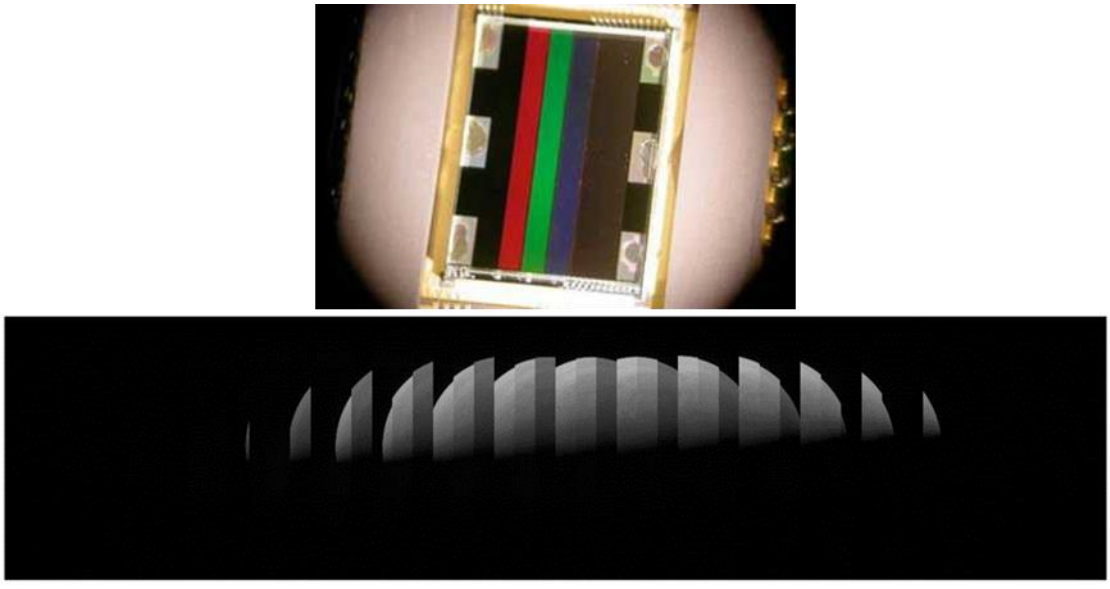

Figure 2.12: (top) JunoCam's detector, with each channel's coloured filters bonded to its surface. Reproduced from Hansen et al. (2017). (bottom) Example raw data from JunoCam, comprising multiple consecutive sets of three framelets, captured simultaneously through the red, green and blue filters (Caplinger, 2014).

Note that the raw JunoCam data of Figure 2.12 (bottom) comprise data from three different filters (the red, green and blue filters shown in the top of the figure). The

framelets are noticeable by the discontinuities in shape and colour. Reconstruction of the scene requires the separation of the three channels' respective data, and the stitching together of the individual framelets belonging to each. These extra steps complicate the process of using push-frame data, as is described by Anderson and Robinson (2009), but are not significant barriers, given the high levels of processing required for all raw scientific image data.

Because push-frame framelets are captured with overlapping footprints, construction of a final image requires mosaicing the framelets together. Re-projection of the framelets based on the camera's imaging geometry may be necessary. The overlap of framelets can assist with determining the camera's scan motion, and recovering an image of high geometric fidelity even when unwanted camera motion occurs (Strojnik & Paez, 1997).

2.3 Imaging Hardware

2.3.1 Detectors

Responsible for converting the light entering a camera into a storable and viewable image, detectors and their characteristics have a significant impact on the performance of an imaging system. For centuries, astronomers had to record, by hand and eye, their observations by means of drawing or carving (Griffiths, 2014). Photographic plates allowed the first physical recordings of astronomical observations in the 19th century, removing the subjectivity from the recording process (e.g. Stroebel and Zakia (1993); Greeley (1994)). From photographic plates followed photographic film, which persisted in consumer cameras into the late 20th century, and was even used in spaceborne cameras (for example, the Luna 3 spacecraft imaged onto film which was developed onboard, before being scanned and transmitted to Earth (Heacock, 1968)). Over the last four decades, detector technology has advanced significantly, and has transitioned to digital sensors, in line with the ubiquitous digitisation of almost all data formats. The improvement of detector technology has been a crucial contributor to significant advances in spaceborne imaging capabilities.

For almost the whole of the electromagnetic spectrum, from radio (e.g. Picardi et al. (2005)) to x-ray (e.g. Burrows et al. (2005)) there are cameras and telescopes aboard spacecraft, utilising digital detector technologies to image solar system objects. For detection in optical wavelengths, two main technologies have emerged and gained significant space heritage: charge-coupled devices (CCDs) and complementary metal-oxide-semiconductor (CMOS) detectors. The purpose of a spaceborne camera's detector is not to simply record the appearance of a scene, but to perform two-

dimensional radiometry, i.e. to make accurate angularly-resolved quantitative measurements of the light's radiant energy as seen by the camera.

2.3.1.1 Signal, Noise and SNR

Later in this thesis (chapter 6, section 6.3) there will be significant discussion of estimating and modelling the signal measured by a camera's detector, as well as the effects it has on derived products and scientific analysis. This section therefore sets out the theory behind image signal, noise and SNR (signal to noise ratio).

In a raw scientific image, each pixel represents a single individual measurement, linked to the number of photons that arrived at the pixel during the measurement period. As with any scientific instrument, its measured values will contain both the desired signal and unwanted noise, limiting the accuracy with which this number of photons can be constrained. In digital detectors, which transduce photons to voltages, there are several steps at which noise is introduced. It is of course desirable to have as strong a signal, and as small a noise contribution as possible, and this is generally characterised by SNR.

SNR is used throughout signal processing, not just in imaging science, and its definition varies. The discussion of SNR here will be tailored to digital detectors whose operation is based on the counting of photoelectrons within pixels (covered further in sections 2.3.1.4 and 2.3.1.5). Photoelectrons are generated by incident photons, with the ratio of photoelectrons to incident photons defining the detector's quantum efficiency, q. The signal N is the 'true' number of photoelectrons that the pixel should count in response to the incoming light (i.e. if no noise were present, repeat measurements would always measure N electrons).

In the presence of noise, the mean value of repeat measurements tends towards N as repetitions increase. The signal N is a function of the viewed scene, the imaging geometry of the detector's FOV, the quantum efficiency, the spectral properties of the camera and detector, and the exposure time (the duration over which signal is measured). Because of various noise sources, repeat measurements do not in reality measure N, but instead measure a random distribution of values around N, whose standard deviation is σ_N. The SNR is then simply given by

$$SNR = \frac{N}{\sigma_N} \tag{2.28}$$

In the detecting chain, from photons arriving at a pixel, to a digital value representing that pixel's measurement being output from the camera, multiple sources of noise are

introduced (Nieke, et al., 1998). In this thesis, these noise sources will be characterised by three main contributions: shot noise, dark noise and read noise.

Shot noise is intrinsic to the light being measured, and results because it arrives as discrete independent photons, rather than a continuous signal (Healey & Kondepudy, 1994). When those photons arrive, and at what intervals, is probabilistic, and subject to variation. Shot noise is the resulting uncertainty in the number of photons (and therefore photoelectrons) arriving at a pixel. Because it is intrinsic to the signal itself, it cannot be eliminated, and it sets a limit for the maximum achievable SNR. Shot noise follows the Poisson distribution, in which the variance is equal to the mean, and the standard deviation N_{shot} is therefore given by

$$N_{shot} = \sqrt{N} \tag{2.29}$$

If the only source of noise is shot noise, the SNR is given by

$$SNR_{shot} = \frac{N}{\sqrt{N}} = \sqrt{N} \tag{2.30}$$

The detector's purpose is to allow photoelectrons to accumulate over a specified amount of time (the exposure time), and then count them to give a measurement of the number of photons to have arrived at the detector. However, during the accumulation of photoelectrons, thermal energy within the sensor leads to the creation of free electrons. The build-up of thermal electrons in the detector is called dark current, because it occurs even when no light strikes the sensor, and it grows with detector temperature. Dark current's magnitude is a property of a detector's design and materials, and the nature of a detector's dark current can be characterised prior to imaging, allowing its subtraction from pixels' final measurements (e.g. Skinner et al. (1998); Widenhorn et al. (2010)).

Dark current, N_d is usually characterised by a rate of electrons per second, meaning over a time period t, a pixel would accumulate $N_d\ t$ thermal electrons. For an image with any given exposure time, the contribution from thermal electrons can be estimated and subtracted, recovering the original signal. However, because dark noise comprises discrete electrons, it too behaves according to Poisson statistics. For an exposure of duration t, with a pre-measured dark current of N_d, there will be an uncertainty (standard deviation) of

$$\sigma_d = \sqrt{N_d\ t} \tag{2.31}$$

in the number of thermal electrons counted.

After the full duration of an exposure, a detector reads out its collected charge by converting it to a voltage and eventually a digital number. The electronics involved in this step introduce additional noise which is often encapsulated by one value, called read noise. Read noise is generally modelled as additive white Gaussian noise (AWGN), with a zero mean, and a standard deviation, N_r (Jingjing, et al., 2007).

With shot noise, dark noise and read noise all present, the resulting total variance in measured signal is given by

$$\sigma_{total}{}^2 = N + N_d\, t + N_r{}^2 \tag{2.32}$$

and the SNR of the pixel's measurement by

$$SNR = \frac{N}{\sqrt{N + N_d\, t + N_r{}^2}} \tag{2.33}$$

Further possible sources of noise exist, such as internal instrumental stray light (e.g. Zhong and Jia (2009); Tubiana et al. (2015)), and equation 2.33 can be readily extended to include the important noise sources of a particular instrument. In this thesis, modelling the major noise sources in a pixel's measurement with equation 2.33 (i.e. shot noise in the measured signal, a time-dependent detector noise and a time-independent detector noise), is sufficient for estimating the expected SNRs of the studied instruments.

2.3.1.2 Dynamic Range

An important property of a scientific camera, which depends closely on its noise behaviour, is its dynamic range. The goal of a camera's detector is to acquire quantitative measurements of a scene's brightness, and differences in the brightnesses measured by each pixel provide structural information about the scene. The dynamic range of a detector is the maximum range of brightnesses to which it is sensitive.

Equation 2.32 shows that even when no signal is present, noise will lead to a random, non-zero measurement in a detector's pixels. This is called the detector's noise floor, the magnitude of which depends on the specific noise characteristics of the detector. Any signal smaller than the noise floor cannot be measured, because it cannot be distinguished from the instrument's noise. The noise floor therefore defines the minimum signal to which the camera is sensitive.

At the other extreme, detectors have a physical upper limit to the amount of signal they can measure. A sensor can only respond to a finite amount of signal before its behaviour becomes non-linear or unpredictable. At this stage the detector becomes saturated, and

it is no longer possible to recover the measured signal from the detector's output. A detector's maximum measurable signal is therefore dictated by its saturation threshold. The dynamic range of a detector is the ratio of its saturation threshold to its noise floor.

A high dynamic range is particularly important when viewing scenes with significant brightness variation. As an example, the OSIRIS instrument on the Rosetta spacecraft captured high resolution images of comet 67P/Churyumov-Gerasimenko containing large brightness contrast. A typical view would capture both fully illuminated and fully shadowed portions of its nucleus, and would also include dust surrounding the nucleus, which scatters sunlight toward the camera, but with a significantly lower brightness than the nucleus' surface. The OSIRIS team estimated that, during Rosetta's encounter, the ratio of the signal received from the dust to that from the nucleus' surface would be as low as 4×10^{-4}. To capture an image in which the signals from dust and nucleus are both accurately measured therefore required a dynamic range greater than 2000 (Keller, et al., 2007).

2.3.1.3 Vidicon

The first spacecraft came at a time when astronomical photography (performed with Earth-based telescopes) was largely film-based. The prospect of capturing and retrieving images from distant spacecraft required the adoption of new detector technology capable of producing a strong, digitisable signal such that images could be transmitted from spacecraft to Earth without introducing additional noise (Smith & Tatarewickz, 1985). During this period, other image capture and broadcasting applications were utilising cathode ray tubes, which allowed cameras' images to be rapidly electrically scanned for redisplaying or storing. A specific implementation of a cathode ray tube, known as vidicon, became popular as the method of image detection on planetary space missions throughout the 1960s, '70s and '80s (Rodgers, et al., 1980).

Vidicon imaging works by scanning an electron beam over a photoconductive surface which has been illuminated with an image of the scene by the camera's optics. The electron beam repeatedly scans the image row by row, and a current of electrons flows out of the photoconductive surface into an amplifying circuit. This output current is modulated by the level of illumination at the point of the electron beam's focus, and combined with the scan rate and number of scans per frame, describes the scene's image. Exposure of the vidicon's sensor surface to the scene is controlled by a mechanical shutter. The amplified signal retrieved from the vidicon can by digitised, stored, transmitted and reconstructed into an image. Spaceborne imagers used slow-scan vidicon, whereby the electron beam is scanned slowly, taking 10s of seconds to

scan a full image (in contrast to television systems which operated at ~30 Hz) (Malling & Allen, 1966).

Vidicon-based cameras were flown on a large number of spacecraft, including the Ranger, Pioneer and Voyager missions. Resolution of slow-scan vidicon images was limited by the focus of the electron beam, and was generally between 200 and 1000 pixels per image side. Vidicon imagers generally have dynamic ranges between 100 and 1000.

2.3.1.4 CCDs

In the 1980s, a desire for greater image resolution, lower image noise and cameras with reduced mass and size motivated the replacement of slow-scan vidicon detectors with an alternative sensor technology. The CCD (charge-coupled device) dates back to 1969 when it was devised and developed at Bell Labs (Smith, 2009). Constructed from semiconductors, a CCD is a 2D array of discrete light detectors (pixels). Light striking a pixel creates a charge within the semiconductor material, where it is then stored. Each pixel accumulates and stores its own, independent charge. Charges can be moved between pixels within the array by manipulating the voltages across the detector. Images are captured by exposing the detector to light for a controlled duration, and then shifting the charges out of the CCD's pixels and into readout electronics. These electronics convert each pixel's charge into a voltage representative of the signal measured, and these voltages are then digitised by an analogue-to-digital (A/D) converter (Boyle & Smith, 1970).

By processing the charges with electronics separate to the sensor, CCDs offer flexibility as to how images are captured, and a single sensor can be used in a variety of ways (e.g. sections 2.3.1.6-2.3.1.8) (Litwiller, 2011). This flexibility is ideal for spaceborne cameras, as it allows multiple imaging modes to be implemented with a single hardware setup. However, the design of CCD electronics limits the speed with which a full image can be retrieved from the detector, because charges must be read out serially. Using external electronics also limits the extent to which a CCD-based camera can be miniaturised.

CCD images are contemporaneous, meaning all pixels' signals are measured simultaneously and for the same duration. Exposure of a CCD can be achieved by mechanical shutters, but can also be controlled by electronic shuttering, whereby the CCD rejects any charge built up during its 'shutter closed' state (electronic shuttering doesn't eliminate charge build up, but reduces it to ~0.001% of normal levels). Electronic

shuttering allows fast ‘opening’ and ‘closing’ of the shutter (on the order of a few nanoseconds) and removes the need for moving parts (Reich, et al., 1993).

In its simplest form, a CCD records the light intensity at each pixel, but does not record any colour information, thus producing a greyscale image. CCDs can typically be designed to be responsive to visible, ultraviolet (UV) and X-ray wavelengths (Burke, et al. (1992); Godet, et al. (2009)). CCDs are also used to image near infrared (NIR) wavelengths, but at longer IR wavelengths thermal noise becomes very strong and photon energy becomes too low to generate signal (Thom, et al., 1980).

A vulnerability in the CCD’s method of readout, and its need for efficient charge transfer across its array, is a sensitivity to radiation damage (Fossum, 1993). The risk of radiation effects generally increases with number of pixels within a CCD. Radiation damage can increase a detector’s dark noise and non-uniformity, and decrease the dynamic range (Gilard, et al., 2008).

2.3.1.5 CMOS and Active Pixel Sensors

Active pixel sensors (APS) are a category of digital detectors that are distinct from CCDs due to their different approach to processing pixels’ measurements. Each pixel of an APS has its own active amplifier electronics, for converting the pixel’s accumulated charge to voltage. This removes the need for pixel-to-pixel charge transfers during the readout stage, helping to increase the radiation tolerance of the device, and can assist the design of compact, low mass cameras. APS are however prone to their own forms of radiation damage, such as single-event latch-ups (Hopkinson, et al., 2004).

Whilst APS were first developed in the 1960s, CCDs were the favoured technology in the early days of digital photography. The development and adoption of APS in both scientific and consumer cameras therefore lagged behind that of CCDs. Over the last couple of decades, APS sensors have caught up in performance and popularity, and their use for spaceborne planetary imaging is increasing. The most common type of APS, the complementary metal-oxide-semiconductor (CMOS), outnumbers CCDs in many consumer applications, such as phone cameras, and is gaining space heritage from the incorporation of CMOS detectors in a number of missions (Soman, et al., 2014; Coates, et al., 2017; Rochus, et al., 2020).

Rather than reading out all pixels serially through a single A/D converter, a CMOS sensor generally has one analogue-to-digital (A/D) converter per row of pixels. The sensor can read out its rows independently, and the process can be completed quicker than an equivalently sized CCD. Commonly, CMOS sensors capture images with a rolling shutter

to maximise their readout efficiency. With such a technique, lines which are not being read out can continue to capture light, allowing rapid acquisition of images whilst maximising signal. The drawback of this is that each line of an image is captured at a different time, which can lead to image distortion if scene motion is present (Chun, et al., 2008). Some CMOS sensors employ a global shutter mode in order to avoid this issue.

Integrating electronics onto the pixels of a CMOS allows the size of imagers to be greatly minimised. Indeed, electronics responsible for amplification, A/D conversion, detector timing and sensor operation can be integrated onto the same chip as the pixels, allowing a complete imager on a single chip (Fossum, 1997). This is beneficial to the design of low-complexity, low-mass compact cameras, providing potential to image from a greater range of spacecraft. Placing the electronics on the chip does however reduce the portion of the detector that is sensitive to light, lowering the quantum efficiency of the device. This can impact the suitability of CMOS to low light conditions (Donati, et al., 2007).

2.3.1.6 Pixel Binning

In the design of an imaging system, image signal strength must be traded off against angular resolution. Spaceborne cameras can encounter a range of brightnesses throughout their operation, and the optimum balance of SNR and resolution can vary significantly. A technique known as pixel binning is often employed by planetary imagers to provide some in-flight flexibility to this trade-off. Pixel binning is the process of adding together the measurements of several individual adjacent pixels. This is usually done for a square number of pixels (i.e. 2 x 2, 4 x 4…). By combining multiple pixels' measurements into a single value, signal is increased at the expense of angular resolution.

Pixel binning can be performed on any image by simple addition of its pixels' values. As a method of controlling the image's SNR, binning is more effective when it is performed by the camera's detector (known as 'on-chip'), because the noise in the final image is minimised. This is discussed further in chapter 6, section 6.4.1, where specific binning methods and their effect on SNR are quantitatively described. Not all detectors can perform on-chip pixel binning, but it is commonly utilised by CCD-based spaceborne cameras (e.g. HiRISE (McEwen, et al., 2007) and OSIRIS (Keller, et al., 2007)). Few APS detectors support the functionality, but there are examples of it within planetary imaging (e.g. the Rosalind Franklin rover's CLUPI (Josset, et al., 2017)).

In addition to providing control over resolution and SNR, pixel binning allows data volume to be reduced. This is beneficial if it is expected that, during a mission, full resolution data

rates will exceed the storage or transmission limits of an instrument or its spacecraft. Indeed, pixel binning may be implemented purely for this reason.

2.3.1.7 Time-Delay Integration

The exposure times of scanning cameras are necessarily limited by their motion relative to the scene they are imaging. As a result, when dwell times are short or a camera is viewing a dark scene, images can suffer from low SNRs. Time-delay integration (TDI) is a method of increasing the effective exposure time, without inducing pixel smear or motion blur (Wong, et al., 1992).

As a push-broom or push-frame camera captures an image, the scene progresses steadily across its detector due to the camera's motion. A detector implementing TDI periodically shifts its pixels' charges in order to match the motion of the scene. The accumulating packets of charge track their respective light sources across the detector, allowing them to accumulate signal for more than the system's dwell time, and resulting in higher SNRs in the measured image (Gibson & Hickson, 1992).

Charges are restricted to travelling along one of a detector array's two dimensions (i.e. either along rows or along columns). It is therefore necessary for scene motion to also follow this same path. Because the charge is shifted at regular intervals, the rate of scene motion should also be uniform. These requirements impact when TDI is and isn't a suitable method to increase a camera's SNR (this is covered in greater detail in chapter 6, section 6.4.3). As with pixel binning, another key factor is that TDI is not a standard function of digital detectors. Space grade CCDs are often capable of performing TDI, but it is much less commonly implemented with APS.

TDI is used or has been used on a large number of space-borne push-broom and push-frame cameras. Whilst it is commonly implemented on high resolution orbital surface imagers, it has also been used on spinning spacecraft (see Keller, et al. (1987) for Giotto's HMC and Hansen, et al. (2017) for Juno's JunoCam). The typical number of TDI steps used in spaceborne imaging ranges from 2 to 128 (see e.g. the HiRISE (McEwen, et al., 2007) and THEMIS (Christensen, et al., 2004) instruments).

2.3.1.8 Frame-transfer Detectors

In a camera's simplest implementation, its optics will focus light to form an image on the detector's 2D array of pixels, which will be used to record and read out the image. Clearly, from sections 2.3.1.6 (pixel binning) and 2.3.1.7 (time-delay integration), the role of the detector can be expanded with additional responsibilities. Frame-transfer detectors are

another example of this, whereby the detector is also responsible for performing short-term storage of its measurements.

Frame-transfer detectors are segmented into imaging regions and storage regions. The storage regions are optically masked, such that they do not measure any signal, whilst the imaging regions are exposed to the camera's image. Charges accumulated in an imaging region's pixels can be shifted into a storage region's pixels, where they cease to accumulate any further charge. Once an image has been transferred to the storage region, it can be read out from the detector. Because the read out is not occurring directly from the imaging region, another image can be exposed whilst the readout takes place. The technique is valuable when the impact of readout time on image exposure needs to be minimised, such as when only a short period of time is available in which to capture multiple images. CCD charge transfer times are significantly faster (up to orders of magnitude) than readout times, meaning use of a frame-transfer detector can facilitate higher frequency imaging (Olson, 2002).

The Descent Imager/Spectral Radiometer (DISR) on the Huygens descent probe utilised a frame transfer CCD to capture images during its rapid descent, during which the probe continually rotated. The use of a large frame transfer CCD allowed several imagers to share a single detector, minimising instrument mass and complexity (Tomasko, et al., 1999). Use of the frame-transfer technique requires either a larger detector (in order to accommodate both an imaging region and storage region large enough for the desired image), or a reduction in image size. DISR images were relatively small format (e.g. 128x256 pixels), leading to coarse angular resolutions on the order of 0.1° (Karkoschka, et al., 2007).

A version of the frame-transfer technique, known as interline transfer, is sometimes used for rapid electronic shuttering of an image, in order to minimise image smear in moving cameras. Interline CCDs have alternating rows of unmasked (imaging) and masked (storage) pixels. After the desired exposure time, charges are shifted from the imaging rows to their adjacent storage rows, and subsequently read out of the detector. Juno's JunoCam utilises this method of electronic shuttering (Hansen, et al., 2017).

2.3.1.9 Pixel Pitch and Fill Factor

On a digital detector, pixels are arranged in a 2D grid with regular spacing. Pixel pitch is the physical distance between the centres of two adjacent pixels' light sensitive regions, as shown in Figure 2.13. Most spaceborne detectors have pixel pitches in the region of 5-40 μm (e.g. Soman et al. (2014), Jaumann et al. (2016), Bibring et al. (2007)). The

pixel pitch, in conjunction with a camera's optics, dictates the angular resolution of an imager.

The pixel pitch of a detector does not define the size of the light sensitive area of a pixel. Light-insensitive components such as electrical connections, electronics and additional infrastructure (for example the masked rows of an interline transfer detector) take up space in each pixel. The ratio of a pixel's light sensitive area to its total physical area is known as its fill factor. Often, microlenses are employed to increase the effective fill factor of a detector. Microlenses are individual optical lenses, aligned over each pixel, which focus a greater amount of light onto the sensitive portion of the pixel. This, in turn with backside illumination, can increase a detector's fill factor to close to 1 (Soman, et al., 2014).

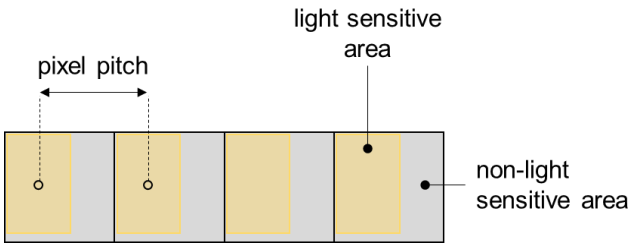

Figure 2.13: Illustration of pixel pitch and fill factor.

2.3.2 Optics

The optical elements of a camera, often called its optics, perform the function of focussing light to form an image on the instrument's detector. The imaging geometry of a camera is dictated by the design of its optics. In the case of the pinhole camera described in section 2.1, the optics are simply the hole that permits light to reach the sensor. Whilst simple pinhole cameras do have some limited applications in spaceborne imagers (e.g. Mobasser & Liebe (2004)), more complex optics are almost always necessary.

Camera optics use either glass lenses, which focus by refraction, or mirrors, which focus by reflection (illustrated in Figure 2.14). In each case, multiple optical elements are required. The first element to be struck by incoming light is known as the primary, or objective element. The advantage of lenses and mirrors versus a pinhole is that a large primary element can capture significantly more light than a pinhole, producing images with higher signal. The total amount of light captured by a camera's optics is dictated by its entrance pupil diameter. The entrance pupil is the image, as seen from in front of the optics' primary element, of the clear aperture through the optics to the focal plane. Only

light that passes through the entrance pupil can reach the detector. For optics with short focal length, the entrance pupil can be significantly smaller than the size of the primary element. As focal length grows and FOV decreases, entrance pupil tends towards the diameter of the primary element.

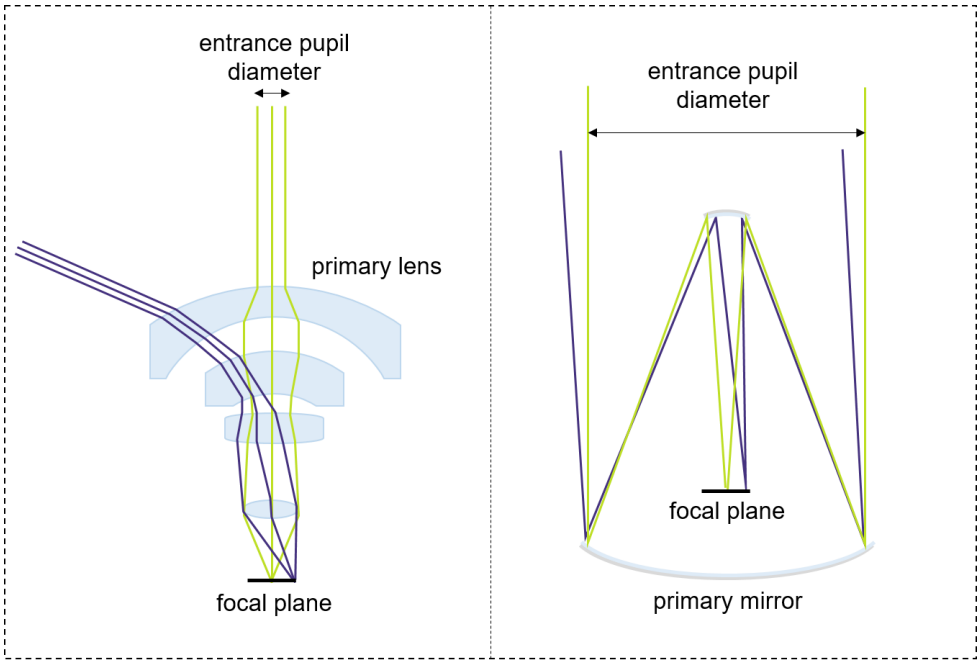

Figure 2.14: Illustration of short focal length glass optics (left) and long focal length mirror optics (right). Two sets of rays and the entrance pupil diameter are shown in each case.

The brightness (energy per unit area) of an optical system's image on its focal plane is described by its f-number $F_{/}$, which is the ratio of the focal length f to the entrance pupil diameter D:

$$F_{/} = \frac{f}{D} \tag{2.34}$$

The f-number provides an easy and convenient way to compare the image brightness of different optical systems. The larger the f-number, the lower the signal per unit area of the optics' image. Optics with small f-numbers are known as 'fast', whilst large f-numbers are known as 'slow' (because the brighter images of fast optics permit shorter exposure times). The f-number is not the sole factor in determining the signal measured in a camera's image. The nature of the detector is also important, and this is discussed in more detail in chapter 6, section 6.3.

Camera optics are designed to produce a desired imaging geometry, such as those in Table 2.1, but distortions arise from, for example, imperfect positioning and orientation of the detector, and physical distortion of optical elements. Planetary imaging, and in particular photogrammetry, require an accurate model of the camera's imaging geometry, including its distortions. Planetary cameras therefore undergo geometric calibration during their construction, achieved by imaging patterns of known geometry with the camera (e.g. Zhang (2000)). Launch stresses, temperature variation and the change in gravitational field can cause planetary cameras to deform once in space. Images of star fields are often used to refine and update cameras' geometric calibration during flight (e.g. Hansen et al. (2017)). Improperly calibrated optics can degrade the spatial resolution of the camera's images, limiting their scientific value (e.g. Farnham et al. (2007)).

2.3.3 Multispectral Imaging

The signals measured in any image are integrated over some finite spectral range associated with the sensitivity of the camera's sensor and the spectral transmission of the camera's optical elements. Characterising this spectral range is important for any planetary imaging setup, because the reflection, scattering, absorption and emission of light by planetary bodies are all functions of wavelength (e.g. Lamy & Toth (2009)).

Planetary imagers are often designed to constrain the range of their spectral sensitivity such that certain wavelengths of interest are targeted by their observations. This is almost exclusively achieved with some form of spectral filter, constructed such that its transmission of light is reduced at wavelengths outside of the range of interest (although active multispectral imaging, with artificial light sources of different wavelengths, has been used too (Jaumann, et al., 2016)). The range of wavelengths transmitted by the filter is referred to as its passband. In cameras observing optical wavelengths, passbands can range from 1-1000 nm in width, with the smaller end of the range described as narrowband, and the higher end broadband. Because filters work by rejecting light at unwanted wavelengths, they reduce the signal arriving at the detector, and the precision with which a certain wavelength can be targeted (i.e. by narrowing the passband) is therefore limited by constraints on signal to noise ratio.

Multispectral imaging is the process of capturing multiple monochrome images of the same scene, each with a different passband. The combination of such images provides information on the relative signals from different wavelengths, i.e. the colour of the scene. Multispectral imaging is often used to identify the presence or absence of certain compositions, with cameras' spectral ranges tailored to the materials of interest (e.g. Bell

III et al. (2021), Hansen et al. (2017), Keller et al. (2007)). Three common methods by which multispectral imaging is achieved in planetary cameras are here described.

2.3.3.1 Filter Wheel

The most common method by which spaceborne planetary cameras capture multispectral observations is the filter wheel. A wheel containing multiple optical filters of different spectral range is positioned between the camera's optics and detector, such that light passes through a single filter before reaching the detector. The wheel can be mechanically rotated to change which filter is positioned in front of the detector, thus allowing the camera's spectral range to be changed.

Filter wheels can accommodate as many as ~10 individual filters, and can be stacked in order to give a camera a large number of discrete spectral bands (e.g. Rosetta's OSIRIS (Keller, et al., 2007), the Rosalind Franklin (ExoMars) Rover's PanCam (Coates, et al., 2017)). A filter wheel is a complex mechanical component, which consumes volume, mass and power, so they are often reserved for larger imagers and unsuitable for miniaturised cameras. The duration required to change between two filters can be as long as ~1 second, making filter wheels unsuitable for multispectral imaging of rapidly-changing or moving scenes (e.g. Keller et al. (2007)). Filter wheels are generally qualified for millions of filter changes, but they have been known to fail, running the risk of trapping the camera on a single filter, or leaving it unable to image at all (Veverka, et al., 2013).

2.3.3.2 Static Filters

Particularly common in scanning cameras, static filters provide a means by which to capture images in multiple spectral bands without the need for moving parts. The filters are mounted in front of the focal plane (generally close to the detector's surface) or physically bonded to the detector (sometimes deposited directly on its surface). Each filter's extent covers a different region of the focal plane, and therefore corresponds to a different portion of the camera's FOV. To obtain a multispectral observation requires capturing an image for each spectral channel, with the pointing of the camera for each image set such that the footprints of the corresponding filters overlap.

This method is particularly well suited to scanning cameras because the scene necessarily moves across their FOV. Multiple filters can be placed over the detector (or detectors) such that the scene passes through all the filters in turn. This mode of multispectral imaging is used on a large number of both push-broom and push-frame cameras including HiRISE (McEwen, et al., 2007), JunoCam (Hansen, et al., 2017), HRSC (Jaumann, et al., 2007) and the LROC WAC (Robinson, et al., 2010).

2.3.3.3 Bayer Filter

Outside of scientific imaging, cameras capable of producing colour images almost exclusively use Bayer filters. A Bayer filter is an array of individual red, green and blue filters mounted in front of the detector's pixels (one filter colour per pixel). The passbands of the red, green and blue filters, and their distribution in the Bayer array, are tailored to mimic the colour sensitivity of the human eye. The resulting array has a red:green:blue pixel ratio of 1:2:1, arranged as shown in Figure 2.15.

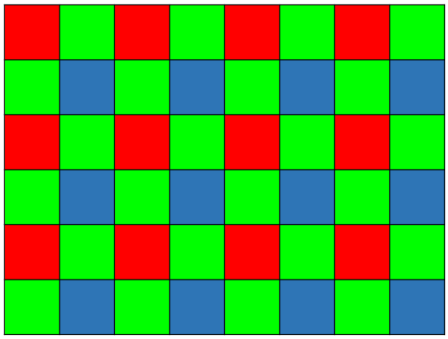

Figure 2.15: The Bayer filter configuration of red, green and blue pixels.

Because no single colour channel is imaged by all the detector's pixels, constructing a full resolution colour image requires estimating, for each pixel, the two unknown channels based on neighbouring pixels. This is known as demosaicing, and can introduce false colours and loss of resolution at small scales, potentially limiting the fidelity of the multispectral product for scientific use (Shortis, et al., 2005). Additionally, because Bayer filters are designed to target the spectral ranges detected by the human eye, they are generally not applicable to scientific multispectral imagers, for which the wavelengths of interest are instead based on the physical processes that the camera intends to study. However, they have been used in a small number of spaceborne cameras where broadband colour imaging was desirable, such as the Mastcam and Mars Descent Imager of the Mars Science Laboratory Curiosity Rover (Malin, et al., 2017). For outreach or engineering cameras, the Bayer filter is a simple and compact method by which to add multispectral imaging functionality, and it also allows colour information to be obtained by the capture of a single image.

3 Spaceborne Planetary Imaging

Sending spacecraft to visit solar system bodies is a challenging and expensive endeavour. Unlike some in situ instrumentation such as plasma analysers (Linder, et al., 1998) or sample collection (Abbey, et al., 2019), cameras can obtain observations of many solar system bodies from large distances, including from the Earth. Yet in spite of this, cameras are a mainstay of planetary space missions, and often lead the science payload (Anderson, et al., 1991). There is of course justification for sending cameras to image solar system bodies from close proximity: the possible viewing perspectives and levels of imaging detail are far superior to those that can be achieved from Earth, and many important discoveries in planetary sciences could not have been made without spaceborne cameras.

Prior to planetary cameras, only the Earth and the Moon could be imaged with resolutions of 10 km or better. Since the first spaceborne camera in 1959, this has been increased to more than 40 objects, including the eight planets, dwarf planets, moons, comets and asteroids. Many solar system objects have been the subject of follow up missions, or have been under prolonged continuous observation, producing valuable datasets such as global maps and high-resolution images with metre and even centimetre resolutions. This chapter provides an overview of the history of planetary observations from spaceborne cameras.

3.1 Imaging Throughout the Solar System

3.1.1 The Moon

Excluding the Earth, in situ imaging of planetary surfaces began, unsurprisingly given its proximity, with the Moon. The first remotely acquired images of its surface were those of the Luna 3 spacecraft in 1959 (Anderson, et al., 1991). Two of Luna 3's images are shown in Figure 3.1. It orbited behind the Moon, providing the first ever views of the lunar far side, which was unexpectedly found to lack the volcanic planes of the near side (Wu, et al., 2019). Necessitated by the need to accurately point its camera, Luna 3 was the first successful implementation of a 3-axis stabilised spacecraft (Heacock, 1968; Siddiqi, 2018). Its images, which provided insight into the role of the Earth's gravity in the

evolution of the Moon's surfaces, and the difference in processes on the near and far sides, could only have been captured by a spaceborne camera.

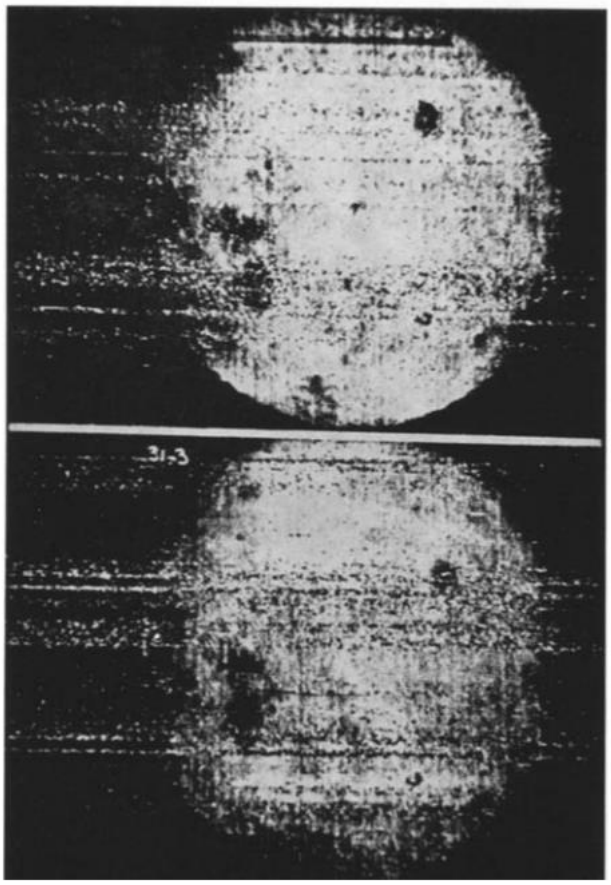

Figure 3.1: The first spaceborne images of another body, and first ever images of the Moon's far side, captured by Luna 3's camera system. Reproduced from Heacock (1968).

The early '60s saw NASA's Ranger program, whose spacecraft (Rangers 7, 8 and 9) imaged the moon down to resolutions of 0.5 m, revealing small craters (~1 m) to be common, and also providing evidence of erosion processes (Jaffe, 1969). Significantly greater resolution required a camera aboard a landed platform, which was first successfully achieved by the Soviet Union's Luna 9 spacecraft (Heacock, 1968). Once landed, a camera is surrounded by the planet's surface (any and all of which can be scientifically valuable to observe), so Luna 9's camera utilised scan imaging to capture a 360° horizontal FOR with a ~30° vertical FOV. Its 1 mm spatial resolution revealed surface roughness on scales too fine to have been resolved from orbit (Heacock, 1968; Jaffe, 1969). Similar images were obtained by NASA's Surveyor landers I, III, V, VI and VII, described by Heacock (1968). Landed cameras provide high levels of detail, but with coverage restricted to the landers' locales. Large-coverage mapping, desirable for

identifying potential landing sites, was achieved with high resolution orbital cameras, first carried out in the 60s by NASA's Lunar Orbiter program (Elle, et al., 1967). Subsequently, the detail and extent of lunar mapping has continued to improve. In 1994 the Clementine spacecraft globally mapped the Moon in 11 colours, by capturing almost 2 million individual images. Two optical instruments, the ultraviolet-visible (UVVIS) and high-resolution (HIRES) cameras captured frame images through spectral filters in filter wheels. An average surface resolution of 200 m per pixel was achieved (Nozette, et al., 1994). The multispectral data provided new insight into the Lunar surface composition (e.g. Jolliff (1999)). The Lunar Reconnaissance Orbiter (LRO) has been orbiting the Moon since 2009, surveying its surface with a goal of supporting future landed missions (Vondrak, et al., 2010). The LRO camera (LROC) comprises three individual imagers, two narrow-angle push-broom cameras (NACs) and one wide-angle push-frame camera (WAC), providing both large-coverage and high-resolution imaging capabilities that could not be achieved with a single camera (Robinson, et al., 2010). The instrument has globally imaged the Moon at 100 m resolution, allowing the extensive cataloguing of craters, mineralogical mapping, and stereoscopic topography measurement, and is expected to support future landed missions (Petro, 2020).

In 2013 and 2019 respectively, the Chang'e 3 (Li, et al., 2015) and 4 (Jia, et al., 2018) missions placed landers and rovers on the Lunar surface, supported by high-resolution global mapping carried out by the earlier Chang'e 1 and 2 orbiters. An array of cameras on each mission captured surface images with decimetre resolution, 360° panoramas, and stereoscopic images for topographic measurement. Images captured during the descent of each lander were used to produce landing site DTMs and precisely constrain crafts' touchdown locations (Liu, et al., 2015; Liu, et al., 2019).

3.1.2 Mars

Mars has seen a similar history and progression of imaging instruments to that of the Moon. In 1965, Mariner 4 captured 22 images, from distances of 17,000 to 12,000 km, during its flyby of Mars. These images were the first of the Martian surface to be captured from a spaceborne camera (Chapman, et al., 1968; Leighton, et al., 1965). Covering ~1% of the surface with a resolution down to 4 km per pixel, they revealed a heavily cratered, old surface, with no evidence of recent ocean or dense atmosphere. In 1971, Mariner 9 (McCauley, et al., 1972) became the first spacecraft to orbit Mars, followed closely by the Mars 2 and Mars 3 spacecraft (Perminov, 1999). From their orbits, their framing cameras captured images of Mars with resolution down to 100 m, showing impacts, volcanism, tectonism, erosion and deposition all to be significant factors in the

shaping of Mars' surface. Unexpectedly, a major dust storm obscured most of the surface from the orbital cameras during the missions' early stages (Thorpe, 1973)).

For Mars, as with the Moon, landed cameras followed soon after orbital cameras. Mars 2 and 3 both deployed landers to the surface, the former of which crashed, whilst the latter became the first successful soft landing on Mars. Its camera failed to return any meaningful images however (Perminov, 1999). Instead, the first image ever captured from the surface of Mars came five years later, in 1976, from the Viking 1 lander. Following the convention of previous planetary landers, its two rotating scanning cameras (one monochrome, one colour), allowed maximum surface coverage from the lander's fixed position. Its images, which resolved down to 1 cm, showed loose rocks with a variety of textures, sediment dunes, and included stereoscopic products (Mutch, et al., 1976). The Viking 2 lander landed 6500 km away, and with the same camera hardware observed a different, more uniform terrain with fewer loose rocks, demonstrating a diversity of the surface at small scales, and the value of multiple, distributed camera stations (Shorthill, et al., 1976).

The Viking landers each accompanied an orbiter, both carrying two framing cameras (called the Visual Imaging Subsystem) for continuous imaging of the surface (achieved by alternately exposing one camera whilst the other reads out) (Wellman, et al., 1976). With a resolution on the surface of <100 m per pixel, and a six position filter wheel, the Visual Imaging Subsystems returned detailed images of fluvial erosion, layered deposits, ice at the poles and clouds in the atmosphere. The orbital cameras also mapped the entire Martian surface at <1 km per pixel (Carr, et al., 1976).

Similarly to imaging of the lunar surface, the 21st century has so far seen the near continuous use of orbital cameras to provide both global coverage and sub-metre resolution imaging of the Martian surface. The Mars Orbiter Camera (MOC) of the Mars Global Surveyor (MGS) operated from 1997 to 2006, providing push-broom images with scale between 1.5 and 20 m per pixel with global coverage (Malin & Edgett, 2001). Since 2004, the High Resolution Stereo Camera (HRSC) aboard the Mars Express orbiter has been imaging the surface with nine push-broom sensors, allowing colour, stereo and multi-phase-angle observations to be simultaneously captured at high resolution (down to 10 m per pixel), eliminating temporal variability in observing conditions. It has imaged more than 25% of the Martian surface (Jaumann, et al., 2007). Similarly, the High Resolution Imaging Science Experiment (HiRISE) has been imaging Mars from the Mars Reconnaissance Orbiter (MRO) since 2006 (Keszthelyi, et al., 2008). The most powerful camera ever flown to another body in the solar system, HiRISE utilises push-broom

imaging in three different spectral bands to provide high fidelity images (down to 25 cm resolution) for a small portion (~5%) of Mars' total surface. Stereo HiRISE observations allow topographic measurement with sub-metre vertical accuracy (McEwen, et al., 2007).

Unique to Mars is the large number of highly successful rovers which have been landed on its surface. The value of landed cameras is their very high spatial resolution relative to orbital data, but their coverage is orders of magnitude smaller. Cameras on rovers, as opposed to static landers, mitigate this issue somewhat thanks to their ability to change position and perspective. Coupled with this is the ability to reposition the rover in response to previous images, allowing a targeted investigation of the surface. Mars Pathfinder's Sojourner was the first Martian rover, landing in 1997, and was equipped with two stereo monochrome cameras and a colour camera at its front and rear respectively (Matijevic & Shirley, 1997). These wide angle cameras were used mainly for the rover's autonomous navigation, but also supported other goals, such as soil analysis. The Mars Exploration Rovers, Spirit and Opportunity, landed in 2004, each carrying 10 cameras. Included in these were two identical framing cameras mounted in a stereo viewing geometry on a rotating mast, allowing multispectral stereoscopic and panoramic images to be captured (Bell III, et al., 2003). The cameras provided insight into the distribution of rocks, dust and fine grain deposits, atmospheric opacity, and even made observations of Mars' moons Phobos and Deimos (Bell III, et al., 2004; Bell III, et al., 2004).

In 2012 the Mars Science Laboratory's Curiosity rover landed in Mars' Gale Crater, carrying on it 17 different cameras (Malin, et al., 2017). The majority of these cameras are dedicated to navigation and hazard identification, with a variety of forward, rearward and downward views. The rover's scientific imagers include two Mast Cams, mounted on a head capable of rotating a full 360° and tilting through ~180°. The high resolution images returned by these cameras have provided insight into the geological history and composition of Mars' surface along a ~25 km long traverse, something which can only be achieved with a moveable camera platform. The Mars 2020 mission builds on the success of the Curiosity rover, with the similarly designed Perseverance rover. Its Mastcam-Z instrument, also mounted on a rotatable tilting platform, comprises two identical cameras arranged with a ~24 cm stereo baseline, and equipped with 4:1 zoomable optics (Bell III, et al., 2021). Using filter wheels with a range of broadband, narrowband and solar filters, the focusable, zoomable cameras provide unprecedented flexibility to optimally image a wide variety of Martian scenes. This flexibility is important on a rover that's expected to operate for years, with serendipitous discoveries and unexpected conditions wholly possible. The rover is accompanied by the Ingenuity

helicopter technology demonstration, which carries two small cameras (Balaram, et al., 2021). Mounting cameras on an aerial vehicle increases the area over which images can be acquired, reduces traversal time, and permits the control of not just the horizontal position, but also the vertical position of the imager. As spacecraft technology continues to improve, new imaging platforms like Ingenuity will provide an important opportunity for novel and valuable planetary imaging.

3.1.3 Venus

Venus is of great interest to planetary science given its many similarities to Earth (size, heliocentric distance, presence of atmosphere) paired with significant differences (high surface temperature and pressure, lack of magnetic field, dominance of carbon dioxide in the atmosphere) (Basilevsky & Head, 2003). Yet images of Venus' surface are scarce in comparison to Mars and the Moon. The first flyby of the planet was performed in 1962 by the Mariner 2 spacecraft, but it was not equipped with any optical imaging instruments. Even if it had been, capturing images of the surface would not have been possible due to the planet's thick atmosphere, which wholly obscures its surface.

The first surface images were therefore not obtained until 1975, when the Venera 9 and Venera 10 landers touched down (Keldysh, 1977). The two identical landers captured ~180° panoramic images of their surrounding surface with rotating push-broom cameras, revealing rock fragments likely originating from lava and distant impact ejecta (Basilevsky, 2012). The harsh Venusian surface conditions limited each lander's lifespan to less than an hour. In 1982, Venera 13 and 14 landed on Venus and captured similar images to those of 9 and 10. Utilising two cameras per lander, panoramas with close to ~360° horizontal extent were captured. The cameras of these latter two landers also included colour filters for the construction of multispectral images (Florensky, et al., 1982).

Visible observations of Venus' surface are, at the time of writing, limited to those that were captured by Veneras 9, 10, 13 and 14. Surrounded by a protective shell, the landers were unable to image the surface during their descents (in any case, the rotational push-broom cameras would have been ineffective during the turbulent parachute descents), preventing the acquisition of larger-coverage images. As a result, only four highly-localised regions have been imaged, and significant scope exists for valuable future observations. Rover-based imaging of the type performed at Mars is unlikely, due to the challenge of operating hardware on the surface for an extended period. Novel imaging approaches and platforms may prove valuable in future studies of the planet's surface.

3.1.4 Titan

As with Venus, the surface of Saturn's moon Titan is optically obscured from observations by its thick, hazy, nitrogen-rich atmosphere. The first optical observations of its surface were made in 1980 with Voyager 1's Imaging Science Subsystem (ISS), comprising both a narrow angle and wide angle camera, as the spacecraft performed a flyby of Saturn (Smith, et al., 1977). These images were captured at wavelengths 590-640 nm, over which the surface signal was heavily contaminated with atmospheric scattering (Richardson, et al., 2004). To overcome this, the next spaceborne images of the surface, obtained by the Cassini spacecraft's ISS in 2004, were captured at longer wavelengths (e.g. 938 nm), at which the atmosphere is known to be optically thinner (Porco, et al., 2005). They achieved resolutions down to ~0.5 km at best, revealing evidence of tectonic, fluvial and aeolian processes.

The Cassini mission had a second method by which to obtain images of the obscured surface: the Huygens descent probe. The probe was deployed from the Cassini spacecraft and, in 2005, descended by parachute to Titan's surface (described by Lebreton et al. (2005)). The aerodynamic design of the lander's structure intentionally caused it to spin as it descended, and three framing cameras repeatedly captured images to build up full coverage of the visible surface using the push-frame technique. Clear views of the surface were achieved from an altitude of 40 km, all the way down to the probe's landing. The images provided strong evidence for fluvial activity, such as the abundance of centimetre-scale rounded pebbles at the landing site. Surface topography of the landing site and surrounding area was measured by applying stereophotogrammetry to the probe's descent images (Soderblom, et al., 2007).

3.1.5 Jupiter

Jupiter has been visited by a number of spacecraft, starting with Pioneers 10 and 11 in 1973 and 1974 respectively (Fountain & Gehrels, 1977). These flyby spacecraft were spin stabilised, and imaged Jupiter with their Imaging Photopolarimeters (IPP), which operated in push-broom mode (with scan motion provided by spacecraft spin) (Pellicori, et al., 1973). Scanning over multiple rotations, the cameras were able to build up images of Jupiter's full disk. They revealed cloud structures in greater detail than ever before, and views for the first time of the planet's north and south poles. Similar flybys of the planet were performed by Voyagers 1 and 2 in 1979 (Smith, et al., 1979b; 1979a). Each spacecraft captured thousands of images of Jupiter during approach and departure, using the narrow and wide angle cameras of their ISS. Their sequences of images were

used to study the formation and evolution of Jupiter's complex cloud structures, and also revealed Jupiter to possess a ring system.

The cameras of Voyagers 1 and 2 also collected observations of many of Jupiter's moons during their flyby of the planet, possible because the framing cameras' pointing could be accurately controlled by the 3-axis stabilised spacecraft. Significant portions of the Galilean moons' surfaces were mapped, at resolutions reaching below 100 km per pixel. Observations included volcanic eruptions on Io, large multi-ring structures on Callisto and evidence of tectonism on Europa and Ganymede. A number of significantly smaller Jovian moons were also discovered and resolved for the first time (Smith, et al., 1979b; 1979a).

In 1995, the Galileo mission placed the first spacecraft in Jupiter's orbit. The Galileo orbiter operated until 2003 (when it was crashed into Jupiter's atmosphere), during which time it captured thousands of images with its Solid State Imager (SSI), a narrow angle framing camera (Belton, et al., 1996b). The SSI used a CCD, improving significantly the spectral range and sensitivity of observations over those of the Voyager cameras and their vidicon sensors. Image resolutions as low as ~50 km were achieved, allowing detailed study of the 3D structure of the planet's Great Red Spot, velocity measurements of Jovian winds, and the detection of lightning within multiple storms (e.g. Little et al. (1999)).

The long duration of Galileo's operation provided ample opportunity to acquire new images of the Jovian moons. A tour of the satellites was conducted through a series of 10 highly elliptical orbits, each one bringing the spacecraft close to a Galilean moon (D'Amario, et al., 1992). Io's volcanic plumes, often extending 100 km from its surface, were observed with 1-2 km per pixel resolutions (Keszthelyi, et al., 2001). Resolutions of the moons' observations were significantly improved, reaching ~10 m per pixel (e.g. Greeley et al. (2000)). A large number of Galileo's observations were suitable for stereoscopic imaging, permitting the measurement of the moons' surface topographies (e.g. Steinbrügge et al. (2020)).

Five months before its arrival at Jupiter, the Galileo spacecraft deployed a probe on a ballistic trajectory into Jupiter's atmosphere. On December the 7th 1995, the probe entered the atmosphere, where it collected and transmitted various measurements for just over an hour as it descended (Young, 2003). The small probe did not carry any cameras, but had it done, it would have had the opportunity to image the complex cloud structures of the upper atmosphere from a unique and changing perspective (though the short lifetime of the probe may have prohibited the transmission of images).

In 2000, Cassini, on its way to Saturn, performed a flyby of Jupiter that lasted several months. Its ISS captured ~26,000 images of the Jovian system over this period (Porco, et al., 2003). Images of Jupiter's atmosphere were used to measure the velocities of Jupiter's zonal winds, observe lightning within storm clouds (Dyudina, et al., 2004), and image Jupiter's northern hemisphere aurora (made possible by the instrument's comprehensive set of spectral filters). Aside from Jupiter, Cassini's ISS was also able to make high resolution measurements of many of the planet's moons, detecting emissions originating from Io and Europa, and observing the shape of previously unresolved moon Himalia. Images of the planet's main ring helped to constrain its phase function in the ~45°-135° range (phase angles that were not observed by the Voyager or Galileo spacecraft).

In 2016, Juno became the second spacecraft to orbit Jupiter (Bolton, et al., 2017). As of 2021, it remains in operation around the planet. Juno's original design did not include a camera, and driven by its science payload and energy requirements, the spacecraft is spin stabilised. During the mission's development, it was decided that sending a spacecraft to Jupiter without a visible imager would be a missed opportunity, and that a camera should be included for the sake of public engagement and education. JunoCam was therefore added to the mission design. The camera uses push-frame imaging, scanning with the spacecraft's spin motion (similarly to the Pioneer cameras), and four spectral filters to obtain colour images (Hansen, et al., 2017). JunoCam's use of push-frame imaging was adopted not because it is necessarily the best way to image Jupiter's atmosphere, but because it was imposed by the spacecraft design.

The concept of JunoCam relies heavily on citizen science. Amateur telescope images assist with observation planning, and processing of raw data is carried out by the public. The public engagement aspect of the instrument has been hugely popular, demonstrating the value that images have in engaging people with space missions. And although the camera was included for outreach purposes, its images of Jupiter's atmosphere, captured from unique perspectives thanks to the spacecraft's polar orbit, have also been of high science value. Images of the poles with resolutions ~50 km have been used to map the wind speeds and evolution of Jupiter's complex circumpolar cyclones (Orton, et al., 2017; Tabataba-Vakili, et al., 2020). Juno's low perijove (~3000 km above cloud tops) allows JunoCam to image Jupiter's equator at unprecedented resolutions of ~5 km per pixel. Sánchez-Lavega et al. (2018) used high resolution images to study the velocities and morphologies of the clouds in Jupiter's great red spot, and Aplin et al. (2020) discuss the possibility that JunoCam has detected lightning flashes.

3.1.6 Saturn

In 1979, NASA's Pioneer 11 flew by Saturn, capturing the first in situ images of the planet, its rings and its larger moons (Gehrels, et al., 1980). In doing so, it discovered the F ring (at 800 km wide, it is too narrow to resolve from Earth). It allowed the first observations of the rings at large phase angles. Through optical thickness measurements, it was inferred that unresolvable structural variations in ring thickness were present. Well defined imaging geometry permitted the measurement of the ratio of Saturn's polar to equatorial radii, linked to its rotation and internal structure. Images of Saturn's atmosphere revealed jet streams, whilst spectral and polarimetric imaging revealed that gasses within clouds extend to higher altitudes than particulates.

A similarly broad range of observations was possible when the Voyager spacecraft encountered Saturn in 1980 (Voyager 1, Smith et al. (1981)) and 1981 (Voyager 2, Smith et al. (1982)). High resolution images of the upper atmosphere revealed fast zonal winds, measured to reach 500 ms^{-1} at the equator, along with convective structures. Comparison of Voyager 1 and Voyager 2 observations, separated by 9 months, revealed evolution of the atmosphere's bands. Saturn's largest moon, Titan, was observed to have multiple thick haze layers, which varied temporally and spatially. Observations of the surfaces of moons Mimas, Enceladus, Tethys, Dione, Rhea and Iapetus resolved geological features, and revealed diverse surfaces, exhibiting a variety of ages and processes. The A, B and C rings were found to contain hundreds of ringlets, confirming the presence of small scale structure that had already been inferred from Pioneer 11's images. Radial structures of micrometre sized particles, called spokes, were observed forming and evolving in the B ring, and multispectral observations revealed significant variation in the rings' spectral reflectance with radius.

After its flyby with Jupiter, Cassini arrived at Saturn in 2004, where it remained in orbit until 2017. Over this time it made extensive observations of the Saturnian system with its ISS (Porco, et al., 2005a; Porco, et al., 2005b). Multispectral mapping of cloud morphology expanded on the atmospheric observations of the Pioneer and Voyager missions, revealing active vortices and an asymmetry in the winds of the northern and southern hemispheres (Vasavada, et al., 2006). Storms were imaged in the atmosphere, and found to correlate with Saturn's electrostatic discharges, suggesting lightning as their origin. Satellites, such as Methone, Pallene and Polydeuces were discovered with long exposure images, and sub-kilometre resolution imaging of the rings revealed new structures within, providing information about the nature of gravitational and kinematic interaction between ring particles and satellites. Stereo imaging was employed on

Cassini's images to derive topographic products for a number of Saturn's moons (e.g. Giese et al. (2006)).

3.1.7 Pluto and Charon

Pluto has an orbit with a semi-major axis of 40 AU, making it difficult to image from Earth, or reach with a spacecraft. Hubble Space Telescope observations of the dwarf planet have a best spatial resolution of approximately one fifth the diameter of Pluto's disk, revealing only the largest of surface features (Buie, et al., 2010). The quality of Pluto image data was drastically improved in 2015 when the New Horizons spacecraft became the first ever to visit the dwarf planet. It performed a flyby of Pluto, with a closest approach of 13,700 km. The encounter is described by Stern et al. (2015).

Visible images captured by two cameras – the Ralph instrument and the Long Range Reconnaissance Imager (LORRI) – covered almost the entire of Pluto's surface between latitudes 90° north and 30° south. The former of these instruments images onto multiple detectors using a mixture of push-broom and framing imaging (Reuter, et al., 2008), whilst the latter is a framing camera (Weaver, et al., 2020). New Horizon's images of Pluto, which at their best resolution had a scale of <100 m per pixel, revealed the surface to have some of the widest variation in radiance factor (reflectance) of any solar system body. Broadband spectral filters also revealed significant spatial variation in the surface colour. Images of Pluto's surface provided evidence of water ice, tectonic activity, glacial flow and possibly also wind. Numerous global haze layers were also observed (Gladstone, et al., 2016).

The New Horizons flyby also provided the spacecraft ample opportunity to observe Pluto's moon, Charon. The moon's radius was measured to be 606 km, more than half Pluto's radius of 1190 km. The moon's young and geologically complex surface displayed many craters alongside evidence of extensional tectonics. Stereophotogrammetry was applied to images of both Pluto and Charon and used to derive localised topographic surface maps of each body (Moore, et al., 2016). New Horizons' images allowed the diameters of Pluto's small moons Styx (~10 km), Nix (~40 km), Kerberos (~10 km) and Hydra (~40 km) to be measured for the first time, and also constituted the first resolved images of these bodies (Weaver, et al., 2016).

3.1.8 Small Bodies

A significant number of asteroids and comets have been visited and imaged by spacecraft. Small body missions differ from the missions described above in that they

generally provide the first and only resolved images of their targets (given that they are too small to resolve from Earth). And whilst planets' disks can be resolved by spaceborne cameras at large distances, missions to small bodies have to get significantly closer if surface details are to be discerned. Furthermore, the gravitational field of a planet is often used to assist with establishing an orbit, allowing an extended period of observations from a variety of viewpoints. This is not possible with asteroids and comets due to their low masses, making it more challenging to achieve a rendezvous, and reducing the flexibility with which observations can be made.

In total, six comets have been imaged at close range by spaceborne cameras capable of resolving the solid nucleus. In chronological order, these are comets Halley (imaged by Giotto and Vega 2), Borrelly (Deep Space 1), Wild 2 (Stardust), Tempel 1 (Deep Impact and Stardust-NExT), Hartley 2 (Deep Impact's extended mission EPOXI) and Churyumov-Gerasimenko (C-G, imaged by the Rosetta mission). The missions and instruments involved in the imaging of these comets are described in detail in chapter 4, sections 4.3 and 4.4 respectively. This section therefore focusses on summarising spaceborne asteroid imaging.

The methods and cameras employed to image asteroids from spacecraft are functionally the same as those used for imaging comet nuclei. The first resolved images of an asteroid came from the Galileo spacecraft and its framing camera in 1991. The asteroid, Gaspra, was revealed to have an irregular shape and a large number of relatively evenly distributed craters (Veverka, et al., 1994). Galileo's images, which covered 80% of the asteroid's surface, allowed the derivation of a shape model for Gaspra (Thomas, et al., 1994). In 1993, Galileo passed closely by another asteroid: Ida, and performed very similar observations, collecting 96 images covering 95% of the surface. From the images, the asteroid was found to have a small satellite: Dactyl. As with Gaspra, the images captured during the flyby facilitated the construction of a shape model of the asteroid (Belton, et al., 1996a).

In 1997, the 3-axis stabilised Near Earth Asteroid Rendezvous (NEAR) spacecraft passed asteroid 253 Mathilde, capturing images with resolution 0.006° per pixel through its multispectral imager (MSI). The disk of the asteroid filled a region of approximately 200 x 200 pixels in the MSI field of view, permitting the detection of large craters as well as smaller features such as faults and layers (Thomas, et al., 1999). After this encounter, NEAR went on to establish an orbit around asteroid 433 Eros. Observing from low altitude (and eventually performing a hard landing on the asteroid's surface in 2001) the spacecraft was able to capture the highest resolution images of an asteroid surface that

had yet been obtained. Images with pixel scales as low as 4 m per pixel revealed unexpectedly-steep slopes, craters, tectonic features and loose boulders (Cheng, et al., 2002).

Prior to its flyby of comet Borrelly, Deep Space 1 encountered asteroid 9969 Braille in 1999, making a closest approach of 26 km. A tracking issue during the encounter meant that the narrow field of view of the spacecraft's MICAS camera lost sight of the asteroid, and only two resolved images of the object were successfully obtained and transmitted (Oberst, et al., 2001).

On its way to comet Wild 2, the Stardust spacecraft performed a distant flyby (closest approach 3000 km) of asteroid 5535 Annefrank. The asteroid was revealed to be irregularly shaped, but images were too low-resolution and captured for too short a duration to characterise surface brightness or construct a shape model (Duxbury, et al., 2004).

JAXA's Hayabusa spacecraft, which rendezvoused with asteroid 25143 Itokawa in 2005, was the first to perform a soft landing on an asteroid, and the first to retrieve and return a sample of the asteroid to Earth. The spacecraft's landing descent, and a number of practice descents, allowed the surface to be imaged at centimetre resolutions, revealing fine-scale regolith and boulders covering the surface's lowlands (Fujiwara, et al., 2006). Following on from the success of Hayabusa, the Hayabusa2 mission performed a similar rendezvous and landings with asteroid 162173 Ryugu in 2018. The asteroid's surface exhibited a similar rubble-pile composition to that of Itokawa, with centimetre-scale descent images revealing fine-grained regolith. Global-coverage imaging of the surface from low altitude facilitated the construction of a shape model with high spatial resolution, characterising the distribution of the asteroid's boulders and scarce craters (Watanabe, et al., 2019).

In 2020, NASA collected a sample of material from asteroid 101955 Bennu with its OSIRIS-REx (Origins, Spectral Interpretation, Resource Identification, Security-Regolith Explorer) spacecraft. Prior to sample collection, OSIRIS-REx orbited and observed the asteroid for almost two years, capturing extensive images of its surface with three optical cameras (together constituting the OSIRIS-REx camera suite, OCAMS). These framing cameras mapped the asteroid globally at a range of scales, down to ~0.5 m per pixel (Barnouin, et al., 2020). This global mapping was essential for deriving sufficiently accurate topographic maps for choosing a safe sampling site. OCAMS' narrow-angle PolyCam (0.8° wide FOV) was capable of detecting Bennu from millions of kilometres away, but thanks to a mechanical focus mechanism, was also able to image the

asteroid's surface from distances as small as 200 m, achieving millimetre resolutions (Rizk, et al., 2018).

3.2 Descent Imaging of Planetary Surfaces

Descent images constitute a small portion of the image data that have been collected by planetary exploration missions, and exist for only a handful of the solar system's bodies. However, descent images provide an opportunity to observe a body from a unique and changing perspective, that neither orbital nor landed images can achieve. The transition from high altitude to surface-level means a sequence of images can be obtained which bridges the gap between, and complements, orbiter and lander observations.

Malin et al. (2001) describe the value of descent imagery in planetary science. Planetary surfaces are shaped by geological processes which occur at a range of scales from many kilometres to a few millimetres. The relative importance and abundance of these processes is dependent on the specific conditions present on the solar system body. The identification of the geological hallmarks of these processes, which themselves occur at all scales, is therefore greatly assisted by the availability of images with a range of resolutions and extents. Malin et al. (2001) demonstrate this with illustrative Earth data. Their Figure 1a shows an orbital image of Antarctic terrain, whilst their Figure 1b shows a high altitude image, representative of one captured at the beginning of a spacecraft descent. The footprint outlines of subsequent, lower altitude images are superposed. Their Figure 2 shows those subsequent images, captured at ever decreasing altitude. Both their Figures 1b and 2 are reproduced below in Figure 3.2.

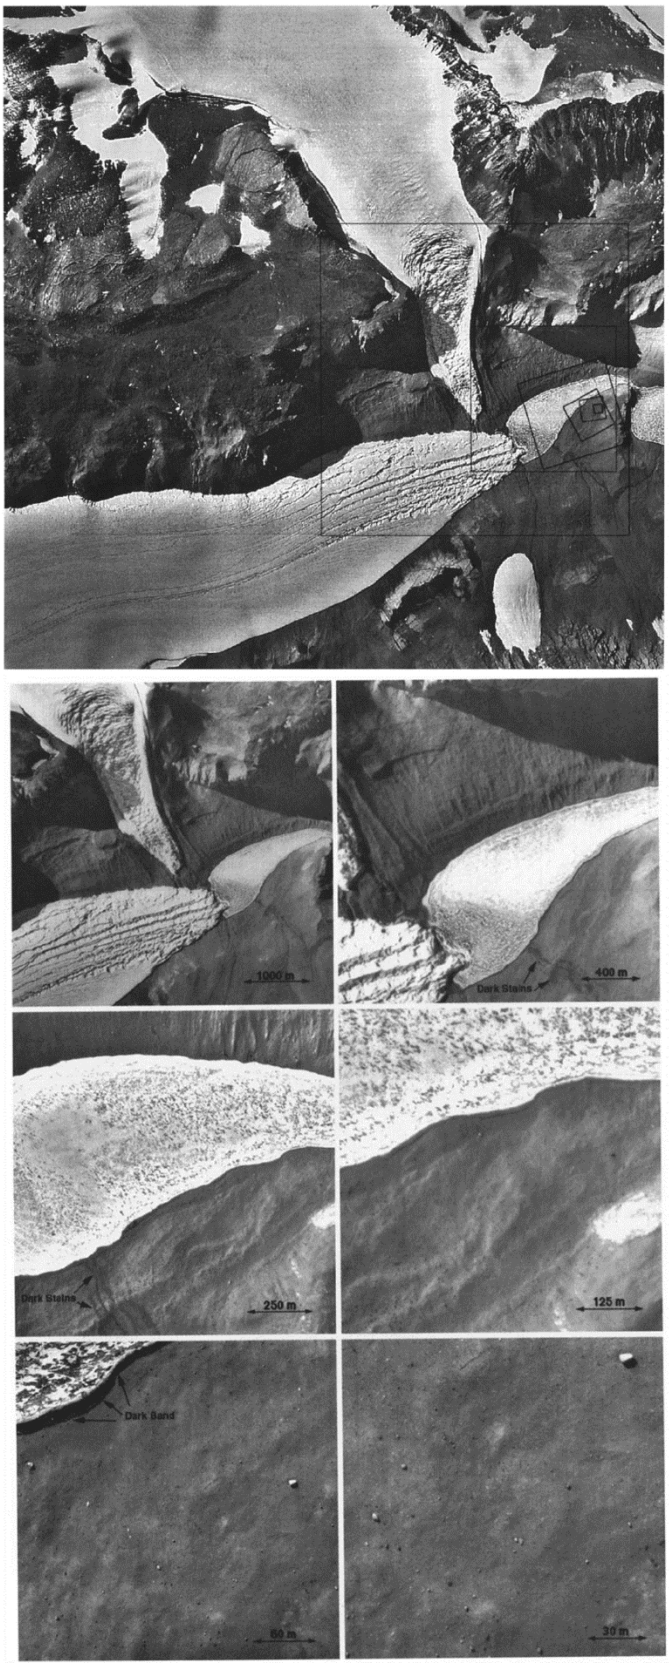

Figure 3.2: Seven images from a simulated descent, with the boundary of each image indicated in the first of the sequence (top). Reproduced from Figures 1b and 2 of Malin et al. (2001).

The first, high altitude image shows the mountainous nature of the terrain, and the presence of multiple glacial flows. Meanwhile, the last two images of the descent sequence reveal deposited material which cannot be resolved from the high altitude images, and indicate the nature of the contributing surface processes (Malin, et al., 2001).

In addition to the direct geological analysis which descent images permit, they also provide the value of landing site context. In isolation, in-situ measurements of the surface or sub-surface are of course valuable, but their interpretation and usefulness is enhanced when they can be placed in context by a wider knowledge of their surroundings and location. This allows, for example, the assessment of how representative the landing site conditions are of the global surface.

Descent images, thanks to their nested footprints and incrementally-increasing resolution, can be used to determine landing site location with accuracies down to less than a metre (Bos, et al., 2018; Karkoschka, et al., 2007; Li, et al., 2002; Liu, et al., 2015; Liu, et al., 2019). Importantly, because of the gradual change in image footprint and resolution, each image in the sequence can be located accurately within the sequence's previous image, and by extension the lowest altitude image can be tied to the highest altitude image. This method can outperform the localisation accuracies possible from the combination of just orbital and landed images (Malin, et al., 2001).

Descent images can also perform the valuable engineering function of recording the motion of the spacecraft as it travels to the surface. This can assist with recovering the platform's lateral motion, rotation, orientation and their variation throughout the descent. These are valuable tools for inferring atmospheric conditions (if an atmosphere is present), for informing modelling of future descent missions (Karkoschka, et al., 2007; Lorenz, 2010), and even for navigating the spacecraft during its descent (Frauenholz, et al., 2008). The unique nature of descent images, which give an onboard perspective of the spacecraft's journey, also makes them an ideal resource for enthusing the public.

The earliest descent images were those of the Ranger impactors, captured by their television camera systems (Schurmeier, et al., 1966). The sequences of images they obtained not only provided increasingly detailed views of the Moon's surface, but also told the story of each spacecraft's descent. Furthermore, Ranger 9's descent television images were the first from space to be viewed live.

Descent imaging was employed during the landing stages of many of NASA's Martian rovers. The two Mars Exploration Rovers (Crisp, et al., 2003) were each equipped with a downward-looking descent camera with 45° FOV, and the Mars Science Laboratory

used a similar descent imager (MARDI) with 90° FOV (Malin, et al., 2017). These cameras captured several hundred sequential images during descent to assist with on-board estimation of descent trajectories, and characterisation of the landing sites.

Similar imaging was carried out during the landings of the lunar missions Chang'e 3 (Li, et al., 2015) and Chang'e 4 (Jia, et al., 2018), with each capturing a few hundred images of the Moon's surface surrounding their landing sites. Stereophotogrammetric processing of these descent images was used to produce localised topographic maps of the landing sites, and constrain landing site location with centimetre precision (Liu, et al., 2015; Liu, et al., 2019).

The three cameras of the Huygens probe's DISR (Descent Imager/Spectral Radiometer) captured and transmitted a total of 376 images during its descent to Titan's surface (followed by 224 images after landing). As it descended, imaging resolution scaled from 150 m to 0.4 mm per pixel, with the coverage of the imagers' FOVs scaling similarly (Karkoschka, et al., 2007). The DISR's three cameras (Side Looking Imager (SLI), Medium Resolution Imager (MRI) and High Resolution Imager (HRI)) were mounted such that their FOVs combined to cover a narrow strip of surface extending from 7° to 96° above nadir. The cameras utilised push-frame imaging with the probe's spin motion to achieve 360° of azimuthal coverage. These descent images contrast to those of the above landers because of their large rather than localised coverage.

3.3 Discussion

Since the first planetary cameras of the 1960s, the technology employed in building and operating spaceborne imagers has improved markedly. The fundamental principle behind planetary imaging remains the same however, with a variety of framing, scanning, narrow-angle and wide-angle cameras used today as they have been throughout the last six decades. Indeed, many contemporary cameras are comparable to those of the early missions, with the main differences being better-performing components, leading to improvements such as lower noise and higher resolutions. As an example, Pioneers 10 and 11 imaged Jupiter and Saturn with rotational scanning cameras in the 70s, and the Juno spacecraft employs the same technique fifty years later.

Scan imaging was employed on these spacecraft because their spin stabilisation necessitated it, but in other imaging scenarios, the use of scan imaging was beneficial in its own right. Scanning cameras on the Luna, Surveyor, Viking and Venera landers helped to maximise the surface coverage that could be achieved from static imaging positions. The Huygens probe similarly used rotational push-frame imaging to observe

as much of the visible surface as possible. Planetary mapping also relies heavily on scan imaging, again because it is conducive to collecting continuous and large-coverage observations. However, the need to maintain a certain camera-scene motion reduces the flexibility of a scanning imager, limiting, for example, its viewing geometry and exposure times.

Voyager, Galileo and Cassini demonstrate the advantage of framing cameras' flexibility. The option to point the camera in any direction allows single instruments to collect a wide variety of observations (e.g. a planet, its moons and its rings), supporting a diverse range of planetary science studies. The ability to capture a variety of observations increases the likelihood of serendipitous discoveries, whilst also allowing the investigation of phenomena already known to be present. At Saturn, for example, the Voyager imaging systems made pre-planned measurements of atmospheric wind speeds, whilst in other observations, they discovered new moons and the existence of previously unseen ring structures (Smith, et al., 1981; 1982). High-resolution framing cameras require strict control of spacecraft pointing in order to obtain useable images, whilst lower-resolution wide-angle cameras can be pointed with lower accuracy.

The flexibility of early planetary imaging systems was essentially necessitated by the lack of prior knowledge of what they would encounter (Anderson, et al., 1991). When a planetary system has not previously been imaged, a more general, less-targeted approach to observations is valuable. Subsequent imaging campaigns can then focus on the areas of most scientific interest. The Moon and Mars have progressed farthest along this route, with global mapping at increasing resolutions, descent and landed images, and moveable imaging platforms on the surface all following on from the first observations of the bodies. A large driver of this progression is the desire for ever-higher spatial and spectral resolutions. But at other solar system bodies, where this progression is less advanced, increasing image coverage or achieving new viewing perspectives, even if at low resolutions, brings new value.

With limited funding sources and long development times, opportunities to observe solar system bodies come relatively infrequently. There have nonetheless been missions such as Galileo's atmospheric probe and the Vega Venus Balloon experiment (which placed two balloons in Venus' atmosphere, (Sagdeev, et al., 1986b)) that have delivered probes to locations from which valuable and novel images could be have been acquired, but no cameras were included in the payload. There is already significant diversity amongst the approaches that have been used to image planets, and developing new ways of

capturing images from planetary probes increases the likelihood that future solar system missions can be utilised to perform valuable imaging science.

4 Comets: Observations, Characteristics and Encounters

A significant portion of the work completed for this thesis focusses on a camera whose objective is to obtain images of a comet and its large scale environment from a flyby spacecraft. The results of this work are the subject of chapter 6. In preparation for this, here, in chapter 4, an introduction to comets and cometary science is given, and an overview of the spacecraft and hardware involved in previous cometary imaging missions is provided.

4.1 Historical Overview of Cometary Observations

With their appearance as bright extended structures in the sky, it is not surprising that comets have been observed with fascination by humans for thousands of years. For much of that time, comets were considered bad omens or harbingers of catastrophe (Sagan & Druyan, 1997). Detailed Chinese records of cometary observations span millennia (as far back as the 7th century BC) (Needham, 1959), whilst perhaps one of the most famous recordings of a comet is the c.1066 depiction of Comet Halley (1P/Halley) in the Bayeux Tapestry (Anderson, 1986). Ancient cometary observations have been used in contemporary studies to calculate the orbits of bright comets, such as 1P/Halley, further back in time than is possible from modern observations (e.g. Hasegawa (1979)).

Aristotle was one of the first to use observations to develop a general theory to describe the phenomenon of comets. In his theory, Aristotle treated comets as atmospheric phenomena (Heidarzadeh, 2008, pp. 1-16), and this view persisted up to the mid-1500s. From the sixteenth century onwards, an increase in the precision with which comets were observed allowed the measurement of their distance from Earth. With this, it could be shown that comets were far beyond the atmosphere, and the theory of comets being meteorological was abandoned. It was recognised that despite their very different trajectories, comets traverse the same space as the planets (Heidarzadeh, 2008, pp. 41-45).

The 17th century saw both Kepler's laws of planetary motion and Newton's theory of gravitation. Observations of comets were important in confirming the validity of both,

though Kepler himself did not consider comets to obey his laws (Heidarzadeh, 2008, p. 67). At the same time, Edmond Halley noted that three comets, observed in 1531, 1607 and 1682, shared very similar orbits, and concluded they were in fact the same comet (known now as comet Halley). With this, he became the first to prove the periodicity of comets, and the first to predict a comet's return (Hughes, 1987).

The orbital mechanics of comets were now better understood, but much about them still remained a mystery. That comets have bright heads and extended tails was well observed (including the anti-sunward orientation of what are now known to be ion tails, and the trailing orientation of dust tails), but with little knowledge of the composition of comets, the formation of the tails was poorly understood.

In the 19th century, Friedrich Bessel observed a bright jet protruding from the head of 1P/Halley in the sunward direction. He concluded that it must be driven by the Sun, and that the nuclei of comets must not be solid bodies, given their volatility (Heidarzadeh, 2008, pp. 154-159). Spectroscopy of comets began in 1864, and with it, comets' compositions became better understood. Huggins (1882) identified the presence of carbon, hydrogen and nitrogen in comet C/1881 K1 (Tebbut). By the early 1900s, cyanide and sodium were also identified as cometary constituents (Heidarzadeh, 2008, p. 235).

Whipple (1950) formalised a cometary model still largely consistent with today's observations and understanding of comets, in which cometary nuclei are solid, and mainly comprise various volatile ices (e.g. H_2O, NH_3, CH_4, CO_2), which are mixed with dust (non-volatile material). Vaporisation of the ices, leading to ejection of gas and dust from the nucleus as the comet approaches the Sun, explained the production of comets' comae, tails and jets, and accounted for non-gravitational forces which had been observed to alter periodic comets' orbits.

The model of comets went hand in hand with the advancement of theories on their origin. It was understood that, since comets' nuclei lose material when near the Sun, they must have a finite life (Whipple (1950) estimated 3000-60,000 years). Öpik (1932) had already discussed the idea of a reservoir from which observable comets are regularly replenished, and Oort (1950) proposed the existence of a large cloud of comets (now known as the Oort cloud) surrounding the outer solar system.

4.2 The Contemporary Understanding of Comets

Comets represent the leftovers from the Solar system's formation, comprising material that was present in the early solar nebula and involved in the formation of the planetesimals that later became the outer planets (e.g. Bockelée-Morvan (2011), Glassmeier et al. (2007), Mannel et al. (2016)). Their formation occurred beyond the snow line, where temperatures are low enough for water ice to condense (approximately 5 AU from the Sun). Many comets were ejected by Jupiter to the outer solar system, where they still reside. At these distances, the Sun's radiation has very little effect on them, and they are cold, dark, solid bodies whose state has gone largely unchanged since their formation ~4.5 billion years ago. For this reason, the study of comets' physical characteristics is of high interest to those wishing to probe the conditions of the early solar system.

4.2.1 Cometary Orbits

The vast majority of the Solar system's cometary nuclei exist within the Oort cloud (Weissman, 1990). This spherical shell surrounds the solar system and occupies heliocentric distances in the range of ~5000 to ~100,000 AU (stretching approximately a third of the way to the Sun's nearest neighbouring star) (Stern, 2003).

The comets within the Oort cloud are only weakly gravitationally bound to the Sun, and subject to a number of perturbations which can significantly alter their orbits. The gravitational effects of nearby stars, giant molecular clouds and the Milky Way's galactic disk disturb comets' orbits, sporadically ejecting some from the Oort cloud into interstellar space, and directing others in toward the inner Solar system (Weissman, 1990).

Comets entering the inner Solar system via this mechanism have orbits with large eccentricity (and potentially also a high inclination). A small proportion have sufficient energy to escape the solar system after their periapsis (travelling along hyperbolic trajectories), and are known as non-periodic comets. Most in fact become long-period comets (by convention, these are comets with an orbital period >200 years (Levison, 1996)), gravitationally bound to the Sun. Long-period comets' orbits generally have perihelia between 1 and 10 AU, whilst aphelia can be as large as 10^5 AU (e.g. Francis (2005)).

In addition to the Oort cloud, the Kuiper belt also provides a continuous source of comets to the inner solar system. The Kuiper belt is a disk (lying in the ecliptic, as opposed to the Oort cloud's spherical shell) of small bodies, extending from Neptune's orbit (~30 AU) to between 50 and 100 AU (Stern, 2003). In a similar fashion to the perturbations of Oort

cloud comets, the gravitational influence of the outer planets on Kuiper belt comets leads to the redirection of some of them into the inner solar system.

Comets with orbital periods shorter than 200 years are known as short-period, of which 1P/Halley (orbital period 75 years) is one of the most famous examples. These comets are thought to derive mostly from the Kuiper belt, and usually have orbits of low inclination (Duncan, et al., 1988). Many short-period comets' orbits have been significantly influenced by the gravity of Jupiter (and to a lesser extent Saturn), and this is often the method by which long-period comets transition to short-period. Comets which reside near to Jupiter's orbit (~5 AU) are known as Jupiter Family Comets (JFCs). JFCs follow prograde orbits (i.e. in the same direction as the planets) with low inclination (Levison, 1996; Volk & Malhotra, 2008).

Another significant bank of comets is interstellar space, which has been estimated in various studies to contain between 10^5 and 10^{13} interstellar comets (larger than 1 km) per cubic parsec (Cook, et al., 2016). These interstellar comets were mostly ejected from their star systems during planetary formation, and if they passed through our solar system, would be identifiable by their hyperbolic orbits with eccentricity significantly larger than 1. However, the occurrence of detectable interstellar comets passing through our solar system is predicted to be on the order of one a century, and indeed only one unambiguous detection of an interstellar comet has been made: Comet 2I/Borisov in 2019 (Ye, et al., 2020). Just two years earlier, the first confirmed interstellar body – 'Oumuamua – was detected passing through the solar system, but observations showed no evidence of a dust or gas coma, and the object did not resemble a comet (Trilling, et al., 2018). It is expected that, once operational, the Vera Rubin Observatory's Legacy Survey of Space and Time (LSST, formerly standing for Large Synoptic Survey Telescope) will increase the detection rate of interstellar objects, and Cook et al. (2016) estimate that in a best case scenario it could detect up to 10 interstellar comets over its 10 year lifetime. The slim possibility exists that Comet Interceptor could be in the right place at the right time, and have the opportunity to conduct the first ever flyby of an interstellar object.

4.2.2 Nuclei

A comet's nucleus, a small, solid body comprising rock, ices and dust, is the source of all the material constituting its many extended structures. Despite the fact that comets can be some of the largest and most spectacular objects in the night sky, their nuclei are some of the smallest and darkest objects in the Solar system. Their sizes range from one to tens of kilometres (Keller & Kührt , 2020; Meech, et al., 2004), meaning they

cannot be resolved by Earth based telescopes. Their detailed study has therefore only been possible since the first comet flybys in the 1980s. Of the six nuclei that have been imaged by visiting spacecraft (see section 4.3), four have a bilobate nucleus. Their masses are not large enough to achieve hydrostatic equilibrium, and often exhibit irregular shapes.

The low reflectivity (geometric albedos on the order of 0.05 (Fernández, et al., 2003; Sagdeev, et al., 1986a)) of cometary nuclei results from their surface composition of primarily dust and rocks (comprising refractories and organic materials), intermixed with volatile ices. These volatile constituents drive the activity of cometary nuclei when sufficiently heated (providing the source of the comet's coma and tails). As a nucleus approaches the Sun, the increasing heat leads to sublimation of the constituent ices, ejecting gas and dust from the surface into the surrounding area. Activity is generally concentrated in discrete regions on the surface, but the entire illuminated surface of the nucleus can be active at low heliocentric distances.

The forces from ejected material are sufficient to affect the rotation of a comet's nucleus (e.g. Neishtadt et al. (2002)), as well as its orbital trajectory (Yeomans, et al., 2004), and in extreme cases lead to the disintegration of the whole nucleus (Hui & Ye, 2020). Short-period comets, which reach perihelion and undergo periods of activity relatively frequently, have heavily altered and evolved surfaces. Long period comets, which have undergone fewer periods of solar heating, often produce more gas and dust when active. Spacecraft observations of cometary nuclei indicate high surface temperatures (e.g. Emerich et al. (1988)), and limited presence of exposed ices on the surface when comets are within 4 AU of the Sun.

4.2.3 Comae

The material ejected from a comet's nucleus spreads into its surroundings, resulting in an approximately spherically distributed 'atmosphere' of dust and gas known as the coma. The material comprising a coma is in a state of weak gravitational attraction with the nucleus, and some ejected material (mainly large dust particles, whose ejection speeds are lower) eventually falls back to the nucleus surface (Thomas, et al., 2015). Nonetheless, the comae of active comets often extend as far out as 10^4-10^5 km from their nuclei, and can exceed the size of the Sun (e.g. Lallement et al. (2002); Lin et al. (2009)). Because they are driven by nucleus activity, comae are not usually seen when comets are far from the Sun (making the detection of distant comets difficult), but they have been observed at heliocentric distances as large as 23 AU (Farnham, 2021).

Whilst the density of coma material is broadly proportional to r^{-2}, where r is the radius from the nucleus, structures within comae result in localised deviations from this at a range of scales (from 1-100,000 km). In the vicinity of the nucleus, narrow collimated jets of dust are common, and they can range in size from hundreds of metres to hundreds of kilometres (Combi, et al., 2012; Farnham, et al., 2007; Lara, et al., 2015). Larger scale jets and broad fans have also been observed (Farnham, 2009).

Larger, more complex structures can form within comae, resulting from variable dust jet activity, nucleus rotation or a combination of the two. Active regions on a spinning nucleus switch on and off as they cycle from solar illumination to shadow. The periodic variation in dust ejection can produce arcs, shells and spirals, which can extend hundreds of thousands of kilometres from the nucleus (Farnham, 2009; Woodney, et al., 2002). These structures are often very faint compared to the surrounding coma, and require careful processing of cometary images to discern. The nature of these structures provides insight into the rotation state of the nucleus, and the location of their source region.

4.2.4 Dust Tails

Dust tails (formerly known as type II tails), along with comae, are among the largest and most visible components of a comet. Comprising dust that was originally ejected from the nucleus by gas drag, they commonly exceed lengths of 10,000 km and can even extend several AU. Solar radiation pressure causes the dust to lag behind the nucleus in its orbit, resulting in an elongated, curved tail of dust. After its ejection from the nucleus, dust particles spread out due to their range of velocities and masses, leading dust tails to become wider and more diffuse at greater distances from the nucleus (though they predominantly remain in the plane of the comet's orbit) (Fulle, 2004).

Visible due to scattering and reflection of sunlight by the constituent dust particles, dust tails exhibit an approximate solar spectrum, and are not generally as bright as the coma (Brandt, 1968). Dust tails can exhibit temporal variation and fine scale structures at a range of sizes. These include narrow rays of increased brightness emanating from the nucleus and oriented along the length of the tail, known as streamers, and fine-scale striated features, known as striae, which generally appear far from the coma and align neither toward the nucleus nor with each other (Price, et al., 2019; Sekanina, 1996). Such structures are indicative of the properties of the comet's nucleus, including its rotation state and activity, and can also reveal the nature of the Sun's influence on dust tail particles.

4.2.5 Ion Tails

Cometary ion tails comprise gas, originally ejected from the nucleus by sublimation, and subsequently ionised by the Sun (either by photoionisation or solar wind charge exchange). The ionised atoms and molecules are carried away from the nucleus and coma by the solar wind, forming large tails which point away from the Sun (Coates, 2004; Cravens & Gombosi, 2004). The size and brightness of an ion tail depends on the activity of its parent nucleus, and they are generally only seen when comets are at heliocentric distances <2 AU, whilst the tails themselves can reach up to a few AU in length (Gloeckler, et al., 2000; Jones, et al., 2000). Ion tails were formerly referred to in the literature as type I tails, and are also sometimes known as plasma and gas tails (though the latter is potentially confusing, as separate neutral gas tails are also sometimes present in comets (Cremonese, et al., 1997; Fulle, et al., 2007)).

Because ion tails are formed by and are sensitive to the solar wind, they often exhibit more spatial and temporal variation than dust tails, and observing these structures assists with probing the solar wind conditions at the comet (Jackson, et al., 2013). Features observed in ion tails include long narrow rays, twisted structures, disconnections and kinks, all of which can exhibit temporal evolution (Mendis & Ip, 1977).

In contrast to a dust tail's reflected solar spectrum, ion tails' spectra contain the emission lines of their constituent species, often allowing their identification through spectroscopy. Detected cometary ions include CO^+, H^+, H_2O^+, H_3O^+, O^{3+} and S^+, amongst others (Balsiger, et al., 1986; Neugebauer, et al., 2007; Ogilvie, et al., 1986; Wilson, et al., 1998).

4.3 Encounters with Comets

Since the 1980s, large leaps in cometary science have been possible thanks to a number of missions that have closely rendezvoused spacecraft with comets, either by means of a flyby or a prolonged escort. These missions have allowed close range and direct measurements of the nucleus and surrounding environment, and observation perspectives not possible from Earth. This section gives an overview of these spacecraft-comet encounters. The set of missions described here is not an exhaustive list of spacecraft to have observed comets, but focusses on missions whose spacecraft closely approached their target comet, just as Comet Interceptor will. To contextualise the design and objectives of EnVisS, details of these mission's imaging instruments are summarised separately in section 4.4.

4.3.1 The Vega Spacecraft

The most recent apparition of comet 1P/Halley in 1986 marked the beginning of cometary exploration by spacecraft. The Vega 1 and Vega 2 spacecraft were two of eight missions which successfully performed some form of flyby of the comet (Stelzried, et al., 1986). 1P/Halley's retrograde orbit limited the missions to fast flybys. With the Vega 1 and 2 spacecraft's closet approaches of 900 km and 8000 km respectively, only Giotto (see below) flew closer to the comet's nucleus.

Both of the Vega spacecraft were equipped with a camera named the television system (TVS). Between the 4th and 11th of March 1986, Vega 1 and 2 captured and transmitted to Earth ~1500 images of 1P/Halley's coma and nucleus. The primary objective of the TVS instruments was to determine whether cometary nuclei are consolidated bodies or loose material. As the spacecraft neared closest approach, the TVS instruments resolved the nucleus, revealing it to be an irregularly shaped solid body approximately 14 km x 7.5 km x 7.5 km (Sagdeev, et al., 1986a). Additionally, jet structures and the active regions from which they could be originating were identified within the Vega TVS dataset. An example false-colour Vega 2 image, from Sagdeev et al. (1986a), is shown in Figure 4.1.

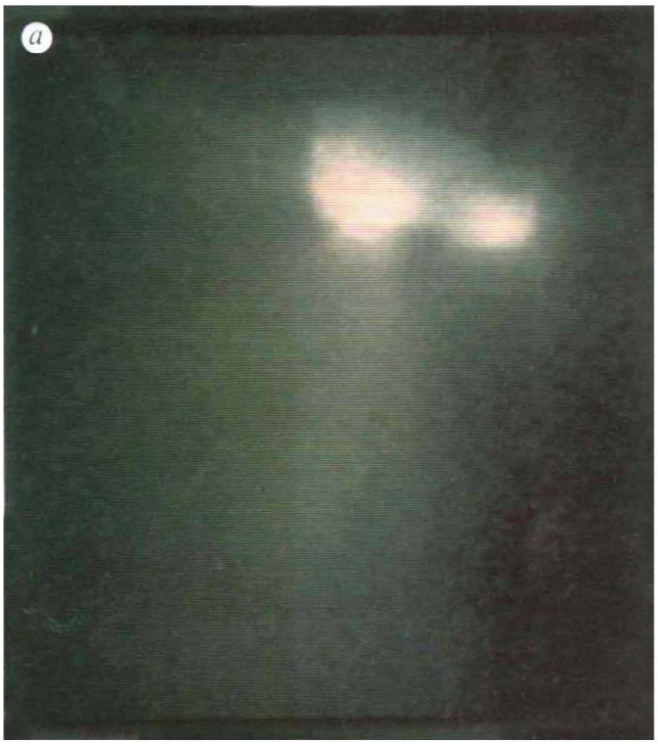

Figure 4.1: False-colour image taken by Vega 2 at a distance of ~8,030 km with a phase angle of 32°. The shape of the nucleus (top right) is clear to see, and dust jets extending toward the bottom left can also be seen. Reproduced from Sagdeev et al. (1986a).

In addition to the TVS imagers, the Vega spacecraft carried dust impact mass analysers, which carried out the first ever in situ dust measurements of a comet. They found the dust of comet 1P/Halley to comprise more low mass particles than predicted, whilst also corroborating, from compositional measurements, the idea that cometary material consists of ices (Kissel, et al., 1986).

4.3.2 Giotto

Giotto is arguably the poster mission of the Halley flybys. It performed the closest approach of any spacecraft visiting the comet, getting within 600km of the nucleus on the 14th of March 1986, and returned observations that were invaluable for advancing knowledge of comets and their formation. The encounter is described by Reinhard (1986).

The spin stabilised spacecraft carried a large scientific payload, including a visible light imager (called the Halley Multicolour Camera (HMC), described in section 4.4.1), a coma brightness experiment (see section 4.4.6), dust detectors and plasma analysers. The images captured by Giotto improved on the resolution of those of the Vega spacecraft

which were captured days earlier, reaching down to 45 m per pixel (Keller & Kührt , 2020). As Keller et al. (1987) point out, relatively little was known about cometary nuclei at the time of the Giotto and Vega missions, and their existence as a solid body was only hypothetical prior to the Halley flybys. A key result of the Giotto mission was therefore the identification of the nucleus' size, shape and surface properties. Figure 4.2 shows an HMC image from Keller et al. (1986) of the seemingly bilobate nucleus and its dust jets, captured when Giotto was ~18,000 km from 1P/Halley's nucleus. Figure 4.3, also from Keller et al., shows one of the last images of the flyby, which captured a jet source on the nucleus from a range of ~2000 km.

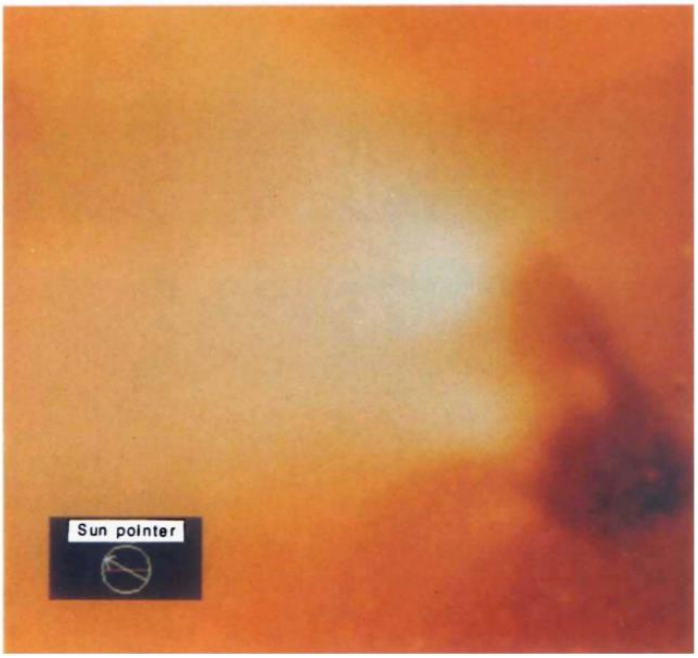

Figure 4.2: The nucleus of comet Halley as seen by HMC from a distance of 18,270km. The frame size is 30km. Reproduced from Keller et al. (1986).

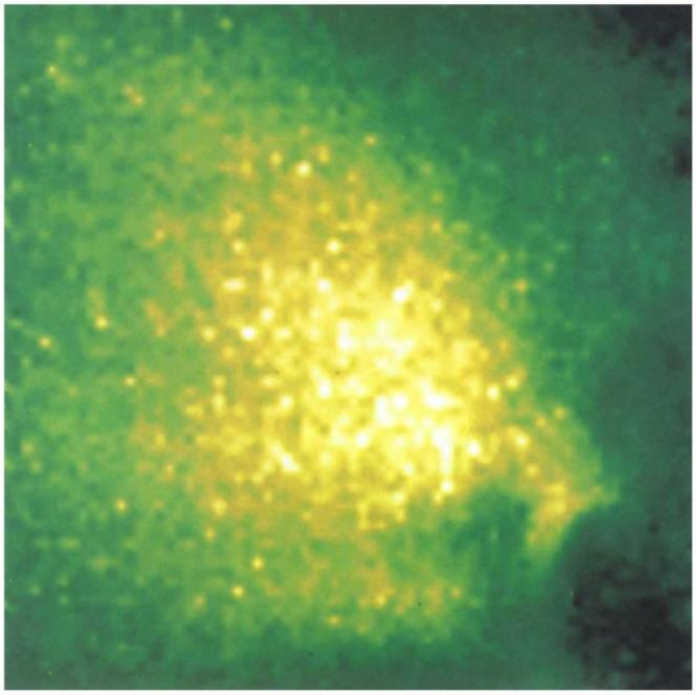

Figure 4.3: HMC image of jet source, taken from 2220 km. The frame size is 3.7km. Reproduced from Keller et al. (1986).

Analysis of the jets within images suggested the dust surrounding the nucleus was optically thin, and the nucleus was found to have a low reflectivity <0.02 (Keller, et al., 1986). Giotto's images of the nucleus indicated that it was dominated by non-volatile dust, interspersed with ices (the reverse of Whipple's model). Whilst Giotto (along with Vega 1 and 2) provided images that far exceeded the detail and resolution of any previous cometary observations, there was still much scope and motivation for higher resolution observations to reveal more about the structure and composition of comet nuclei. Indeed, several missions, discussed in the following sections, have since done this.

Aside from its images, Giotto's magnetometer and plasma analysers conducted the first ever detailed and in-situ study of the spatial structure of a comet's interaction with the solar wind. It detected bow shocks and a magnetosheath that were among the most complex observed in the solar system, and revealed the importance that cometary ions have in shaping these structures (Coates, et al., 1997; Neubauer, et al., 1986).

Due to comet 1P/Halley's retrograde orbit, Giotto passed it with a high flyby speed of 68 kms^{-1}. A dust shield on the spacecraft's front facing surface protected it and its instruments from high speed dust impacts. Nonetheless, approximately 14 s before closest approach, a large dust impact knocked Giotto off its spin axis, leading to

intermittent communications and dust damage to some sensors. Giotto's camera (which used a periscope to view past the dust shield) sustained damage to its baffle and mirror, and became unable to capture further images (Reinhard, 1986).

In spite of this, fuel remaining in the spacecraft after the Halley flyby was used to successfully redirect Giotto, via Earth, to comet 26P/Grigg-Skjellerup. It passed within 200 km of the nucleus on the 10th of July 1992 at a speed of 14 kms^{-1}. Numerous instruments were able to operate at partial or full capacity throughout the flyby (Grensemann & Schwehm, 1993). Optical measurements revealed a difference in the optical properties of dust grains in the inner coma versus jets and the outer coma (Levasseur-Regourd, et al., 1993), whilst impact measurements suggested large particles dominated the dust environment (McDonnell, et al., 1993).

4.3.3 Deep Space 1

The next spacecraft encounter with a comet came in September 2001, when the Deep Space 1 mission performed a flyby of comet 19P/Borrelly (Borrelly), with a closest approach of 2170 km (Boice, et al., 2002). With its MICAS camera (Miniature Integrated Camera and Spectrometer, see section 4.4.2) it captured visible wavelength images and near-infrared spectra of the nucleus and dust coma, whilst its Plasma Experiment for Planetary Exploration (PEPE) performed ion mass spectrometry of the coma.

High resolution images revealed complex dust structures emerging from the nucleus. Both collimated jets and wider fans of dust were observed, along with a transient loop dust feature on the nucleus' night side, which indicated that surface regions could remain active in shadow due to their thermal inertia. From stereoscopic images, the orientations of dust structures were measured. Thus, jets were found to emanate from regions of smooth terrain, and a nucleus rotation period of 26 hours was derived (Boice, et al., 2002). As with Halley, Borrelly's nucleus was found to be very dark (bond albedo ~0.006), with a surface heavily dominated by non-volatile material (Keller & Kührt , 2020).

PEPE's in situ ion measurements exhibited an asymmetry in the ion density distribution relative to Deep Space 1's closest approach. Young et al. (2004) identified this as being consistent with the spacecraft having passed through a dust jet, observed by MICAS, prior to closest approach (Boice, et al., 2002; Nordholt, et al., 2003). PEPE measured the dominant constituents of Borrelly's coma to be O^+, OH^+, H_2O^+, CH_2^+ and CH_3^+ (Nordholt, et al., 2003).

4.3.4 Stardust

A logical next step in cometary exploration was to retrieve some material from a comet and return it to Earth for study. This was the task of the Stardust mission, which encountered comet 81P/Wild 2 (Wild 2) in January 2004. During its flight through the comet's coma, which reached a closest approach to the nucleus of 236km, the spacecraft passively collected dust using aerogel and foil collectors. At the same time, observations were collected by three active experiments: The Stardust Imaging Camera, CIDA (Comet and Interstellar Dust Analyzer) and DFMI (Dust Flux Monitor Instrument) (Tsou, et al., 2004).

The Stardust Imaging camera captured over 70 images of the nucleus and its immediately surrounding region. Craters, landslides, cliffs and overhangs were observed on the nucleus' surface with spatial resolutions down to as low as 20 m pixel^{-1}. Both old and young surface was observed, and the high level of cohesiveness and roughness of the surface suggested that the nucleus was less evolved than those of Halley and Borrelly (Keller & Kührt , 2020; Tsou, et al., 2004).

A number of narrow dust jets were identified within the Stardust images, and the orientations and source locations of twenty of them were derived. Seven were found to coincide with bright spots on the nucleus. As with Borrelly, some of the dust emissions were found to emanate from the nucleus' dark side, driven by stored thermal energy rather than direct solar heating (Sekanina, et al., 2004).

DFMI observations found dust impacts, unexpectedly, to occur in bursts. Clark et al. (2004) suggest that this supports a model of dust being released from the nucleus in large clumps, which later disintegrate. Spacecraft crossings with dust jets, as determined from images, were found to coincide with strong peaks in the number of particles recorded by DFMI.

After its encounter with Wild 2, Stardust ejected its sample return capsule, which entered the Earth's atmosphere and safely landed in 2006. The capsule's collected material suggested many of the Comet's particles were weakly bound aggregates, dominated by olivine, low-calcium pyroxene or iron-sulphides. A large amount of variability in the presence of S, Ca, Ti, Cr, Mn, Ni, Cu, Zn and Ga in the collected particles indicated compositional diversity in the comet's ejected material (Flynn, 2008). From the Stardust sample return, Elsila et al. (2009) achieved the first detection of a cometary amino acid with their identification of extraterrestrial glycine.

After ejecting its sample return capsule, the Stardust spacecraft was redirected to perform a flyby of Comet 9P/Tempel 1, but this comet would first be visited by the Deep Impact mission.

4.3.5 Deep Impact

The Deep Impact mission entailed a flyby of comet 9P/Tempel 1 in July 2005. As part of the mission, an impactor released from the flyby spacecraft was purposefully collided with the nucleus' surface in order to produce a crater and ejecta cloud (Blume, 2005; Hampton, et al., 2005). The 372 kg impactor struck the nucleus' surface at a speed of 10 kms^{-1}, releasing a bright plume of material and creating a crater with an estimated diameter between 130 and 220 m (Schultz, et al., 2007). Little difference in the ejected material's composition from that of the pre-existing coma was measured (Keller & Kührt , 2020).

Three cameras aboard the Deep Impact spacecraft observed the nucleus and its surrounding environment before and after the impact, with a maximum spatial resolution of 8 m $pixel^{-1}$. Prior to impact, craters, terrain flows and eroded material were observed on the surface. Strong evidence of layering was also found (Thomas, et al., 2007). Observations of the nucleus' surface after the impact were obscured by the cloud of ejected material, though light scattered from this material illuminated a portion of the nucleus' night side. Small bright regions of exposed water ice were identified on the nucleus' surface (Sunshine, et al., 2006).

Comprehensive image sequences captured during the spacecraft's approach to the comet allowed the study of the dust coma's pre-impact state (Farnham, et al., 2007). Three major dust jets were observed, along with an associated arc-like dust feature. Temporal variation in the nucleus' activity due to its rotation was observed to result in large periodic outbursts of material. As with Borrelly and Wild 2, some dust jets appeared to originate from shadowed surface regions.

4.3.6 EPOXI

After its successful encounter with Tempel 1, the Deep Impact mission's flyby spacecraft was redirected for an encounter with comet 103P/Hartley 2, under a mission name of EPOXI (Extrasolar Planet Observation and Deep Impact Extended Investigation). In November 2010, a closest approach of 694 km was achieved with a flyby velocity of 12.3 kms^{-1}. The encounter is described by A'Hearn et al. (2011).

EPOXI's images revealed Hartley 2's nucleus to comprise two rough lobes, connected by a smooth narrow waist. Spectral analysis of the coma showed that H_2O vapour was

concentrated around the waist, whilst CO_2, water ice and organic species were concentrated at the nucleus' smaller lobe (A'Hearn, et al., 2011). Unlike comets Wild 2 and Tempel 1, widespread craters were not seen on the nucleus, which was covered, particularly at its waist, with non-volatile back fallen material.

With data from the MRI camera, Kelly et al. (2013) identified sizeable particles, as large as 20 cm, in the comet's coma, thought to consist either of bright icy material or dark dust. The production rate of H2O observed by EPOXI was greater than should be possible by sublimation from the surface, given its nucleus' surface area. A'Hearn et al. (2011) suggest this is indicative of solid chunks of water ice being ejected from the nucleus by subliming CO_2, after which they themselves sublime (a model which is consistent with the identification of large particles within the coma).

4.3.7 Stardust-NExT

The extended Stardust mission, Stardust-NExT performed a flyby of comet Tempel 1 in 2011, providing the first opportunity to revisit a comet, and a chance to better study the crater left by Deep Impact (Veverka, et al., 2013). The flyby occurred at a speed of 10.9 kms^{-1} and passed the nucleus with a closest approach of 178 km.

Photometry measurements of the nucleus returned similar results to those of Deep Impact, indicating a generally photometrically homogenous surface. Over an 8-minute period centred on closest approach, Stardust's navigation camera captured 72 images of Tempel 1's nucleus with resolutions ranging from 158 to 11 m $pixel^{-1}$. The region of the nucleus where Sunshine et al. (2006) identified bright areas of water ice in Deep Impact images was out of view for Stardust-NExT, and no such bright regions were identified within its images (Veverka, et al., 2013).

With a clearer view of the Deep Impact crater, its estimated diameter was refined by Richardson and Melosh (2013) to 49±12 m, with an ejecta blanket 85-120 m in diameter. From these measurements, they likened the surface's material properties to those of lightly packed, dry mountain snow.

Farnham et al. (2013) studied 11 dust jets within the Stardust-NExT images, and found that jet activity was largely associated with steeply sloped terrain and boundaries between rough and smooth surface. Some jets were observed to continue to emanate from regions of surface that had been in darkness for up to 4 hours. Unlike at Wild 2, the dust analyser (DFMI) showed no evidence of the spacecraft passing directly through Tempel 1's jets.

4.3.8 Rosetta

By the 2010s, flybys of six separate comets had been performed, five of them with high resolution imaging of the nucleus. Coma material had successfully been returned to Earth for laboratory analysis, and a nucleus had been impacted in order to facilitate the study of its surface and sub-surface.

Even before the first of these flybys took place, planning was underway by NASA and ESA for a joint mission to collect material from a comet's nucleus and return it to Earth. Various budget issues for NASA led to their pulling out from the project, and ESA had to temper the ambitious original goals, culminating in the design of the Rosetta mission (Keller & Kührt , 2020). Rosetta would perform the first rendezvous with a comet, and escort it for an extended period, rather than perform a flyby. Instead of collecting and returning nucleus material, it would deploy a lander to the surface in order to perform in-situ measurements (Glassmeier, et al., 2007).

The Rosetta spacecraft rendezvoused with comet 67P/Churyumov-Gerasimenko (67P/C-G) in August 2014, after a ten-year journey. The mission was originally targeted at comet 46P/Wirtanen, but a launch delay required the change of comet.

The rendezvous occurred at a distance of 3.5 AU from the Sun, and Rosetta orbited the comet for two years, following it to its perihelion distance of 1.25 AU and back out to 3.8 AU, when the spacecraft, in September 2016, was purposefully crashed into the nucleus (Accomazzo, et al., 2017). This extended period allowed the spacecraft to monitor the comet's evolution as it approached the Sun, and then compare its post-perihelion and pre-perihelion states (Taylor, et al., 2017).

Rosetta's imaging suite OSIRIS (see section 4.4.5) captured extensive observations of the comet's nucleus, dust jets and wider coma. It revealed the bilobate nucleus to be covered in surface features at a range of scales. Cliffs and mountains, comparable in appearance to those on Earth, were widespread, whilst metre-scale fractures indicated the effect of gravitational stresses and seasonal thermal contraction on the surface (El-Maarry, et al., 2015). The nucleus has been significantly evolved by its approaches to the Sun, with an abundance of granular back fallen material (dust). Material was observed to emanate from all illuminated surface, rather than being limited to isolated active regions (Keller, et al., 2017; Kramer & Noack, 2016).

El-Maarry et al. (2019) found that all of the previously visited comets' nuclei possessed some surface features that were also present on 67P/C-G. The prolonged period that Rosetta spent orbiting and imaging 67P/C-G allowed the OSIRIS cameras to image its

entire surface down to scales as small as 0.2-0.3 m pixel^{-1}. This permitted the construction of accurate (1-1.5 m horizontal sampling) 3D shape models of the surface (e.g. Preusker et al. (2017)), and was used to identify safe landing sites for the mission's lander Philae.

On the 12th of November 2014, Rosetta released its Philae lander, which seven hours later made contact with the surface. The lander itself is described by Bibring et al. (2007), whilst the details of its landing are described by Biele et al. (2015). Due to the nucleus' low surface gravity, the lander required thrusters and harpoons to keep it on the surface, the failure of which led to Philae bouncing several times before coming to rest on the surface two hours later. Biele et al. used data from Philae's engineering sensors to estimate, from its bounce characteristics, the properties of the surface material. They estimated the location of the lander's first touchdown to have a 20 cm deep layer of granular soft material, whist Philae's final touchdown location had a hard, stiff surface. From Rosetta observations, O'Rourke et al. (2020) identified water ice in the small craters exposed by Philae as it bounced.

Images captured during the lander's descent showed the initial touch down site to be covered in granular debris, with the presence of Aeolian-like erosion features (Mottola, et al., 2015). Panoramic images captured after Philae's final touchdown provided images with millimetre and centimetre pixel scales, and revealed the presence of cliffs, pebbles and fractured boulders on a surface of very low average reflectivity (Bibring, et al., 2015; Boehnhardt, et al., 2017).

Rosetta's long term orbiting of 67P/C-G permitted extensive study of its dust and gas environment. Vincent et al. (2016) found dust jets to predominantly arise from rough, fractured terrain, and they suggest that this generally drives erosion of cliffs on comet nuclei. Night time activity (an outburst from the nucleus' night side) was observed by Knollenberg et al. (2016).

Mass spectrometry of the comet's coma identified the presence of almost all pre-established cometary species, with an abundance of CO_2 and C_2H_6. A large difference in the relative abundances for the comet's summer and winter hemispheres suggests that surface evolution was a factor (Le Roy, et al., 2015).

4.4 Imaging Comets from Spacecraft

This section gives an overview of cameras which have observed comets from spacecraft. It includes only cameras whose main purpose was to take images of a comet during an

encounter, but it is worth noting that valuable observations of comets have been provided by cameras on the solar-observing SOHO and STEREO spacecraft (e.g. Knight & Battams (2014) and Hui (2013)) and the Hubble Space Telescope (e.g. Weaver et al. (1995)).

4.4.1 Giotto's Halley Multicolour Camera

The Halley Multicolour Camera (HMC) was an imager aboard the Giotto spacecraft, which captured high resolution images of the nucleus and surrounding coma of comet Halley. The instrument is described in detail by Keller et al. (1987).

HMC was a line scan (push-broom) camera, with its scan motion provided by Giotto's spin stabilisation. As Giotto rotated, the narrow ~0.5° FOV of the detector described an annulus in the sky, centred on the spacecraft's spin axis. A turret mirror controlled the radius of the annulus, such that it always intercepted the nucleus of Halley, even as the viewing geometry changed with flyby progression. Only a small portion of the annulus, that which contained the comet's nucleus and surrounding area, was imaged by the camera. Keller et al.'s illustration of this imaging method is reproduced here in Figure 4.4.

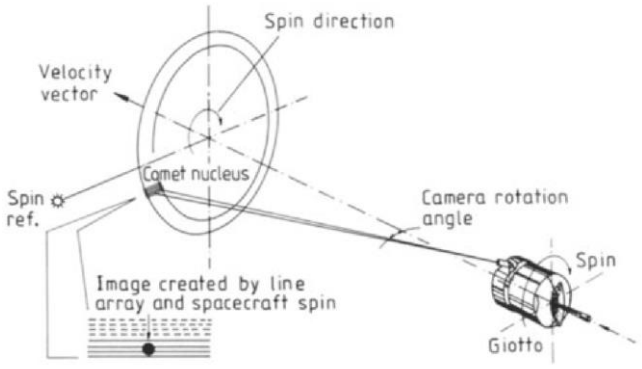

Figure 4.4: HMC's principle of line scan imaging from a spin stabilised spacecraft during a comet flyby. Reproduced from Keller et al. (1987).

The constant motion of HMC's FOV across the sky limited its possible exposure times. At some imaging geometries, exposure time could be limited to as low as 14 μs. Such a small exposure time could not be achieved with a mechanical shutter and line scanner, and a frame transfer CCD was instead used. After a line of the detector was exposed, it was shifted to a masked region of the detector. Lines built up to produce a 2D image, which was then read out from the detector during the non-imaging portion of the spacecraft spin. Use of a CCD also permitted the implementation of TDI in order to further

assist with maximising signal. Because of the curved motion of the comet's image across the camera's detector, use of TDI led to a small loss of resolution. Multiple degrees of on-chip binning were also used to balance resolution and signal (Keller, et al., 1987).

HMC was fitted with a filter wheel containing nine spectral filters and two polarisers. These included two high pass filters admitting wavelengths above 300 and 700 nm respectively, orange and blue broadband filters centred on 650 and 420 nm respectively (each with FWHM ~150 nm), and narrow band filters centred on OH (315 nm), C_3 (408 nm) and C_2 (509 nm) emission (Thomas & Keller, 1990).

4.4.2 Deep Space 1's Miniature Integrated Camera Spectrometer

The Miniature Integrated Camera Spectrometer (MICAS) was a compact, low mass camera aboard the Deep Space 1 spacecraft. It imaged the nucleus and inner coma of comet 19P/Borrelly during the spacecraft's flyby. The camera was intended as a test and demonstration of a capable, compact, low mass imaging system (Beauchamp, et al., 2000). This is relative however; at a total mass of 12kg, it is half the mass of the OSIRIS instrument (section 4.4.5), but approximately ten times more massive than EnVisS.

Two visible light detectors (a CCD and a CMOS) and two spectrometers (IR and UV) all shared the same optics (Soderblom, et al., 2002). The UV detector failed soon after launch, and neither spectrometer will be discussed here. Electronic shuttering of the CCD and CMOS meant that the instrument contained no moving parts. The 1024x2048 pixel frame transfer CCD integrated with half its pixels, and read out from the other, masked half, to which images were quickly transferred after exposure. The CMOS imaged with a square of 256x256 pixels.

The optics had an aperture diameter of 100 mm, projecting narrow FOVs of 0.69°x0.78° and 0.26°x0.26° onto the active regions of the CCD and CMOS respectively. Each detector was sensitive over the 500 nm to 1000 nm wavelength range.

4.4.3 Stardust Imaging Camera

The stardust spacecraft, which performed a flyby of comet 81P/Wild 2 in January 2004 followed by a flyby of comet 9P/Tempel 1 in February 2011, included in its scientific payload an imager called the Imaging Camera. Its main purpose was to assist the spacecraft's navigation as it approached the comets, but its images were also used for scientific study of the nucleus. The camera is described in detail by Newburn et al. (2003).

The 3.5° x 3.5° FOV of the instrument was focussed onto a CCD of resolution 1024 x 1024 pixels through a 58 mm diameter aperture. A scan mirror, centred on the instrument's optical axis at an angle of 45°, was used to control the pointing of the camera (in conjunction with spacecraft attitude). A second, fixed periscope mirror was used to image the nucleus during the encounter, when the dust shield obscured it from the scan mirror (in order to protect the scan mirror from dust impacts).

Eight filters, installed in a single filter wheel, were selected to address the navigation and science goals of the instruments. However, after launch but before imaging of comet Wild 2 began, the filter wheel became stuck in the position of the broadband optical navigation filter, through which all images of the comet had to be acquired. That the filter became stuck on the broad band filter was somewhat fortuitous, as its wide spectral range of ~500 nm to 900 nm ensured a high signal. However, the FWHM of the instrument's point spread when using this filter was 1.6 pixels (~6 times larger than the scientific filters) leading to a reduced angular resolution.

4.4.4 Deep Impact's High Resolution and Medium Resolution Instruments

The High Resolution (HRI) and Medium Resolution (MRI) Instruments on the Deep Impact spacecraft were imagers which observed comet 9P/Tempel 1 over the course of the mission's rendezvous with the body, followed by 103P/Hartley during the EPOXI extension of the mission (Hampton, et al., 2005). The Deep Impact mission involved impacting a spacecraft with the comet surface and imaging the resulting crater and ejecta plume. As described by Hampton et al. (2005), the challenge that this posed to the mission's instruments was that they needed to be designed to obtain useful data for any of the very large range of possible outcomes of the mission (given that impacting a comet's nucleus was unprecedented). Comet Interceptor's instruments, EnVisS included, face a similar challenge, with the conditions of the yet-to-be discovered DNC unknown.

Deep Impact's HRI and MRI were mounted on the flyby spacecraft, from which they imaged the nucleus of Tempel 1 before and after impact. The impact spacecraft was equipped with an Impactor Targeting Sensor (ITS) which was identical to the MRI, but for its use of an unfiltered CCD, providing high SNR images for assisting with navigation. The HRI performed visible imaging with a CCD, and IR spectral imaging with a long-slit IR spectrometer, though only the visible imaging capabilities will be discussed here.

The HRI imaged through an aperture of 30 cm diameter, with a focal length of 10.5 m. Its square 0.118° FOV projected onto a 1024 x 1024 pixel CCD. The MRI used the same format CCD, with an aperture diameter of 12 cm and focal length of 2.1 m resulting in a

0.587° FOV. Each instrument had a 9-position filter wheel whose spectral filters are shown in Table 4.1

Table 4.1: Deep Impact HRI and MRI filters. Reproduced from Hampton et al. (2005).

Filter wheel position	MRI centre [nm]	MRI filter width [nm]	MRI filter target measurement	HRI centre [nm]	HRI filter width [nm]
1	650	>700[a]	Context	650	>700[a]
2	514	11.8	C_2 in coma	450	100
3	526	5.6	Dust in coma	550	100
4	750	100	Context	350	100[b]
5	950	100[c]	Context	950	100[c]
6	650	>700[a]	Context	650	>700[a]
7	387	6.2	CN in coma	750	100
8	345	6.8	Dust in coma	850	100
9	309	6.2	OH in coma	650	100

[a]Filters in positions 1 and 6 are uncoated and not band limited.
[b]The coating on the 350 nm filter is shortpass, the substrate limits the short wavelength throughput.
[c]The 950 nm filter is longpass.

A stipulated requirement (no. 18) of the instruments in Hampton et al. (2005) was that coma observations at the expected brightness of approximately 1.5×10^{-6} $Wm^{-2}sr^{-1}nm^{-1}$ should provide an SNR of at least 100 through dust filters and a signal level at least 2 DN higher than background through gas filters. Additionally, requirement 7 stated that an SNR of at least 30 should be achieved on the inner coma and ejecta cloud (up to a distance of two nucleus radii from the surface). Using calibrated models of the instruments' signal throughputs, required exposure times for observations to achieve their desired SNRs were calculated, with an exposure time of 80 ms being required for HRI to achieve an SNR of 70 when observing the nucleus surface through filter 1 for example (whilst an exposure time of 4 ms was modelled to achieve an SNR of 10 for the same observation).

4.4.5 Rosetta's Optical, Spectroscopic and Infrared Imaging System

The Optical, Spectroscopic and Infrared Imaging System (OSIRIS) was a camera system aboard the Rosetta spacecraft, which intercepted and escorted comet 67P/Churyumov-Gerasimenko from 2014 to 2016. Comprising a Narrow Angle Camera (NAC) and Wide Angle Camera (WAC), the instrument is described in detail by Keller et al. (2007).

These two large cameras imaged the surface of the nucleus at high resolution (down to a few cm per pixel by the last stages of the Rosetta mission), along with dust and gas jets in the inner coma. Objectives of the cameras included determining the size, shape and density of the nucleus, identifying regions of activity on the surface, observing temporal variation in the surrounding dust environment and measuring the dust's optical properties. Polarimetry was considered, but omitted from OSIRIS' capabilities due to cost and complexity. The entire OSIRIS instrument had a mass of 35 kg.

The NAC imaged a 2.2° by 2.2° FOV with a CCD comprising 2000 by 2000 pixels, whilst the WAC had a FOV of 12° by 12°, which projected onto an identical detector to that of the NAC. Their optics comprised mirrors, and with aperture sizes of 90 mm and 25 mm respectively, they were significantly larger cameras than EnVisS. Both cameras utilised an electromechanical shutter capable of exposures of any multiple of 0.5 ms, no smaller than 10 ms. On the 3-axis stabilised platform of Rosetta, exposure times as long as tens of seconds were employed to achieve high signal, high dynamic range (14 bit) images of the dust and gas environment. This permitted the study of a variety of narrow spectral ranges as well as broad, using the most comprehensive set of optical filters to have ever been sent to a comet. Each camera used two filter wheels to store and select their filters, the details of which are listed in Table 4.2.

Table 4.2: The filters of OSIRIS' NAC and WAC cameras (Keller, et al., 2007).

Camera	Name	Wavelength [nm]	Bandwidth [nm]	Peak transmission [%]
NAC	FFP-UV	250-850		>99
	FFP-Vis	250-1000		>95
	FFP-IR	300-1000		>99
	NFP-Vis	300-1000		>98
	Far-UV	269.3	53.6	37.8
	Near-UV	360.0	51.1	78.2
	Blue	480.7	74.9	74.6
	Green	535.7	62.4	75.8
	Neutral	640.0	520.0	5.0
	Orange	649.2	84.5	92.4
	Hydra	701.2	22.1	87.4
	Red	743.7	64.1	96.0
	Ortho	805.3	40.5	69.8
	Near-IR	882.1	65.9	78.4
	Fe_2O_3	931.9	34.9	81.6
	IR	989.3	38.2	78.1
WAC	UV245	246.2	14.1	31.8
	CS	259.0	5.6	29.8
	UV295	295.9	10.9	30.4
	OH-WAC	309.7	4.1	26.0
	UV325	325.8	10.7	31.6
	NH	335.9	4.1	23.6
	UV375	375.6	9.8	57.3
	CN	388.4	5.2	37.4
	Green	537.2	63.2	76.8
	NH_2	572.1	11.5	60.9
	Na	590.7	4.7	59.0
	VIS610	612.6	9.8	83.4
	OI	631.6	4.0	52.4
	R	629.9	156.8	95.7

4.4.6 Giotto's Optical Probe Experiment

Although not an imager, the Optical Probe Experiment (OPE, originally the Halley Optical Probe Experiment, HOPE) aboard Giotto is included here because of its use of polarimetry. Mounted on the rear of Giotto, viewing antiparallel to its direction of motion (i.e. directly away from Halley's nucleus), OPE collected multispectral brightness and polarisation measurements of comet Halley's coma (Levasseur-Regourd, et al., 1984).

The instrument collected light from a circular 2.6° diameter FOV, centred on Giotto's anti-velocity vector. A spectral filter mosaic on the objective lens and a segmented detector enabled multispectral measurements without the need for a filter wheel (this was driven by a need for reliability, and led to an instrument of only 1.3 kg). The bandpasses imaged by OPE are listed in Table 4.3. Each filter (except for CO+, with which no polarimetry was performed) imaged through a single fixed polaroid foil (mounted in front of the objective lens) (Giovane, et al., 1991). Giotto's spin changed the polarisation axis of the polaroid relative to the scene. Over the course of half a spin, individual measurements could be taken with polariser angles of 0° to 180° (in 45° steps), and the linear Stokes parameters could be derived (Levasseur-Regourd, et al., 1999).

Whilst the design of EnVisS' imaging polarimetry method is different to that of OPE, there will be a small portion of the EnVisS' field of view that images antiparallel to the spacecraft velocity vector, and OPE's method of polarimetry could be replicated over this portion of sky, as is discussed further in section 6.5.

OPE's measurements of both comets 1P/Halley and 26P/Grigg-Skjellerup provided high quality calibrated data on the brightness of their comas, and they are a valuable resource for EnVisS (Levasseur-Regourd, et al., 1999). Later (section 6.3), these data will be used to help to constrain the imaging conditions EnVisS will encounter.

Table 4.3: Characteristics of the individual filters comprising OPE's filter mosaic (Giovane, et al., 1991).

Filter	Central wavelength (nm)	Bandpass (nm)
UV	368.1	5.6
Blue	441.4	4.7
Green	576.1	9.8
Red	718.1	3.35
OH	310.3	5.4
CN	386.5	3.7
CO^+	423.9	3.6
C_2	513.5	3.6

4.5 Discussion

The cameras which have flown to comets to date are mostly a variation on a theme (with the exception of HMC) in that they were designed to image a comet's nucleus and immediate surroundings with relatively narrow FOVs, from a pointable 3-axis stabilised spacecraft. The major advantage of this pointing stability is the scope to obtain long exposure images which can reveal the faint details within cometary environments. HMC had to overcome the challenge of imaging from a spinning spacecraft, and did so with the push-broom technique, but its maximum possible exposure time was restricted as a result.

The narrow FOVs of these cameras were designed to resolve structures on the nuclei's surfaces, and in their immediately surrounding dust environments. As is apparent from the discussion of section 4.2, comets extend far beyond this region, and processes occurring at large scales have important links with the state of the comet's nucleus and the evolution of cometary material. Earth-based cometary observations typically focus on these larger scales, mainly because telescopes are too distant to resolve scales as small as the nucleus. But Earth-based observations of comets are restricted to a small range of observational geometries limiting, for example, the phase angles from which a comet can be observed. Understandably, when cameras have been flown to comets, the unique opportunity to focus on the nucleus with great detail has been exploited. The result of this however is that the larger coma, dust and gas structures of comets are relatively poorly observed from these close perspectives. There is valid motivation therefore to pursue different approaches to spaceborne cometary imaging that can

address these gaps in the observations, and chapter 6 presents work that was carried out for such an instrument.

5 A Scanning Descent Camera for a Planetary Penetrator

This chapter describes a study, carried out for this thesis, into the concept of performing descent imaging from a planetary penetrator using a rotational scanning camera. Penetrators are introduced in section 5.1 and the instrument concept is described in section 5.2. This is followed by a detailed analysis of its imaging geometry in section 5.3. Examples of the camera's expected data, produced through software simulation, are presented in section 5.4, and their processing for scientific analysis is simulated in section 5.5. Section 5.6 details a study of descent stereophotogrammetry using the camera's images. Much of the material in this chapter, particularly section 5.6, is also published in Brydon et al. (2021). The findings of this chapter are summarised and discussed in section 5.7.

5.1 Planetary Penetrators

Penetrators and hard landers provide a method by which to deliver hardware to a planetary surface, which is in contrast to the more common approach of soft landers. By design, they descend to a surface at high speed, and are hardened to survive the impact, either remaining on the surface or penetrating down to sub-surface depths of a few metres (Lorenz, 2011). Onboard instrumentation, designed to withstand the rapid deceleration of impact, thus gains access to the surface/sub-surface where in-situ measurements can be made.

Similar to but distinct from penetrators are impactors, which are designed to collide with a planetary surface but not survive, and not perform any measurements or other functions post-impact. The first such example is the USSR's Luna 2, which impacted the Moon in 1959, a few months after the failed attempt to do the same with Luna 1 (Huntress, et al., 2003). These were followed by NASA's Ranger program, which included 7 (Rangers 3-9) spacecraft designed to crash into the lunar surface. Five of them succeeded, whilst Rangers 3 and 5 both missed the moon due to propulsion errors (Schurmeier, et al., 1966). Each spacecraft carried television cameras, which were designed to capture sequences of surface images during their approach to impact. Also

carried aboard Rangers 3, 4 and 5 were hard landers in the form of spherical capsules, surrounded by balsa wood to cushion their high speed impacts and allow their instruments to operate after landing. None successfully landed. These three probes are often omitted from reviews of penetrator technology, but with their science payloads, which were hardened to survive a violent landing, they represent the first attempts to emplace instrumentation on a planetary surface at high speed.

Numerous penetrator concepts and missions have been studied over the past three decades, targeting not just the Moon (Gao, et al., 2008; Shiraishi, et al., 2008), but also Europa (Gowen, et al., 2011), and short period comets (Boynton & Reinert, 1995). The only planetary penetrators to make it through development to launch were aboard Mars bound missions Mars 96 (Surkov & Kremnev, 1998) and Deep Space 2 (Smrekar, et al., 1999). The Mars 96 mission failed to escape Earth's orbit, and crashed into the Pacific Ocean a day after its launch in 1996 (Rich, 1996). The Deep Space 2 penetrators (of which there were two), which were launched alongside the Mars Polar Lander, successfully arrived at Mars in 1999, and began their atmospheric entry, losing contact with Earth as planned. The expected post-landing resumption of contact never came, and all three craft were lost. Whether the penetrators successfully impacted the Martian surface and operated their science payloads, is unknown (Albee, et al., 2000). LUNAR-A, the only other penetrator mission to undergo significant development with the expectation that it would be launched and deployed, was abandoned after over a decade of development due to funding issues (Shiraishi, et al., 2008).

As a result, there exists no track record of successful deployment of planetary penetrators. There has however been development and testing of the technology on Earth, in preparation for the two failed Mars missions, the abandoned LUNAR-A mission, and for mission concepts (Gowen, et al., 2011). It should also be noted that penetrators have their origins in terrestrial applications, particularly military, from which the survivability of high speed impactors is well studied. Lorenz (2011) provides a more complete history and discussion of planetary penetrators, from their terrestrial origins to their inclusion in space missions.

The major advantage of penetrators over soft landers is the relative simplicity of their landing method. They don't require complex attitude control or autonomous landing systems and therefore provide a lower cost, lower mass method by which to access planetary surfaces and sub-surfaces (Lorenz, 2011). Achieving an impact trajectory for a penetrator requires less fuel than establishing a spacecraft in orbit, and so penetrators also provide potential opportunities to observe bodies which cannot be reached by

orbiters. Their low cost and mass additionally presents the option of incorporating multiple penetrators into a single mission, to achieve multipoint sampling or reduce the likelihood of failure.

By their nature, penetrators are compact and self-contained, restricting their capacity for science hardware, and so they will always carry a smaller payload than can be accommodated by a large orbiter, or soft lander such as the Mars Science Laboratory Curiosity rover (Grotzinger, et al., 2012). Whilst some of the penetrator designs discussed above include cameras for imaging of the surface or sub-surface after impact (such as Mars 96's TV-camera), an often-omitted instrument is a descent camera, for surface imaging during the penetrator's descent, the value of which is outlined in section 3.2.

5.2 Camera Concept

A spinning penetrator presents a challenging opportunity to perform surface imaging of a target body. Over the short period of the descent, the viewpoint is constantly changing, potentially at a high rate. Little time is available for cycling through a filter wheel or mechanically changing the pointing of the camera. Coupled with the low mass and volume constraints of a penetrator, this necessitates a small, simple and fast imager.

The camera concept investigated and presented in this chapter aims to address and overcome these challenges, and provide the opportunity to acquire valuable scientific data from a difficult imaging opportunity. The camera has its origins in the Akon concept, an M-class proposal submitted to ESA in 2016 comprising a spin stabilised penetrator which would have impacted Europa's surface (Jones, et al., 2017). Included in the proposed payload was a multispectral descent camera. This chapter investigates the performance and capabilities of a camera based on its design and operating principle, which is now described.

The instrument concept requires a physically compact camera, such that it could be accommodated on a penetrator (it would also be applicable to other low mass probes). Previously flown instruments, to which this camera would likely need a comparable size, include the Rosetta mission's CIVA-P cameras of dimensions 70 x 52 x 36 mm (Bibring, et al., 2007), and the 113 x 136 x 81 mm MasCam of the Hayabusa 2 mission (Jaumann, et al., 2016).

The penetrator on which the camera is mounted is assumed to be spin stabilised and descending vertically. The spin axis is fixed parallel to the penetrator's velocity vector

(survival of the penetrator requires that this be satisfied to within ≈5° (Lorenz, 2011)), and the penetrator's nadir is therefore also its landing site. Spinning spacecraft are often stabilised with a spin period of a few seconds (e.g. Hansen et al. (2017); Reinhard (1986)), but a penetrator may require a higher spin rate (e.g. 1 Hz) to maintain a sufficiently precise orientation, especially in the presence of an atmosphere (Lorenz, 2010). For a free fall, descent velocity would start close to 0 ms^{-1} and for most planets and moons of interest would likely reach between 100-500 ms^{-1} at impact.

Descent imaging occurs through wide angle, fisheye optics to achieve a 90° x 90° FOV. The camera is mounted on the side of the penetrator, either externally or viewing from internally through a window, such that its principal axis views 45° below horizontal, and its vertical FOV is in plane with the penetrator's spin axis. This results in the camera's FOV extending from the penetrator's nadir all the way to horizontal. Whether mounted internally or externally, surviving impact would not be a requirement of the descent camera.

The camera employs scan imaging (both push-broom and push-frame methods could be utilised, based on other driving factors, and analysis throughout this chapter is kept general enough to apply to both), whereby imaging is carried out with a narrow sensor spanning the full 90° vertical FOV, but a smaller portion of the horizontal FOV. The footprint of such a sensor on the surface approximately describes a thin circular sector centred on nadir and extending to the camera's visible horizon. As the penetrator spins, this footprint sweeps across the surface through a full 360° of azimuth. Images from the narrow sensor are captured at a frequency linked to the penetrator's spin rate such that they combine to produce a scan of the entire visible surface, as illustrated in Figure 5.1.

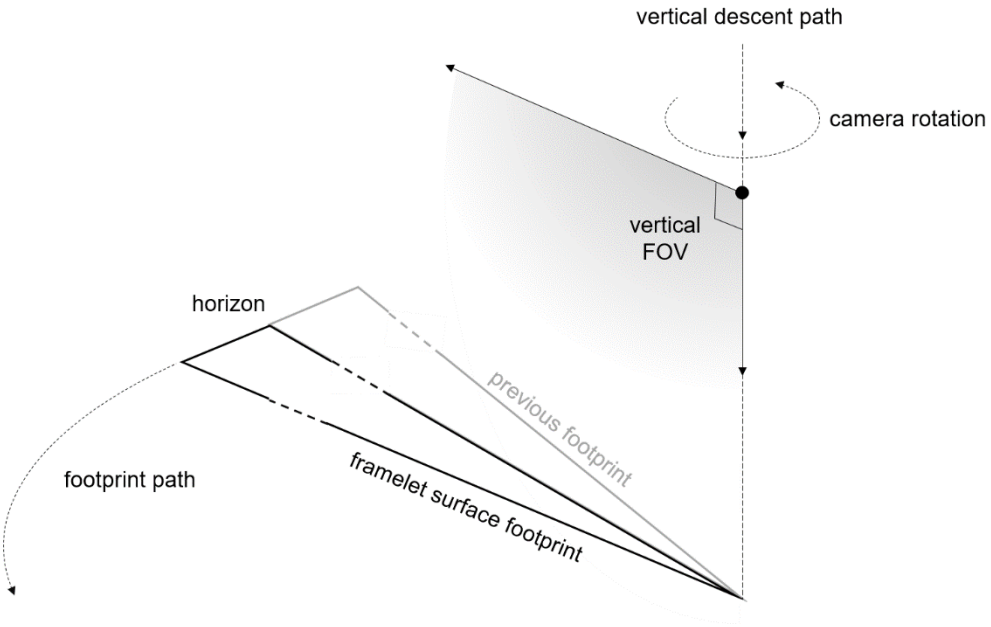

Figure 5.1: Penetrator descent camera imaging concept.

As the penetrator descends, a series of nested images is captured with spatial resolution increasing with each successive image. Thus a multiscale dataset is obtained with large surface coverage and high spatial resolution proximal to the impact site. A sequence of images over the course of the descent provides a gradual and continuous transition from orbital to landed viewing perspective (Brydon & Jones, 2018).

To achieve multichannel imaging, a CCD or CMOS array detector is used, with different portions of it dedicated to their own imaging channels, which image through different optical filters, each spanning the detector's full vertical FOV, as shown in Figure 5.2. Under each filter, a column, or a small number of adjacent columns of pixels is exposed in order to capture framelets to build up full-scan images. Columns under more than one filter can be exposed and read out together, allowing multiple channels to image simultaneously. This is the mechanism by which multichannel (multispectral and polarimetric) imaging can be performed in such a short lived, fast moving descent.

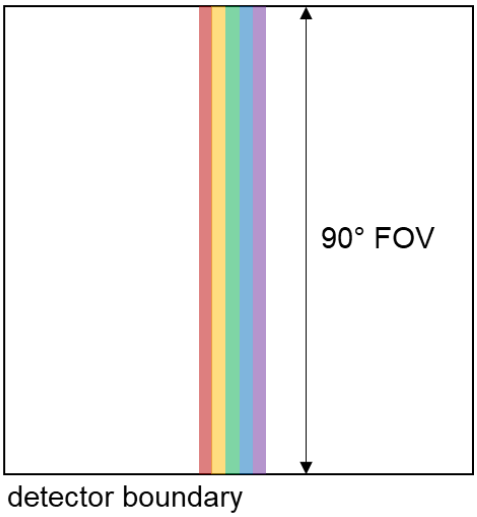

Figure 5.2: Illustration of the detector and filter arrangement for the scanning descent camera concept.

Next, in section 5.3, a more detailed model of the camera's imaging geometry and operation will be developed and presented, whilst section 5.4 shows the sort of data the camera would be expected to return. It is important to note that whilst the theme of this chapter is this penetrator descent camera, with its specific mode of imaging and intended spacecraft, much of the analysis could readily be applied, with minimal modification, to other platforms (including the descent sequence of a soft lander, ascent/descent balloons or an aerial vehicle similar to the Ingenuity helicopter (Balaram, et al., 2021)), and camera designs (such as side- and downward-looking framing cameras with wide or hemispherical FOVs).

5.3 Camera Model

5.3.1 Imaging Geometry

As described above, the imaging concept comprises a camera mounted on a spin stabilised penetrator, and its images are the result of scanning as the penetrator rotates around a fixed, vertical spin axis. To achieve a 90° FOV conventionally requires fisheye optics, and an equidistant geometry (see section 2.1.2) is assumed throughout this chapter to define the camera's mapping function. Let $(\hat{x}, \hat{y}, \hat{z})$ denote the camera's Cartesian coordinate frame (as described in section 2.1.1), whose origin is assumed to be on the penetrator's spin axis (a reasonable approximation given the small size of a penetrator and the large distances at which the camera would be imaging a surface). The camera is mounted such that its $\hat{y}$ axis, its $\hat{z}$ axis (its principal axis) and the

penetrator's spin axis are all in the same plane. The angle between the nadir direction and principal axis of the camera is denoted A, and henceforth referred to as the camera's principal off-nadir angle.

As the camera spins, a point on the surface moves, in the camera frame, along a circular path centred on, and in a plane (here called the rotation plane) perpendicular to, the spin axis. The spin axis intersects the rotation plane a distance h' from the camera's optical centre. The intersection of the camera's y-z plane and the rotation plane defines a camera reference direction, from which the point's camera-centric azimuth θ is measured (anti-clockwise around the spin axis). The camera has a principal azimuth Θ measured from a fixed planetocentric reference direction, from which the surface point has a planetocentric azimuth θ_p. The point is a distance r from the spin axis (measured in the rotation plane). Figure 5.3 shows the described geometry.

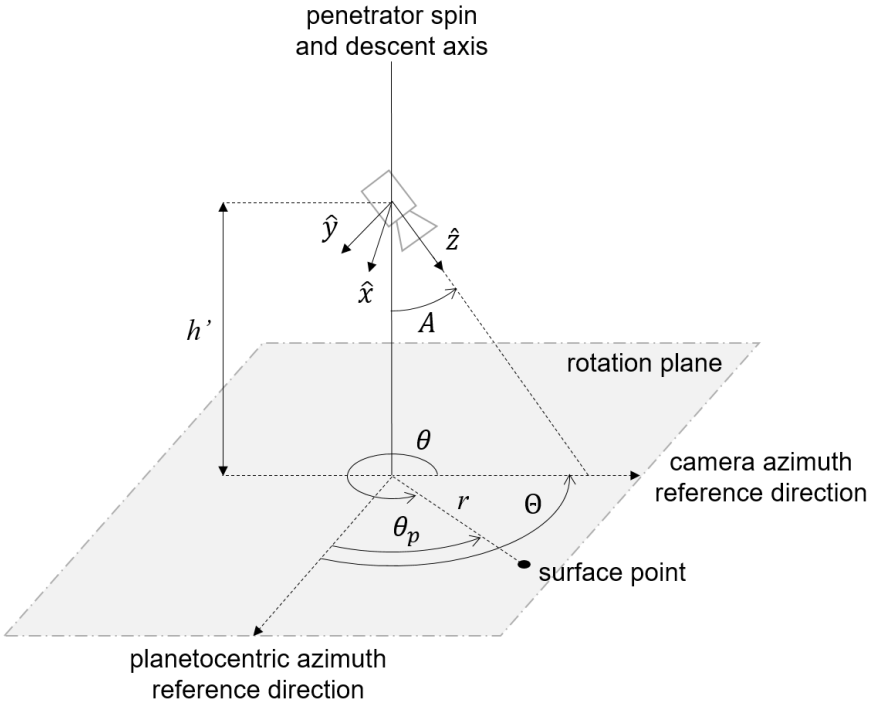

Figure 5.3: Diagram of the coordinate systems described in the main text.

A point is transformed from its cylindrical coordinates $[r \quad h' \quad \theta]^T$ to its Cartesian position in the camera frame by

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \vec{p} = \begin{bmatrix} -r\sin\theta \\ h'\sin\mathrm{A} - r\cos\theta\cos\mathrm{A} \\ h'\cos\mathrm{A} + r\cos\theta\sin\mathrm{A} \end{bmatrix} \tag{5.1}$$

Equation 5.1, in combination with equation 2.12, therefore defines the mapping of a point on the rotation plane to its position on the camera's focal plane. Because the camera captures images whilst moving, it is important to also describe the imaging geometry's time dependence. The spin of the camera leads to a change in the value of θ, whilst the descent motion results in a changing value of h'. The motion of a point within the camera frame is thus described by

$$\frac{d\vec{p}}{dt} = \begin{bmatrix} -r\cos\theta\dfrac{d\theta}{dt} \\ \sin\mathrm{A}\dfrac{dh'}{dt} + r\sin\theta\cos\mathrm{A}\dfrac{d\theta}{dt} \\ \cos\mathrm{A}\dfrac{dh'}{dt} - r\sin\theta\sin\mathrm{A}\dfrac{d\theta}{dt} \end{bmatrix} \tag{5.2}$$

This equation's use need not be restricted to a spinning penetrator descent camera, and could in fact be applied to a wide range of spacecraft mounted cameras. Examples include descent cameras on stabilised and non-spinning platforms like the Mars Science Laboratory Entry Descent and Landing System (Prakash, et al., 2008), or cameras on board aerial probes such as ascending/descending balloons (e.g. Rand & Phillips (2002) and Chassefiere et al. (2009)) or the Mars helicopter Ingenuity on the Perseverance rover mission (Balaram, et al., 2021). More complex motion of the camera platform, such as atmosphere-induced oscillations, could be approximated by the expansion of equation 5.2 with inclusion of time dependence in A, for example. The rate of image motion due to the camera's rotation and descent can readily be found from equation 5.2 in combination with equations 2.24-2.27 in section 2.1.4. This is addressed further in section 5.3.3.

Whilst the cylindrical coordinates used here are suitable for describing the location and motion of a single point on a surface, they are not quite sufficient for describing a real, extended planetary surface. This is because, due to both surface curvature and topography, different points on a planetary surface will correspond to different values of h', and will move in their own rotation planes. Conversion between each point's cylindrical coordinate frame and a common planetocentric frame is achieved by the following equations, with the underlying geometry illustrated in Figure 5.4.

$$h' = h + R_p - \left(R_p + e\right)\cos\frac{R_g}{R_p} \tag{5.3}$$

$$r = \left(R_p + e\right)\sin\frac{R_g}{R_p} \tag{5.4}$$

$$\theta = 2\pi + \Theta - \theta_p \tag{5.5}$$

where the camera altitude h and the point's elevation e are measured relative to a reference spheroid (datum) of local radius (assumed to be constant over the extent of the imaged surface) R_p.

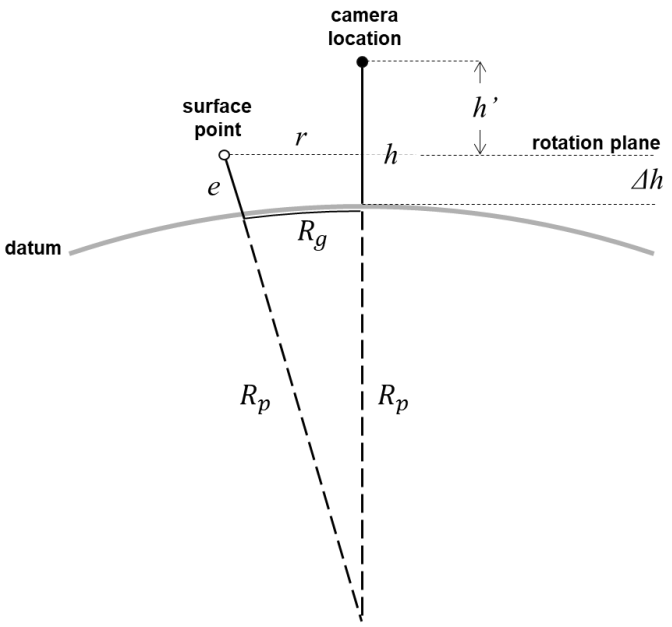

Figure 5.4: Planetocentric representation of the cylindrical coordinates of equation 5.1.

An illustration of the camera's imaging geometry is shown in Figure 5.5. It presents an orthographic bird's eye view of the surface, with the instantaneous projection of the camera's detector (assumed to be 1024 x 1024 pixels covering 90° x 90° FOV) overlaid for a penetrator altitude of 1 km. For clarity, individual pixels are not shown, and the grid instead splits the detector into 32 x 32 pixel segments. The penetrator's nadir is at (0, 0) km. Only the detector columns near the camera's principal axis (illustrated by the black grid), where distortion is at its lowest, would be used for imaging (c.f. Figure 5.2). Near to the nadir, pixel footprints are close to square in shape. At greater distances, they very evidently become elongated in the radial direction.

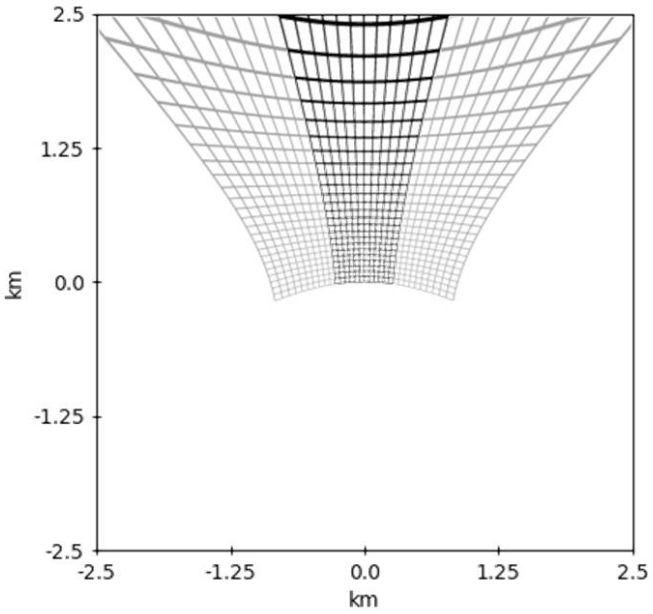

Figure 5.5: Orthographic view of the surface, from directly above the penetrator's nadir (located at 0, 0 km), showing the instantaneous projection of the camera's full 1024 x 1024 pixel detector on the surface, when imaging from 1 km altitude. Each grid square represents 32 x 32 detector pixels. Only a central portion (black) of the detector would actually be used by the penetrator descent camera for imaging.

The geometry of an image-scan captured by the penetrator camera depends on which columns of the detector were used to capture it (i.e. an image captured through the red filter of Figure 5.2 would have a different geometry to that of one captured through the blue filter). This is evident from Figure 5.5, where it is clear to see that different columns of the detector have differently shaped footprints on the surface. Figure 5.6 further illustrates this, by showing the footprints of two narrow strips of the detector on the surface at four different times during a rotation. The footprints, corresponding to columns 488-535 (black) and 440-487 (red) represent two different channels of the camera, for example two filters of different spectral range. The central yellow dot shows the position of the penetrator's nadir. Note that the black footprint, whose region on the detector encompasses the principal point, reaches all the way to the nadir. By contrast, the red footprint does not extend to the nadir, and its field of regard therefore omits an approximately circular portion of the surface. The radius of the omitted region of surface, here called the omission radius, is plotted for each column of the camera's detector in Figure 5.7.

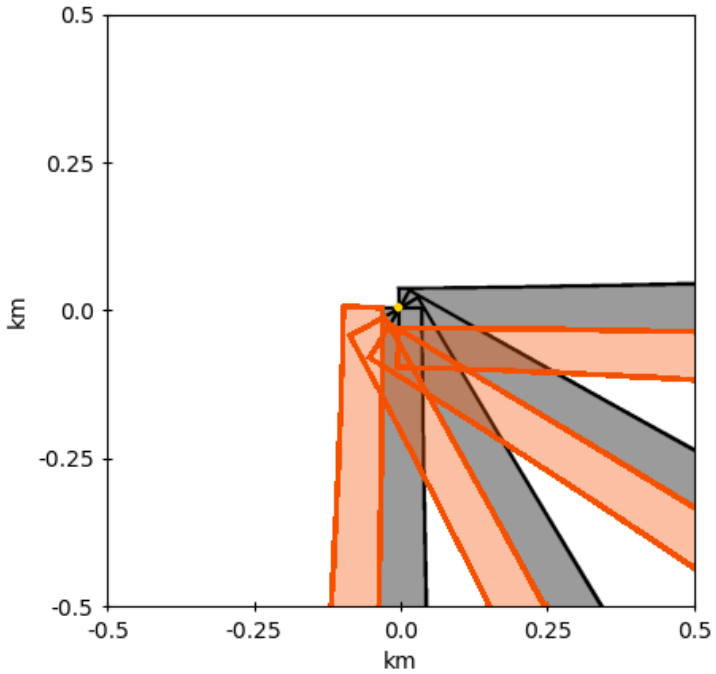

Figure 5.6: Footprints of columns 488-535 (black) and 440-487 (red) of the camera's detector at four different times during a rotation of the penetrator, projected onto the surface for a penetrator altitude of 1 km. The yellow dot indicates the penetrator's nadir.

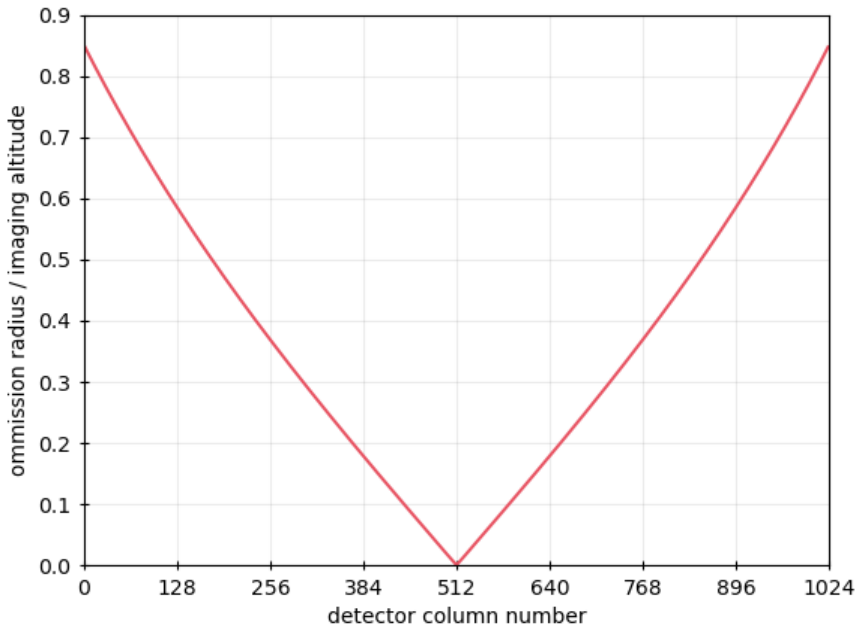

Figure 5.7: Omission radius normalised by imaging altitude as a function of detector column.

Assuming only columns 256-768 of the detector are used for imaging, an image's blind spot could be as large as 22 km across (11 km radius) at the beginning of a descent (assuming 30 km altitude). For the goals of providing both landing site context and location, such a large un-imaged region surrounding the landing site could be detrimental. As an example, determining a penetrator to have landed within a crater would be important for correctly interpreting its in situ measurements. Craters of diameter significantly less than 20 km are common throughout the solar system (Bierhaus, et al., 2005; Robbins & Hynek, 2012; Xiao & Werner, 2015). For images captured with detector columns offset from the principal point, low altitude imaging would therefore be important not just for achieving high spatial resolution, but also extending coverage as close to the landing site as possible. Capturing an image at 200 m altitude for example would reduce the diameter of the blind spot to ~150 m.

The omission of some surface by non-central detector columns is an unavoidable consequence of the imager design, and prevents multispectral or other multichannel images from covering the landing site. Given that the central portion of the detector (black footprint in Figure 5.6) does cover the landing site, the effect of the blind spot can be mitigated by careful selection of which detector columns are used for what imaging goal. An imaging filter intended for stereophotogrammetry would not benefit from seeing the

nadir, as no parallax occurs for that region of surface (see section 5.6 for more on stereophotogrammetry).

The equations introduced in this section are used in the following section to characterise the surface resolution of the camera, and in section 5.3.3 to describe the nature of motion in its imaging.

5.3.2 Ground Sample Distance

In surface remote sensing and photogrammetry, an important and often-used metric of a camera's imaging geometry and capability is its ground sample distance (GSD) (Felipe-Garcia, et al., 2012). This is the distance, measured on the imaged surface, between adjacent pixels' centres, and closely relates to the spatial resolution of a camera (though other factors such as modular transfer function also affect this (e.g. Lindler et al. (2013))). Whilst variations in a scene on the scale of a single pixel can be captured by a camera, identifying a distinct feature generally requires it to span at least two pixels, and the GSD does not therefore indicate the minimum scale of resolvable objects (e.g. Robinson et al. (2010)).

In orbital satellite imagery, the GSD is generally relatively uniform across an image due to the narrow field of view and close-to-vertical viewing perspective. In contrast, a wide angle descent camera would view the surface from a large range of perspectives and distances, and GSD would vary significantly. A useful demonstration of this is given by Karkoschka et al. (2007), in their analysis of DISR images. Their figures 2 and 3 show the wide variety of resolutions with which Titan's surface was imaged during the Huygens probe's descent. A penetrator descent camera's images would exhibit a similar variation.

In this section, the nature of the penetrator descent camera's GSD is characterised. Given the rotating perspective of the camera, its imaging of the surface can be described by three separate GSDs. Two (radial GSD and tangential GSD) are based on the instantaneous imaging geometry of the camera, whilst the other (scan GSD) is based on the camera's scan motion. Radial and tangential GSDs are the distances between the centres of the footprints of two adjacent pixels on the detector, in the radial and tangential directions respectively (relative to camera nadir). The scan GSD is the distance travelled by the centre of a pixel's footprint due to scan motion between two consecutive exposures. These GSDs are illustrated in Figure 5.8.

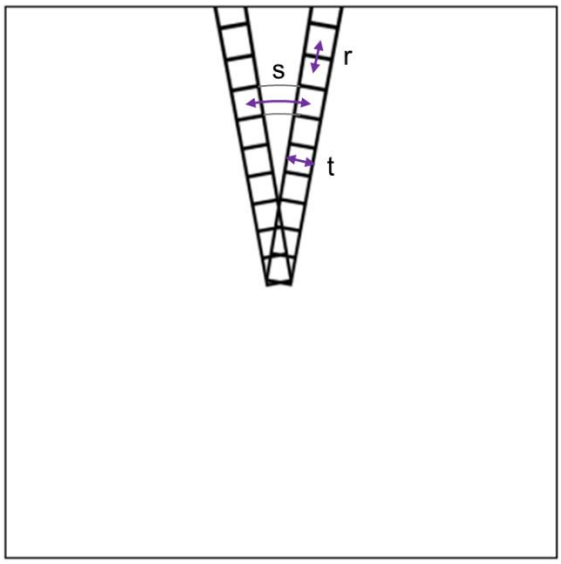

Figure 5.8: Orthographic view of the surface showing two footprints of the detector's central column of pixels, at the beginning of two consecutive exposures. The radial (r), tangential (t) and scan (s) GSDs are labelled.

Consider the geometry formed by a single pixel's line of sight (LOS), viewing from an altitude h, with an off-nadir angle of a_{nad}, the surface of a spherical body with radius R_p, as shown in Figure 5.9.

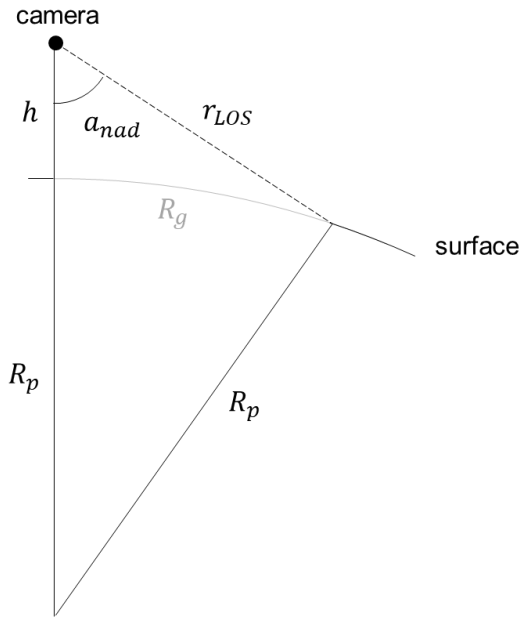

Figure 5.9: The geometry of a pixel's LOS (dashed line) to a viewed point on the surface.

The ground radius of that pixel's viewpoint on the surface is given by

$$R_g(a_{nad}) = R_p\left(\sin^{-1}\left(\frac{h+R_p}{R_p}\sin a_{nad}\right) - a_{nad}\right) \tag{5.6}$$

whilst the line of sight radius is given by

$$r_{LOS}(a_{nad}) = R_p\frac{\sin\left(\sin^{-1}\left(\frac{h+R_p}{R_p}\sin a_{nad}\right) - a_{nad}\right)}{\sin a_{nad}} \tag{5.7}$$

From these two equations, the GSDs of the penetrator camera can be described. Figure 5.5 illustrates that the GSD of a pixel depends on its position on the detector. But for the central columns of the detector (the ones which would be used for imaging), pixel geometry does not vary significantly as a function of column. Therefore, for the characterisation of the camera's GSDs, it is sufficient to consider just the central detector column as being approximately representative of the whole sensor.

It is clear to see from equation 5.6 that surface curvature is a contributing factor to viewing geometry, and of course limits the maximum viewable distance. However, surface curvature varies significantly throughout the solar system, including amongst the bodies most suited to penetrator descent imaging. In the interest of keeping the analysis of this section general, many of the plots presented will be normalised to planetary radius R_p.

5.3.2.1 Scan GSD

Each pixel's footprint sweeps along a circular path as the camera rotates. The length of the arc swept by a single pixel between two consecutive exposures defines its scan GSD, and is given by

$$GSD_s(a_{nad}) = R_g(a_{nad})\,\Delta\theta_\tau \tag{5.8}$$

where $\Delta\theta_\tau$ is the angle through which the camera rotates during a single exposure time τ, and is given by

$$\Delta\theta_\tau = \tau\frac{d\theta}{dt} \tag{5.9}$$

Scan GSD can therefore be controlled by choice of exposure length, and resolution and signal can be traded off as appropriate for the particular imaging conditions and objectives of a given camera. Figure 5.10 shows the behaviour of scan GSD as a function of both camera altitude (top) and pixel LOS off-nadir angle (bottom), for a $\Delta\theta_\tau$ of 0.2°. Distances are normalised to the target body's radius.

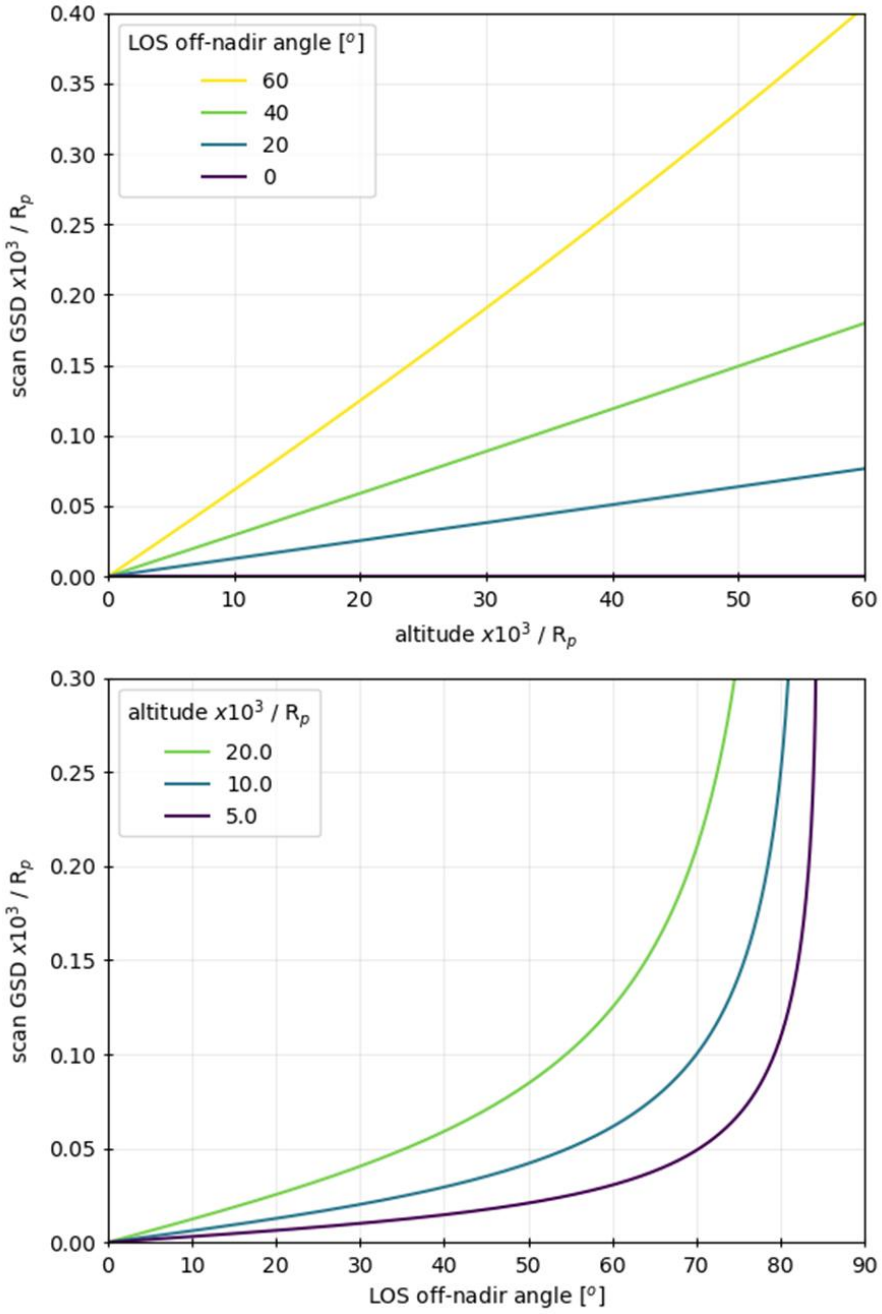

Figure 5.10: Scan GSD as a function of imaging altitude (top) for four different off-nadir angles, and LOS off-nadir angle (bottom) for three different altitudes. This is for an example 0.2° of rotation between consecutive exposures.

Scan GSD is approximately proportional to altitude, with surface curvature causing deviation from this only at high altitudes and large off-nadir viewing angles (in Figure 5.10 (top), only the 60° line is visibly not straight). A pixel viewing the nadir (off-nadir viewing angle of 0°) has a scan GSD of 0, given that the nadir coincides with the rotation

axis, but its resolution is of course still finite and is limited by its radial and tangential GSDs (discussed below). As off-nadir angle increases, scan GSD rapidly grows due to the observed point's increasing distance from the spin axis.

5.3.2.2 Tangential GSD

Equation 5.7 gives, for the off-nadir angle of a specific pixel, the distance r_{LOS} from the camera to the point of surface viewed by said pixel. The tangential GSD of a pixel is then closely approximated by

$$GSD_t = r_{LOS}\, IFOV_u \tag{5.10}$$

where $IFOV_u$ is the angular resolution of the pixel, measured parallel to the detector's rows. Plots of tangential GSD as a function of altitude and viewing angle are shown in Figure 5.11, normalised to target body radius.

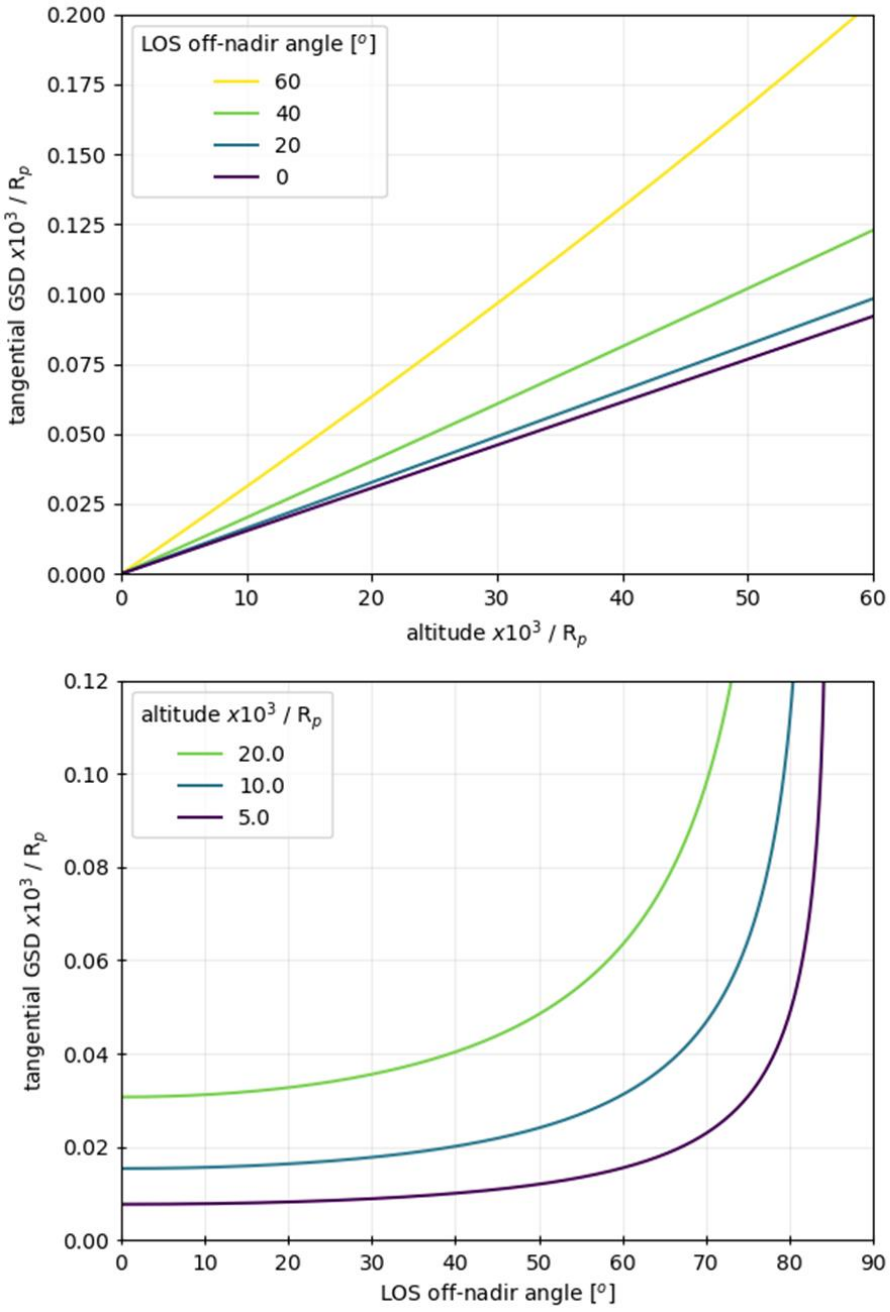

Figure 5.11: Tangential GSD as a function of imaging altitude (top) for four different off-nadir angles, and LOS off-nadir angle (bottom) for three different altitudes. This is for a pixel angular resolution of 0.088°.

The behaviour of tangential GSD is very similar to that of scan GSD, with the exception that tangential GSD tends towards a finite minimum size, rather than 0, when viewing the nadir. The relative magnitude of tangential GSD and scan GSD is a function of the

exposure time, but resolution will always be limited by at least the tangential GSD, even if scan GSD is smaller.

5.3.2.3 Radial GSD

Radial GSD is the difference in viewed ground radius due to the differing off-nadir angles of two adjacent pixels' LOSs. Thus, in conjunction with equation 5.6, it is described by

$$GSD_r = R_g(a_{nad} + IFOV_v) - R_g(a_{nad}) \tag{5.11}$$

where $IFOV_v$ is the angular resolution of the pixel, measured parallel to the detector's columns. Radial GSD, normalised to target body radius, is plotted in Figure 5.12 as a function of attitude (top) and viewing angle (bottom).

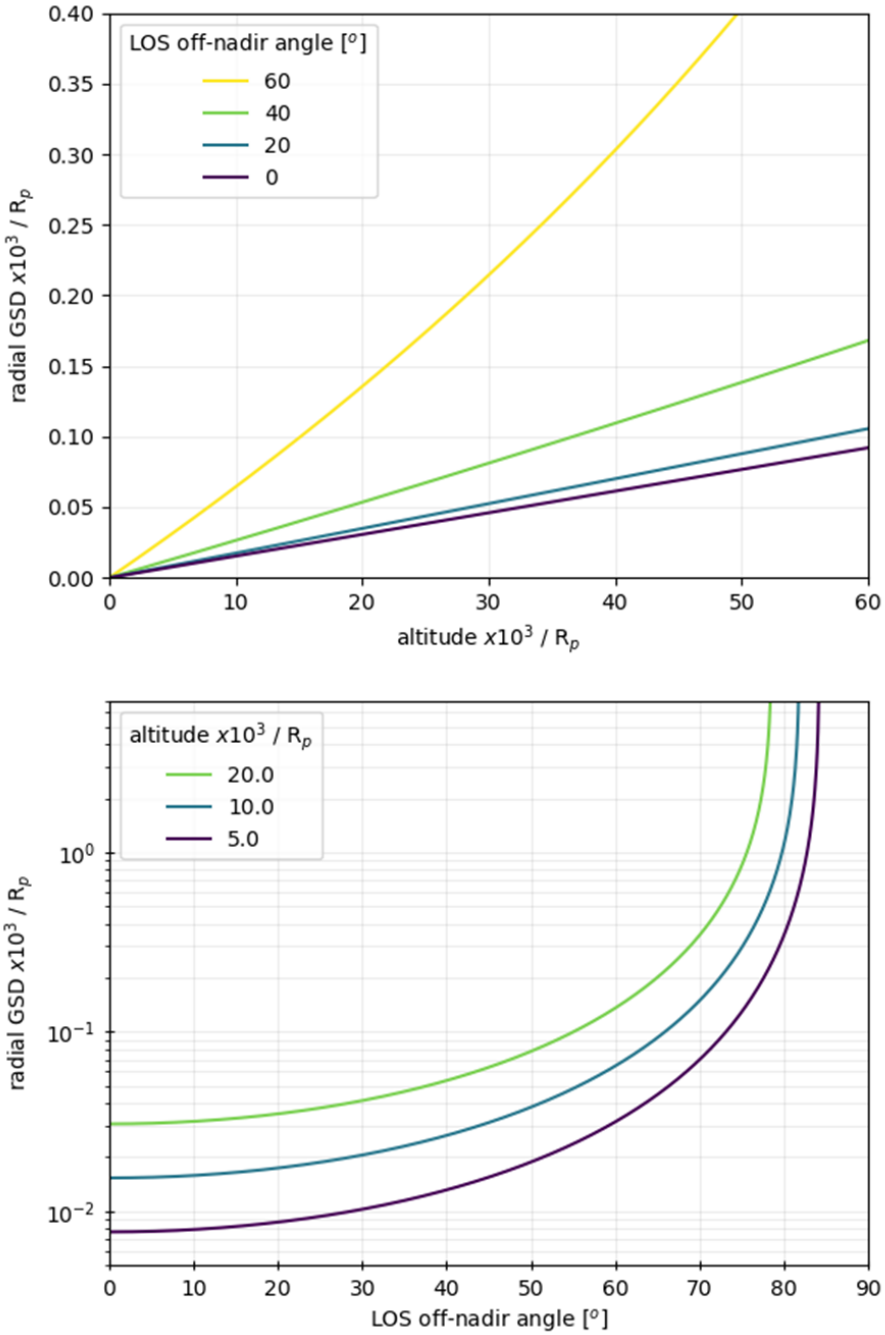

Figure 5.12: Radial GSD as a function of imaging altitude (top) for four different off-nadir angles, and LOS off-nadir angle (bottom) for three different altitudes. This is for a pixel angular resolution of 0.088°. Note the use of a logarithmic vertical scale on the bottom panel, which is in contrast to Figure 5.10 and Figure 5.11

Figure 5.12 illustrates that, as with scan and tangential GSDs, radial GSD increases approximately proportionally to altitude, but the effects of surface curvature when viewing more distant surface are evident in the 60° line of the top plot. It is clear from the bottom

plot that the effect of off-nadir angle on radial GSD is far stronger than for scan and tangential GSDs (note that Figure 5.12 has a logarithmic y-axis, whilst Figure 5.10 and Figure 5.11 have linear y-axes). The increase in both viewing distance and viewing obliquity as off-nadir angle increases results in a rapid growth of radial GSD toward the horizon of two orders of magnitude from 50° to 80°. The radial GSD quickly grows significantly larger than either scan or tangential GSD, leading to pixel footprints which are radially elongated.

5.3.2.4 GSD Comparison and Optimal Camera-Surface Geometry

To unify the above analysis of the camera's three GSDs, Figure 5.13 plots them together for two situations. The top panel shows the GSDs when the camera's imaging height is fixed, and the ground radius at which the GSDs are measured is varied. The bottom panel considers a single point on the surface with a fixed ground radius, and plots the GSDs at this point for a range of camera altitudes. In each case, the GSDs are normalised in order to generalise them to any angular resolution. The scan GSD is calculated based on an angular rotation equal to the camera's angular resolution.

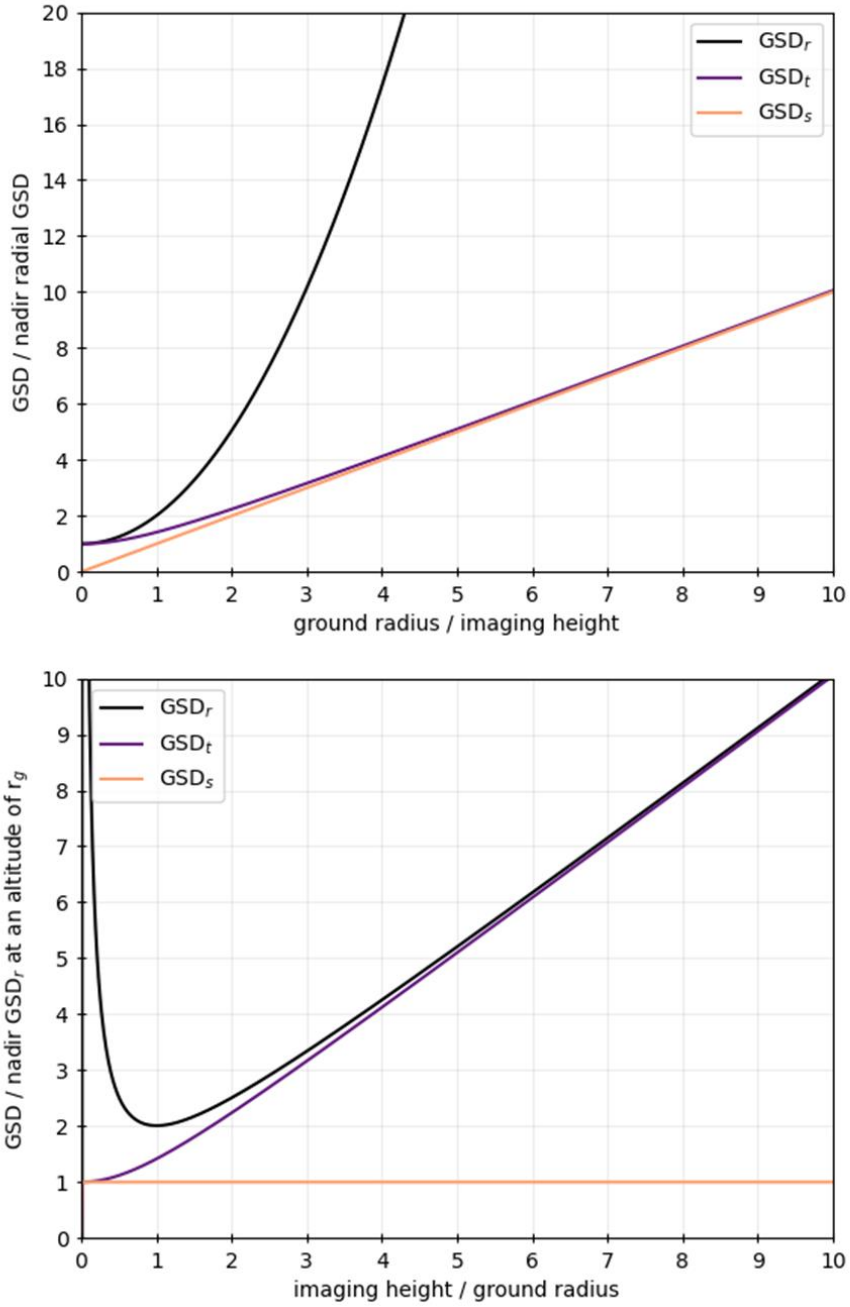

Figure 5.13: Camera GSDs as a function of viewed ground radius for a fixed altitude (top) and as a function of imaging altitude for a fixed ground radius (bottom). In both cases, GSDs are normalised in order to generalise to any angular resolution.

Figure 5.13 (top) demonstrates the contrast of radial GSD versus tangential and scan GSDs, the latter of which quickly approaches tangential GSD as ground radius increases. For ground radii of at least double the imaging height, radial GSD is more than twice as large as the other GSDs, and the disparity quickly widens with increasing

ground radius. Radial GSD is always the largest of the GSDs for any given pixel. Figure 5.13 (bottom) reveals that, whilst tangential GSD of a fixed ground radius monotonically increases with altitude, and scan GSD remains constant, radial GSD is affected by the interplay of decreasing distance between camera and surface, and increasing viewing obliquity. The result is that to image a fixed point on the surface with its minimum possible radial GSD, the camera should be at an altitude equal to the surface point's ground radius (i.e. the camera should have an altitude-to-ground-radius ratio, $h/r_g = 1$). This corresponds to viewing the surface with an off-nadir angle of 45°. This of course means a single image cannot achieve maximum resolution over the entirety of its coverage, and the entire surface cannot feasibly be imaged at its best possible radial GSD.

Instead, descent images can be timed such that every point on the surface is imaged with a radial GSD no greater than some multiple of the minimum possible. It can be seen from Figure 5.13 (bottom) that defining an acceptable factor by which radial GSD can differ from its minimum provides a range of altitude-to-ground-radius ratios from which imaging can occur, defined by its minimum, $\left(h/r_g\right)_{min}$ and maximum $\left(h/r_g\right)_{max}$. For example, a requirement that the surface be imaged with a radial GSD no more than twice its minimum possible dictates that the camera should have an altitude-to-ground-radius ratio in the range 0.3-3.7. This provides a height range from which to image a single point, but given that a descent camera images an extended surface from individual locations, it is more useful to associate a ground radius range with single heights. For an image captured at height h, the region of surface imaged with suitably high radial GSD is given by

$$\frac{1}{\left(h/r_g\right)_{max}} \leq \frac{r_g}{h} \leq \frac{1}{\left(h/r_g\right)_{min}} \tag{5.12}$$

From this, it is possible to define the necessary relationship between descent image altitudes in order to cover the entire surface with images of suitably optimal radial GSD. Let $h_1, h_2, \dots, h_n$ denote the decreasing heights of a set of images taken during a descent. For continuous coverage of the optimally imaged region:

$$\frac{1}{\left(h/r_g\right)_{max}} h_n \leq \frac{1}{\left(h/r_g\right)_{min}} h_{n+1} \tag{5.13}$$

and therefore

$$h_{n+1} \geq \frac{\left(h/r_g\right)_{min}}{\left(h/r_g\right)_{max}} h_n \tag{5.14}$$

The result of equation 5.14 is that images should be acquired at increasing frequency as the camera descends, and acceleration of the penetrator serves to exacerbate this. The majority of the surface is captured at its optimal radial GSD early in the descent, when viewing geometry changes slowly, whilst the late stages of the descent contribute the smallest absolute GSDs (highest resolutions), covering a small area.

This has further implications for how imaging throughout a descent should be planned. Once a portion of the surface has been imaged at its optimal altitude, there is potentially little benefit to continuing to image it from lower altitudes. This could yield benefits if the camera were designed to, for example, reduce power consumption and data rates by using a smaller region of its detector when imaging at low altitudes, or immediately discarding portions of images, so as to only capture the surface whose optimal imaging has not yet occurred. This is a feature of the wide-angle, large-coverage imaging of the camera concept discussed here, and is in contrast to other descent cameras with narrow, downward looking FOVs.

As an example, the Rosetta Lander Imaging System (ROLIS) aboard the Philae lander of the Rosetta mission, imaged in a ring buffer mode, continuously capturing an image every 10 seconds and only storing 7 at a time. Each time a new image was captured, the oldest in the buffer was discarded. This approach to imaging, whilst appropriate for ROLIS, would be sub-optimal for the penetrator descent camera, as valuable images of the distant surface would be discarded. This demonstrates the importance of image acquisition and data handling being tailored to the specific imaging geometry of a camera.

5.3.3 Scene Motion

With the camera both descending and rotating, the motion of the scene, as observed by the camera, is an important factor in its view of the surface. Equation 5.1 provides the position of a point in the descent camera's coordinate frame as a function of its polar surface coordinates, and equation 5.2 gives its time derivative as a function of the camera's spin rate and descent speed. Equations 2.24-2.27 (because the wide angle optics of the camera are here modelled as equidistant) therefore, in conjunction with equation 5.2, describe the scene motion across the camera's detector as it spins and descends.

Figure 5.14 plots this scene motion across the camera's detector for an example descent to Europa's surface, with a detector capturing images of 1024 x 1024 pixels at an altitude of 5 km. To demonstrate their relative effects, the scene motion due to penetrator rotation

(with a period of 1 s) and descent velocity (100 ms^{-1}) are shown in isolation in the left and right boxes respectively. See Figure 5.16 for an example of the magnitude of the scene motion when both rotation and descent velocity are considered simultaneously.

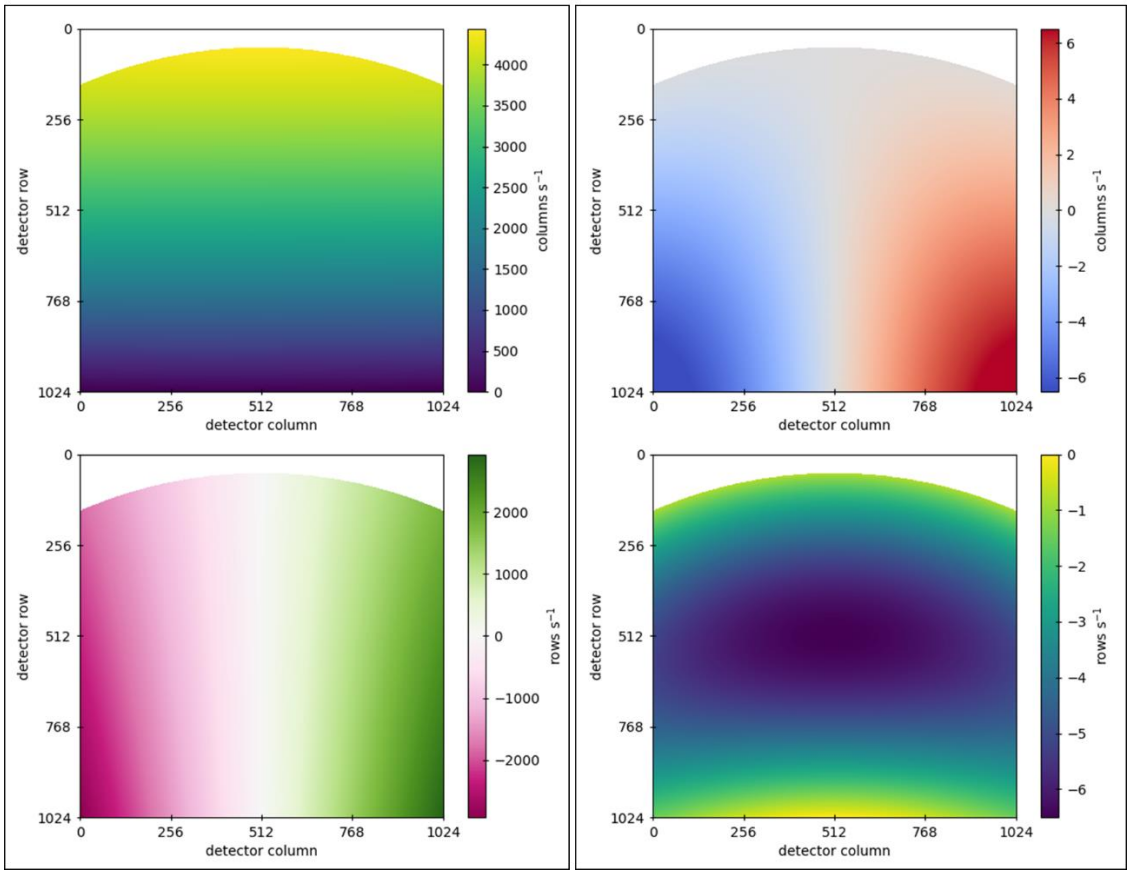

Figure 5.14: Example rate of scene motion across the descent camera's detector's columns (top) and rows (bottom) at an altitude of 5 km during a descent to Europa. (left) Motion is due only to penetrator rotation, with a period of 1 s. (right) Motion is due only to penetrator descent velocity, with a value of 100 ms^{-1}.

Figure 5.14 shows that, for the example spin period and descent velocity used, the rotation of the penetrator is a much larger contributor (~3 orders of magnitude) than its descent velocity to the motion of the scene across the detector. The contribution of descent velocity to scene motion does however become more significant as the penetrator's altitude decreases (due both to the free fall acceleration and reduced distance between camera and scene), whilst the effect of rotation does not.

To illustrate this, Figure 5.15 shows the rate of scene motion across the detector's rows due solely to descent velocity (i.e. there is no spin motion), as a function of the penetrator's altitude. For this illustrative case, acceleration is ignored and a constant

descent velocity of 100 ms^{-1} is used. For each altitude, the value plotted is the maximum magnitude value occurring anywhere on the detector's central column (number 512).

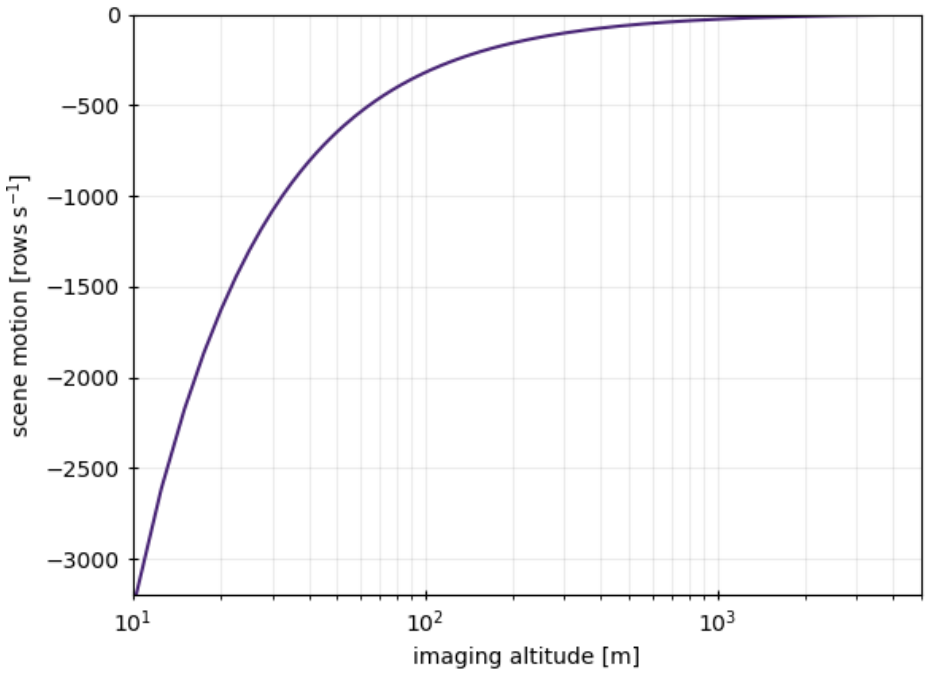

Figure 5.15: Rate of motion of scene across detector for a non-spinning camera with 100 ms^{-1} descent velocity. For each altitude, the plotted value is the value of maximum magnitude occurring anywhere on the detector's central column (number 512).

The rate of scene motion due to the penetrator's descent only becomes significant in comparison to the rate of motion due to the penetrator's spin when the penetrator is at low altitudes. For example, it does not exceed a magnitude of 500 rows s^{-1} until below an altitude of 100 m, i.e. the last seconds of descent. For a freefall descent to Europa from 30 km, this represents less than 0.5% of the available imaging time. In contrast, 98% of the available imaging time would be at altitudes above 1 km, when scene motion due to the penetrator's descent is below 50 rows s^{-1}. The vertical shift of the scene on the detector over a full 360° rotation would therefore, assuming a spin period on the order of a few seconds, not be more than a few 10s of pixels.

The penetrator's spin motion is therefore the constraining factor for exposure times and motion blur. Figure 5.16 shows the magnitude and direction of the scene's motion at an altitude of 5 km, with a penetrator spin period of 1 s and descent velocity of 100 ms^{-1}. Because of the camera's wide FOV, there is a large range of speeds observed.

To ensure no motion blur occurs anywhere within an image, an exposure time could be selected based on the fastest moving region of the scene, such that the scene does not move farther than 1 pixel's length across the detector. Formally,

$$\tau = \frac{1}{\left|\frac{d\vec{p}}{dt}\right|} \tag{5.15}$$

where $d\vec{p}/dt$ is the rate of scene motion across the detector, in pixels s^{-1}. For the example plotted in Figure 5.16, this would restrict the system to an exposure time of just ~0.2 ms. Such a short exposure time strongly limits the signal that the camera can collect.

The areas of the detector where scene motion is at its fastest are the same areas which view the surface obliquely. As addressed in section 5.3.2, this means the surface is imaged with poorer radial resolution, and pixels have rectangular footprints which are stretched radially away from the penetrator. This raises the question whether it is worth restricting the camera's exposure time to such a limiting value in order to prevent blur across pixels which already have poor spatial resolution due to their radial GSDs.

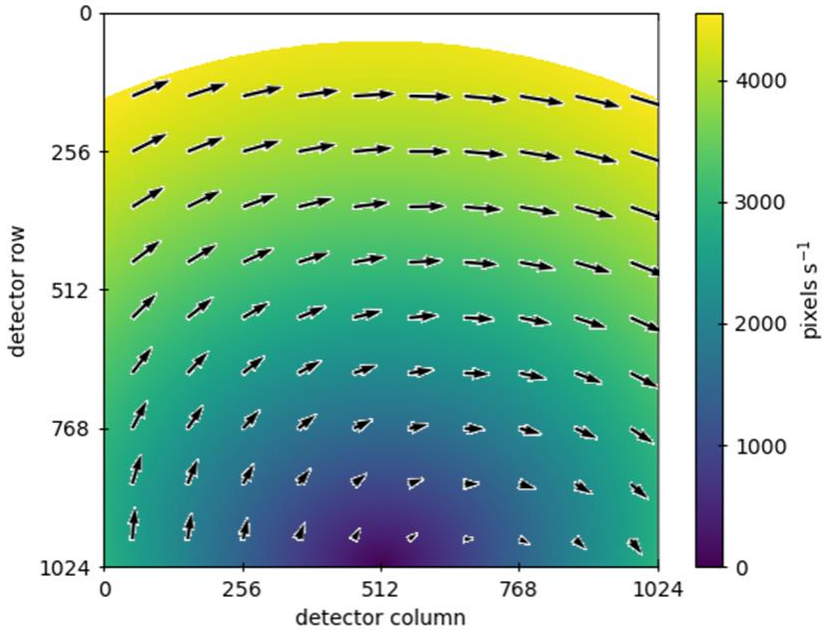

Figure 5.16: Total magnitude of scene motion across the camera's detector when the camera's motion comprises both rotation (1 s period) and descent velocity (100 ms^{-1}), for the same example case as Figure 5.14. Black arrows indicate the direction of scene motion.

An alternative metric by which to determine a suitable exposure time is the aspect ratios of the regions on the surface swept by pixels during an exposure. With the radial and tangential GSDs of pixels being fixed properties of the imaging geometry, an exposure time can be selected such that the scan GSD matches the radial GSD, meaning a pixel samples a square region of surface during its exposure. This is illustrated in Figure 5.17.

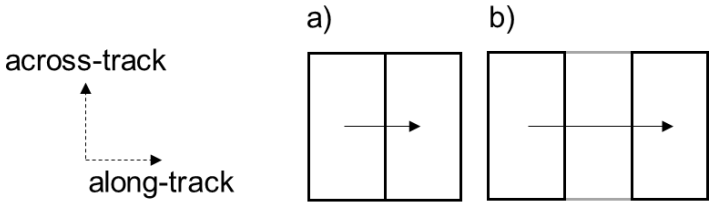

Figure 5.17: Simplified illustration of two approaches to setting exposure time based on pixel footprint (black rectangles). (a) the exposure time is set such that scan GSD matches tangential GSD. Two consecutive footprints of the same pixel butt together. (b) the exposure time is set such that the pixel moves, in the along-track direction, a distance equal to its radial GSD, i.e. scan GSD equals radial GSD. The pixel sweeps across a square region of surface.

The method of selecting exposure time to achieve square spatial resolution (Figure 5.17 b) permits longer exposure times, as the camera's viewing geometry is such that rectangular pixels are always longer in the across-track direction than the along-track direction. This is shown in Figure 5.18 (bottom), along with, in the top panel, the exposure times corresponding to the motion plotted in Figure 5.16, calculated using equation 5.15.

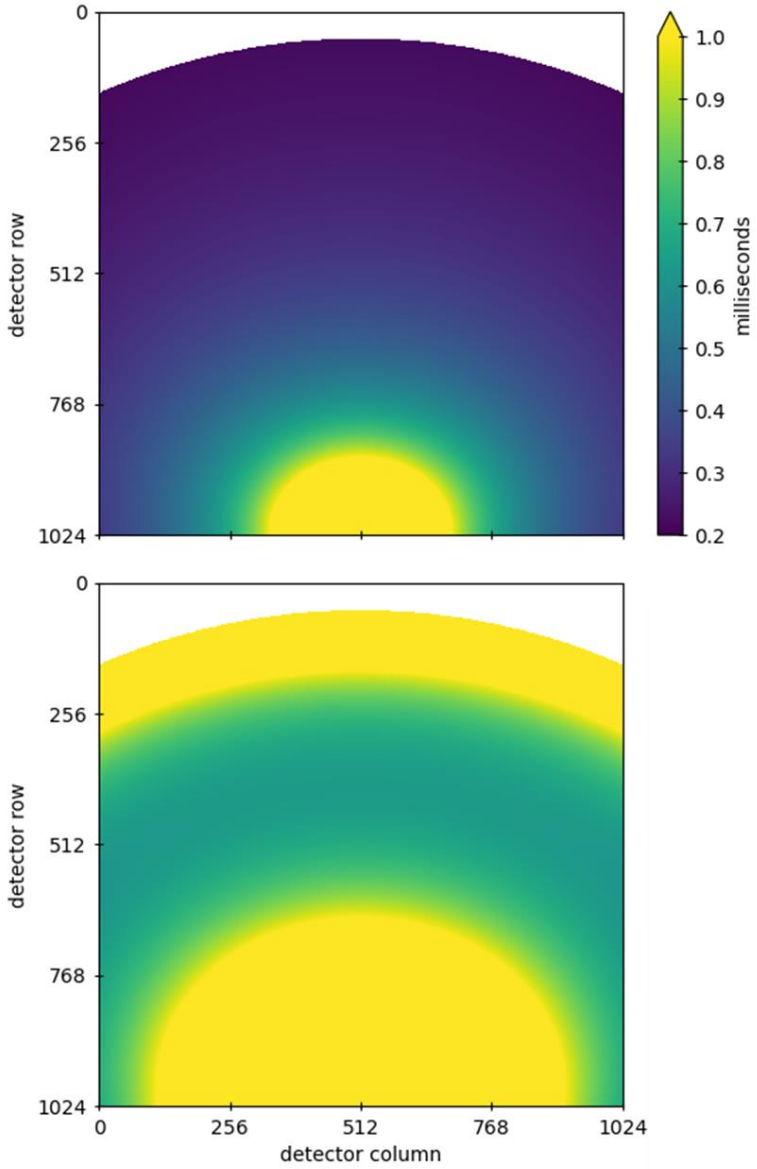

Figure 5.18: (top) exposure time for each pixel calculated from equation 5.15 and scene motion of Figure 5.16. (bottom) exposure time to achieve square pixels for same descent scenario as top.

Figure 5.18 shows that when square pixels are the criterion for exposure duration, the limiting exposure time (i.e. the smallest exposure time for any pixel) is increased to ~0.6 ms, approximately 3 times larger. In the context of planetary imagers, this is still a short exposure time, likely to limit measured signal. Figure 5.18 (bottom) shows that the exposure time required to image the surface with square pixels is not uniform across the focal plane, reaching a minimum around the centre of the detector (where the surface is imaged optimally) and growing towards the low and high numbered rows. An imaging

approach could be adopted whereby different portions of the surface are observed with different images whose exposure times need not be the same. The downward looking portion of the detector (high row numbers) observes a slower rate of scene motion, and images with higher signal could be obtained from this portion of the FOV. The exposure times presented here are indicative values, calculated for an example camera and penetrator system. A detailed signal analysis of the type carried out in chapter 6, section 6.3 of this thesis, based on both camera performance and scene properties, would allow an assessment of the signal limits of the instrument, and could inform design choices such as resolution, aperture size, spectral bandpasses and exposure durations.

5.4 Simulated Data

This section presents simulated images which represent the data that the penetrator descent camera would be expected to obtain. The following section (5.5) discusses the typical processing that would be required for detailed analysis of the camera's images, whilst section 5.6 describes an example of such analysis: stereophotogrammetry, and simulated data also appear throughout these sections.

The descent camera images displayed in the remainder of this section were produced with bespoke camera simulation software, developed as part of, and throughout this thesis. The software, which simulates cameras' views of real surfaces, is discussed in chapter 7. Important to mention here is that it utilises orthorectified images (ORIs) to reproduce the albedo variation of a surface, and digital terrain models (DTMs) to replicate its 3D shape. The simulated data presented in this section show the descent camera's view as it descends to Mars, because of the availability of high quality, high resolution Martian datasets. All of the datasets used for producing the simulated images shown throughout this chapter are summarised in Table 5.1.

Table 5.1: Summary of the data used for simulation of penetrator descent camera images presented throughout this chapter.

Dataset ID	Constituents	Source
5.1	Gale Crater CTX[a] and HRSC[b] blended DTM[f] mosaic	Persaud et al. (2019)
	Gale Crater CTX[a] greyscale ORI[g] mosaic	
5.2	Victoria Crater HiRISE[c] DTM[f]	Paar et al. (2014)
	Victoria Crater HiRISE[c] SRR ORI[g]	Tao and Muller (2016)
5.3	Victoria Crater CTX[a] DTM[f] and ORI[g]	Persaud (2018)
5.4	Mars HRSC[b] and MOLA[d] blended global DTM[f] mosaic	Fergason et al. (2018)
	Mars MOC[e] global ORI[g] mosaic	Caplinger (2002)

[a]Context Camera
[b]High Resolution Stereo Camera
[c]High Resolution Imaging Science Experiment
[d]Mars Orbiter Laser Altimeter
[e]Mars Orbiter Camera
[f]Digital Terrain Model
[g]Orthorectified Image

As introduced in section 5.2, the camera concept envisages a 90° vertical FOV, whilst rotating horizontally to scan a full 360° FOR. The result is a rectangular image which represents the entire visible surface, as demonstrated in Figure 5.19. The vertical axis covers the vertical FOV's 90° range from nadir up to horizontal, whilst the horizontal axis corresponds to the azimuthal rotation of the camera. All simulated penetrator camera images presented in this section will, unless otherwise stated, use this same representation (which is shown in Figure 5.19), and labelled axes will be omitted.

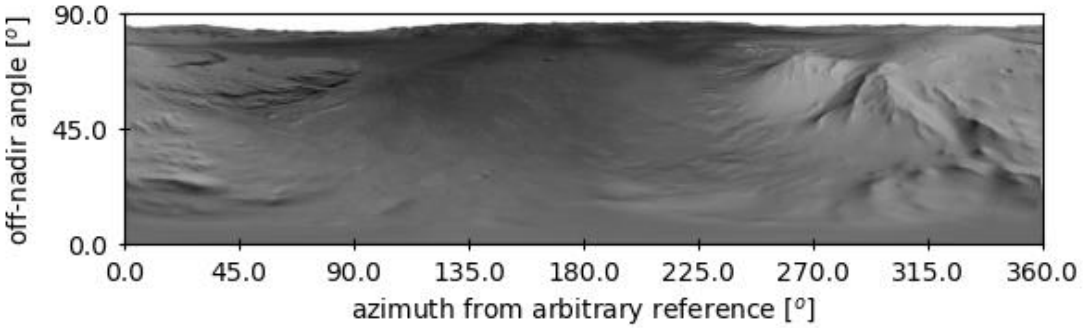

Figure 5.19: An example full scan image from the camera (dataset 5.1).

The penetrator's camera represents the surface with significant distortion. Features near to the penetrator's nadir are stretched horizontally, making them appear larger when compared to the distant surface than is the reality. Distant features have little to no

distortion, and are represented similarly to if they were imaged with a conventional framing camera. Distortion of surface features is further demonstrated in Figure 5.20.

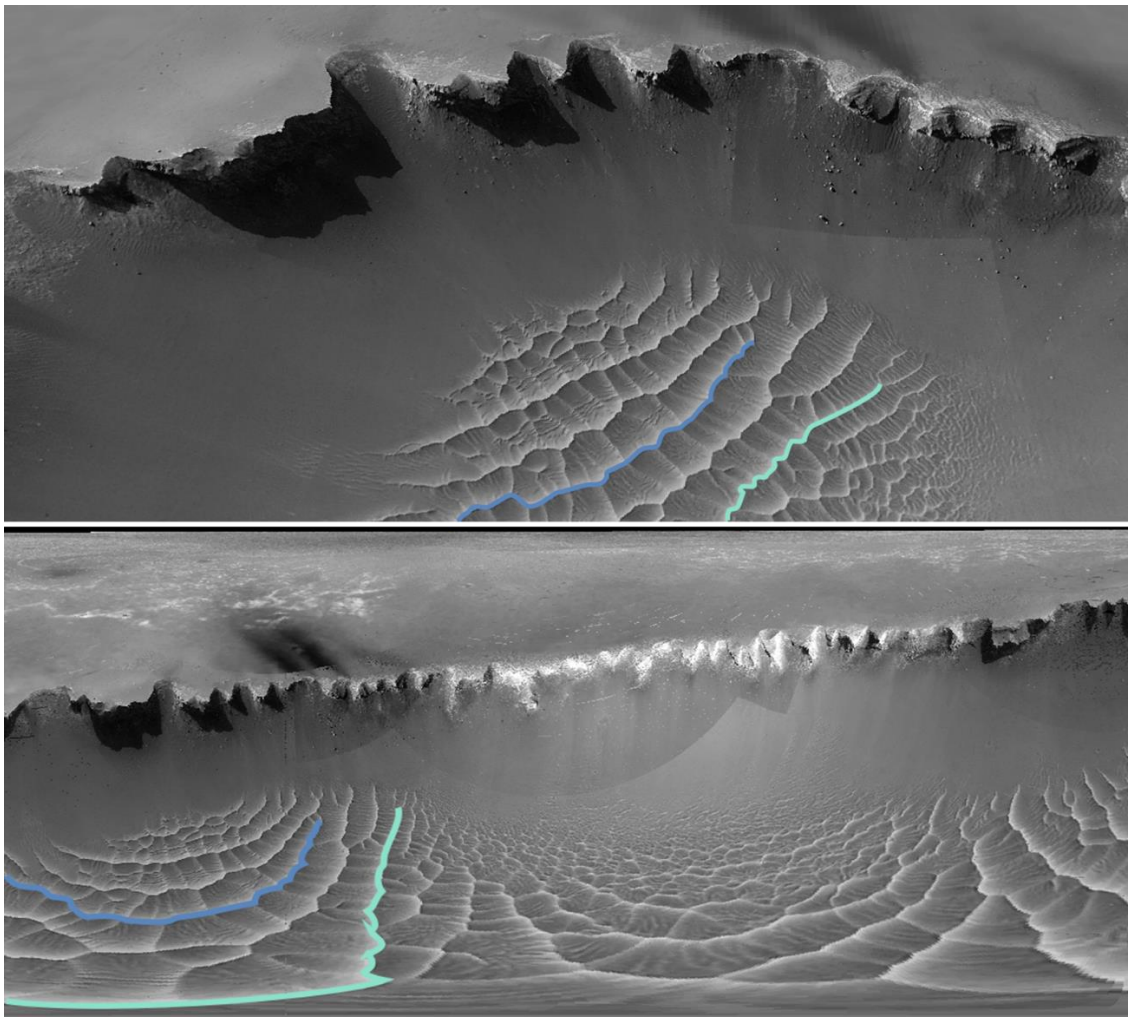

Figure 5.20: Simulated images of Victoria crater on Mars (datasets 3.2 and 3.3) from a framing camera (top) and the penetrator camera (bottom). The same two surface features are outlined in each image.

In Figure 5.20, the same scene (Mars' Victoria crater (Squyres, et al., 2009)) is shown as viewed by a framing camera (top) and the penetrator descent camera (bottom), for comparison of their representation of the surface. Two distinct features, the ridges of dunes on the crater floor, are outlined to highlight the difference in their shape between the images. Whilst the dunes' ridges describe two approximately straight, parallel lines in the frame image, they have significant curvature in the scanned image. If viewing the penetrator camera's image in isolation, it would not be immediately or intuitively obvious that the dune's ridges are approximately straight.

The radial symmetry of the camera's imaging method means that lines of equal radius on the surface map to horizontal lines in the captured images. This is most obvious from the crater's rim, which forms an approximately straight line in the scanned image.

As an example of how the camera's view of the surface changes throughout a descent, Figure 5.21 shows a simulated descent image sequence. The images within the sequence are simulated as if captured from a penetrator freefalling to Mars' Gale crater from an altitude of 20 km, with a surface nadir point of 5.886° south, 138.170° east. The images are simulated as being captured every 10 s, and the penetrator is simulated to have a spin period of 1 s. By the time of the final image capture (at 1.4 km altitude), the penetrator would be travelling approximately (ignoring Mars' atmosphere) 370 ms^{-1}, leading to significant altitude change over the course of scanning the image.

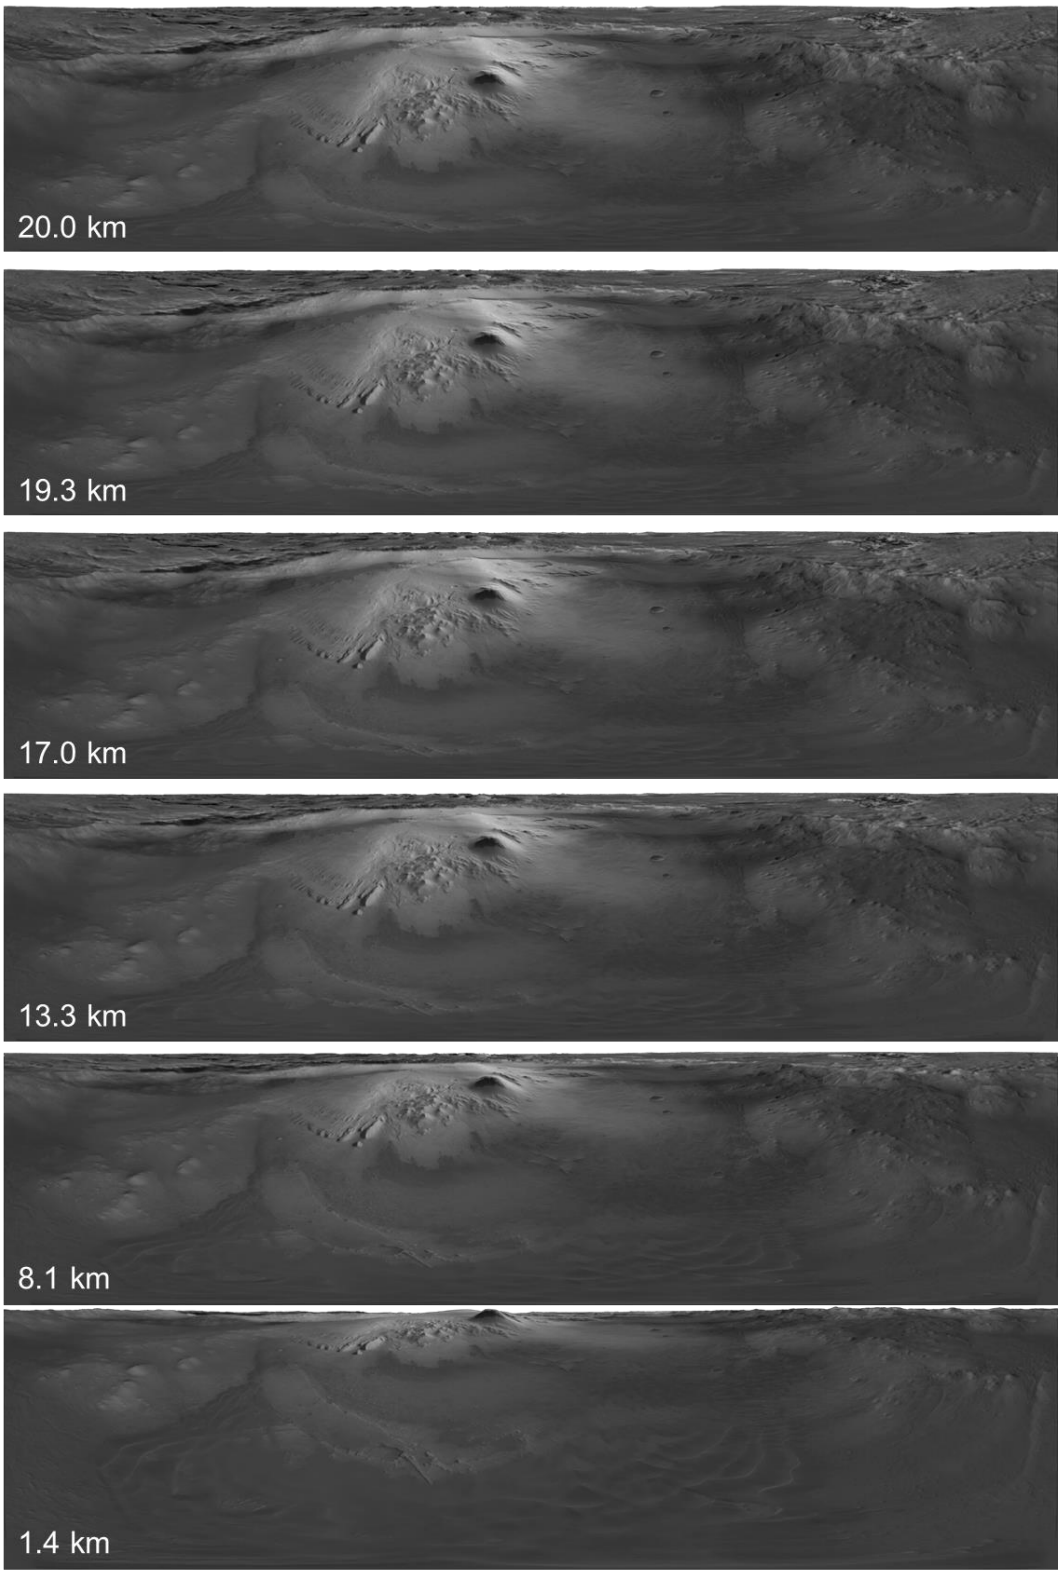

Figure 5.21: A simulated descent image sequence for a freefall to Mars' Gale crater from an initial altitude of 20 km (datasets 5.1 and 5.4). Images were simulated as being captured at 10 s intervals, with the 360° scan of each image lasting 1 s. Each image is labelled with the altitude at which its scan began.

The significant change in the camera's view of the surface is evident from Figure 5.21. At the beginning of the descent, a large area of surface is visible, and surface structures and detail can be seen throughout the images. In the late stages of the descent (e.g. the images at 8.4 and 1.4 km), much of the initially visible distant surface is obscured behind the horizon, whilst the significantly reduced GSD around the nadir reveals previously unresolvable surface detail.

The change in scene between consecutive images is minimal at the beginning of the descent, but considerable by the end, due to the combination of increasing descent speed and reduced distance between the surface and the camera. When images are taken at a fixed regular time interval, as in Figure 5.21, this results in an increasing degree of change between consecutive images as the descent progresses. Simple inspection of the images of Figure 5.21 suggests an alternative approach of increasing the imaging frequency as the descent progresses may be more appropriate from a scientific stand point (the 17.0 and 19.3 km images provide little more value than the 20.0 km image, whilst additional views between the 8.1 and 1.4 km images could provide useful extra information), and could also potentially alleviate the instrument's data and power requirements. This is corroborated by the quantitative analysis of section 5.3.2.

Indeed, the concept of optimal imaging height described in section 5.3.2.4 is well demonstrated by Figure 5.21. Note that the two small isolated craters in the right of the images are well resolved at 20 and 19 km altitude, but at and below 8.1 km are poorly resolved. However, an extenuating factor to this analysis, which Figure 5.21 clearly demonstrates, is the visibility of surface topography in the upper region of the camera's FOV when imaging from low altitudes. The profile of Mount Sharp's peak is seen most clearly in the 1.4 km altitude image. Whether to discard or retain portions of images as in the discussion of section 5.3.2.4 should therefore factor in the images' aims. For example, it may be more important for an imaging channel intended for topographic measurement to keep the full extent of its images' vertical FOVs at low altitudes. Where present, prior knowledge of the surface's topography would inform this consideration.

5.5 Simulated Data Processing

Before scientific cameras' images can be used for quantitative analysis, they must undergo several steps of processing to remove peculiarities of the instrument and allow physical properties of the scene to be recovered independently of the camera's individual response. This entails correcting for non-uniform detector response, characterising the radiometric and spectral behaviour of the camera's elements, and calibrating the

camera's optics, which will contain distortions from the ideal geometries discussed in section 2.1. These processes are required in one form or another for every spaceborne scientific camera, and the details of the steps are well described in both general and instrument specific papers (e.g. Datta et al. (2009), Thomas and Keller (1990), Tubiana et al. (2015)). This section will instead focus on the penetrator descent camera's major processing steps required which are peculiar to the camera and its imaging technique.

5.5.1 Geometric Correction

Scanning cameras such as the penetrator descent camera discussed here rely on motion to construct their images and define their imaging geometry. Any deviation from the expected movement of the camera introduces geometric errors in the images. In the case of the penetrator camera, where the spin stabilisation provides the scanning motion, variations or uncertainty in spin rate would lead to geometric errors.

To maximise accuracy of image analysis, and minimise the propagation of errors through an analysis pipeline, it's important for images to be corrected, as much as is possible, for geometric errors. As a demonstration, data analogous to descent camera images are here presented to illustrate the effect of rotationally induced geometric errors, and an investigation into their correction is outlined. Instead of synthetically generated data, this section (5.5.1) uses simulated descent images derived from drone-mounted camera footage. The drone performed a 360° rotation whilst its camera captured high resolution (4096 x 2160 pixels) video at 23.98 frames per second. The camera's vertical FOV, whilst large, was smaller than the 90° of the penetrator descent camera concept, and approximately extended from horizontal down to an off-nadir angle of 40°.

The scan imaging mode of the penetrator descent camera concept was simulated by extracting the same single column of pixels from each individual video frame (see Figure 5.22), as if captured by a linear (push-broom) sensor, and concatenating them in order to build up a scanned image. An example of such a simulated scanned image, constructed from the column of pixels coinciding with the centre of the camera's FOV (and the optics' principal axis) is shown in Figure 5.23.

Figure 5.22: A sample frame from the drone video. A single column, such as is illustrated by the green box, was extracted from each frame, and the rest discarded, to replicate the vertical scan line of the penetrator camera.

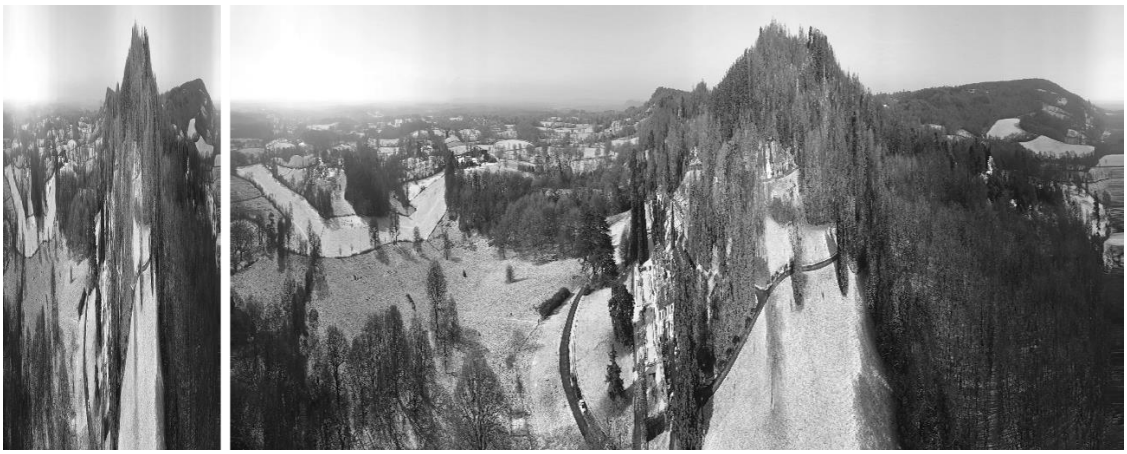

Figure 5.23: A 360° scanned image produced from the drone video footage (as described in the main text) in its raw form (left) and horizontally stretched (right) to more closely match the angular resolutions of the horizontal and vertical axes.

In its raw form, the image of Figure 5.23 is tall and narrow. This is because the number of scan lines captured by the ~30 frames per second video over the course of the 360° scan was smaller than the number of rows on the camera's detector. The right image of Figure 5.23 is uniformly horizontally stretched for better viewing, but is not geometrically corrected.

The drone's rotation rate during the video capture was not constant, resulting in a varying horizontal angular resolution in the derived images. This distorts the surface, and results in unclear and misleading representations of the size, scale and shape of surface features.

Knowledge of the rotation rate during image capture can be used to correct for motion errors and recover the correct imaging geometry. For a penetrator or other descent platform, spacecraft rotation rate can be measured by for example magnetometers, sun sensors and gyroscopes (Karkoschka, et al., 2007; Lorenz, 2010). For the analogue data presented here, accompanying rotation rate data was derived from the original drone video itself. Features were tracked within the video frames, and their rate of horizontal motion measured in order to derive a curve of relative rotation rate. The measured rotation rate is shown in Figure 5.24.

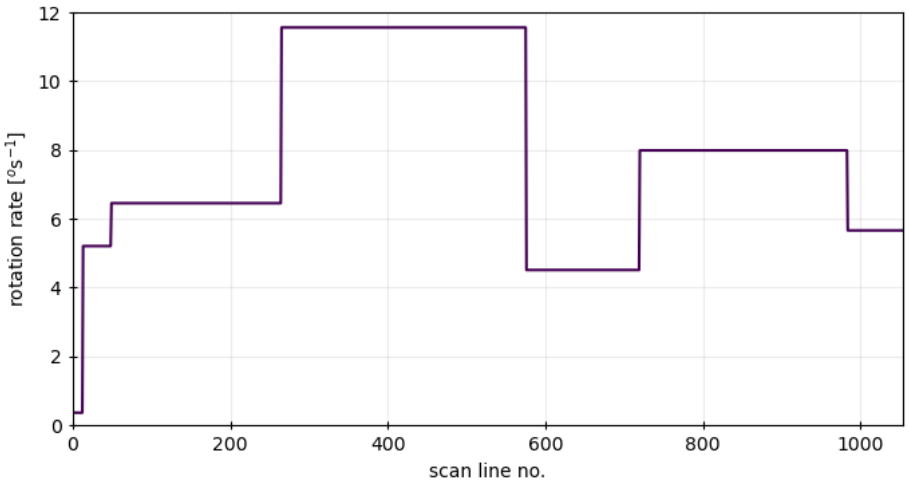

Figure 5.24: The drone's rotation rate, measured by tracking features in the captured video footage.

From the measured rotation rate, the imaging geometry is corrected by stretching each individual column of the scanned image to achieve a uniform angle per pixel of 0.0234° (in order to match the optics' vertical angular resolution). The resulting corrected image is shown in Figure 5.25 (bottom). Visual comparison of the corrected push-broom image with two frames from the original video (Figure 5.25 top) illustrate the improvement in imaging geometry, with features' positions and sizes comparing well.

Figure 5.25: (bottom) The scanned image of Figure 5.23, geometrically corrected for rotation rate. (top) Two frames from the original video are shown, aligned with their corresponding position in the push-broom image, for visual comparison of imaging geometry.

The rotation rate measurement method of tracking features in the camera's images, used here because no other sensors were available for collecting orientation data, could itself be used for a penetrator descent camera. This would require ancillary image data (i.e. not just a series of narrow image strips, but instead larger 2D images) to be captured and transmitted, so may be limited by data and telemetry budgets.

5.5.2 Image Rectification

For direct quantitative comparison of two or more descent camera images, or in order to combine them to produce another science product, it is necessary to determine the relationship between their imaging geometries. It may be that two descent images, each captured with a different region of the detector (images through two different spectral filters for example), are to be combined. Having both been imaged with a different region of the optics' FOV, they will inherently differ in their views of the surface. Alternatively, two images captured at different times during a descent might need to be compared, to measure topography for example. Whether or not they are captured with the same channel, their imaging geometry will differ due to the different viewpoints from which they were captured. In both these example cases, the correspondence (that is, which pixels or regions in one image correspond to the same pixels or regions in the other image) between the images must be found.

Image correspondence cannot be determined solely from a priori knowledge of the camera system and imaging mechanism, because it is also a function of the viewed scene's three-dimensional structure. It's for this reason that determining image correspondence allows the measurement of surface topography (see section 5.6).

Achieving full image correspondence potentially requires the comparison of every pixel within one image with every pixel in the other image. Numerous algorithms exist for measuring image correspondence, often based on windowed searches of each image, with comparison of pixel intensity distribution (e.g. Bhat and Nayar (1998); Saad and Kirakawa (2020)), though manual matching can also be employed (e.g. Soderblom et al. (2007)). Particularly for high resolution images, the problem is not trivial, manual matching is very time consuming, whilst automated matching can be computationally intensive and suffer from both false positives and negatives.

The correspondence problem can be streamlined, and its robustness increased, by reducing its degrees of freedom. Commonly, especially in the context of stereoscopic imaging, this is achieved through the process of rectification. Consider two cameras, A and B, viewing the same scene. A ray projected out the centre of one of camera A's pixels into the scene appears as a point to camera A, and a line to camera B. Such a line is known as an epipolar line, and for every one of camera A's pixels there is a corresponding epipolar line seen by camera B (providing it enters its FOV). Image rectification transforms images such that all epipolar lines are parallel with either the horizontal or vertical image axis, and share the same coordinate in all the rectified images. Searching for correspondence to a pixel in another image therefore requires searching only along its epipolar line, rather than the whole image, decreasing computation time, and reducing the occurrence of errors (Papadimitriou & Dennis, 1996). Importantly, rectification requires no knowledge of the scene.

This section presents an analogous rectification process for the penetrator descent camera's images, whereby they are aligned along one axis by transforming them to a common azimuth coordinate, requiring knowledge of the viewing geometries, but not the surface. Two scenarios are discussed: rectifying two images captured simultaneously but with different imaging channels, and rectifying two images captured at different locations along a descent trajectory.

5.5.2.1 Distinct Imaging Channels

Two images captured through different filters or imaging channels of the penetrator descent camera would be scanned by different regions of its detector, and would therefore view the surface with different imaging geometries, even if captured simultaneously. The difference in their imaging geometries would be a function of the camera's optics, and it would be important that this is well characterised by geometric calibration of the instrument (e.g. Douxchamps and Chihara (2008)). Images which are captured simultaneously have the advantage of viewing the surface from the same

position, scanning the scene with the same spin axis, and containing the same rotationally geometry (as in previous section 5.5.1). Their rectification can be achieved by transforming them to common azimuth coordinates, where azimuth is measured around the spin axis with which the images were captured.

As an example, Figure 5.26 shows two simulated descent camera images. These images are simulated as having been captured simultaneously, using two different columns (numbers 512 and 576) of the camera's 1024x1024 pixel detector.

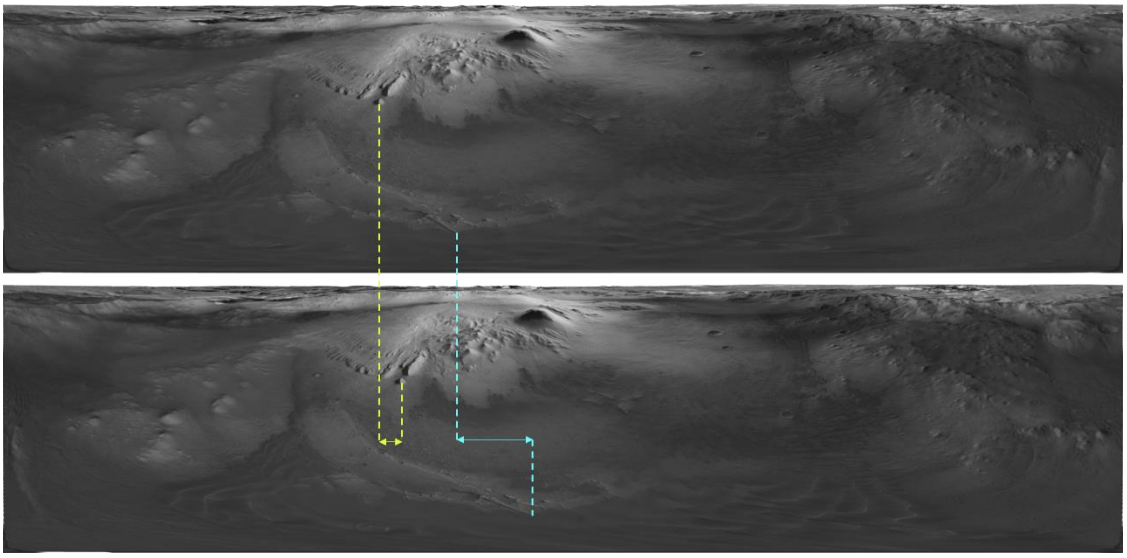

Figure 5.26: Two simulated images, simultaneously captured through different filters (datasets 5.1 and 5.4). With a 1024x1024 pixel detector, the images were simulated as being captured with columns 512 (top) and 576 (bottom). The dashed lines indicate the positions of the same two surface features within each image, along with their horizontal offsets.

Two surface features are indicated within each image by the vertical dashed lines. With the images in their original scan coordinates (as in Figure 5.19), the features appear in different places in each image, with their horizontal offset indicated by the horizontal arrows. This is a result of the different imaging geometry by which each scan was obtained. Note that the horizontal offset is not the same for each feature. Figure 5.27 demonstrates why this difference in imaging geometry occurs. The top panel shows the azimuth viewed by each of the camera's pixels relative to the camera's principal azimuth. The entirety of the central column of pixels (no. 512), which coincides with the principal point, views an azimuth of 0° (i.e. the same as the camera's principal azimuth). Detector columns which do not coincide with the principal point view not a single azimuth but a range of azimuths. The size of this range increases with the detector column's distance

from the principal point. The bottom panel of Figure 5.27 plots the difference between the azimuths viewed by columns 576 and 512, as a function of detector row number. At row 0, column 576's view direction is separated by an azimuth angle of magnitude ~5° from that of column 512. At the bottom of the detector (row 1024), this has grown to 90°.

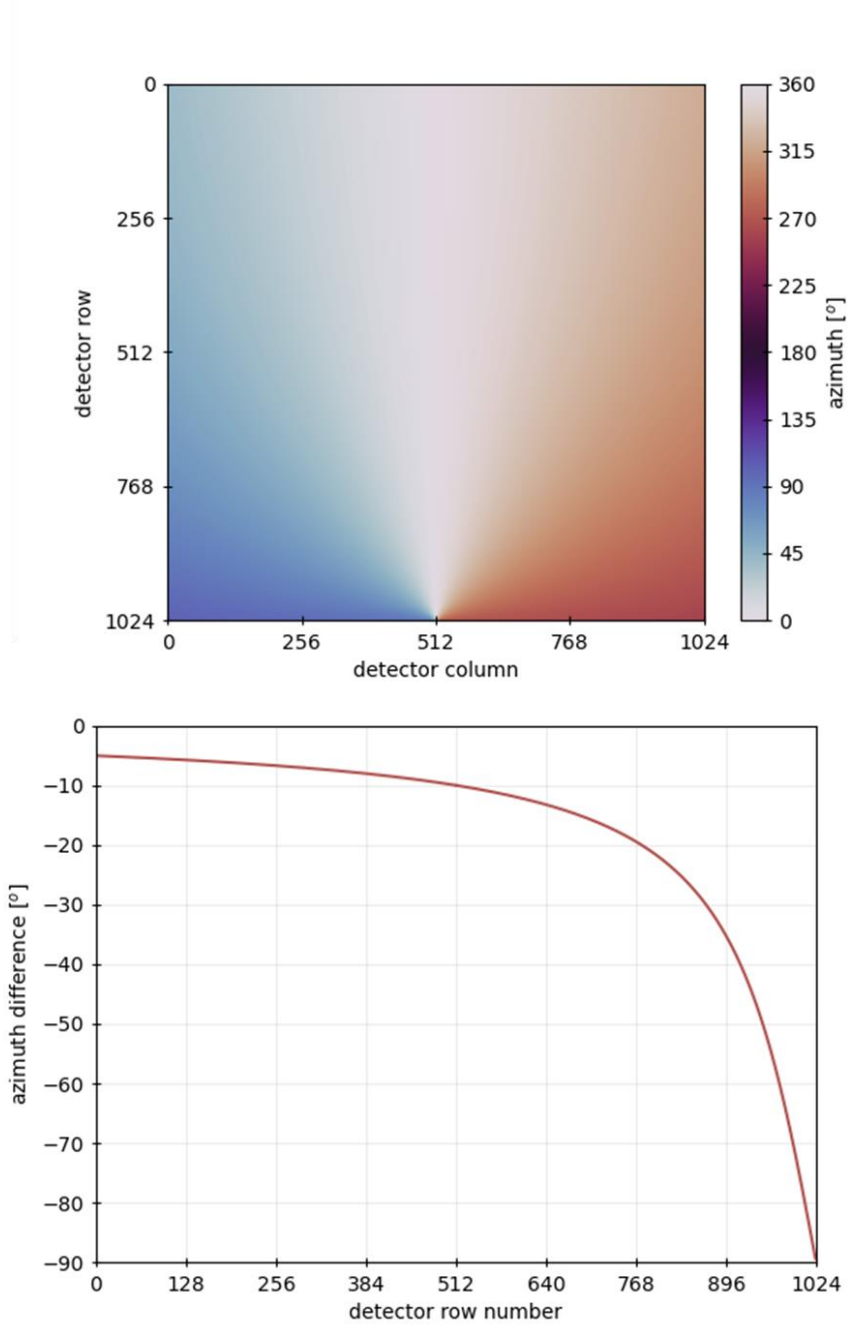

Figure 5.27: (top) The azimuth viewed by each of the camera's pixels relative to the azimuth viewed by its principal axis. (bottom) The difference between the azimuths viewed by columns 576 and 512 as a function of detector row.

The relationship of Figure 5.27 (bottom) describes the transformation required to rectify the two images of Figure 5.26. Horizontally translating the pixels of the column 576 image by their azimuth offset brings the image to the same azimuthal coordinates as the 512 image, aligning the scene to the same columns within each image. Because the horizontal axes of the images correspond to a full 360° azimuth range, translations wrap around from one end of an image to the other. Note from Figure 5.27 (bottom) that the transformation is not uniform, and distorts the original image.

The result of applying this transformation to the column 576 image is shown in Figure 5.28. The two surface features, illustrated to be unaligned in the original images of Figure 5.26, can be seen to now be aligned along the horizontal axes of the two images. Note that the image has only been transformed horizontally, and any vertical misalignment of features is retained. This is illustrated by taking the difference of the two rectified images, which shows in Figure 5.28 (bottom) that surface features still do not fully coincide.

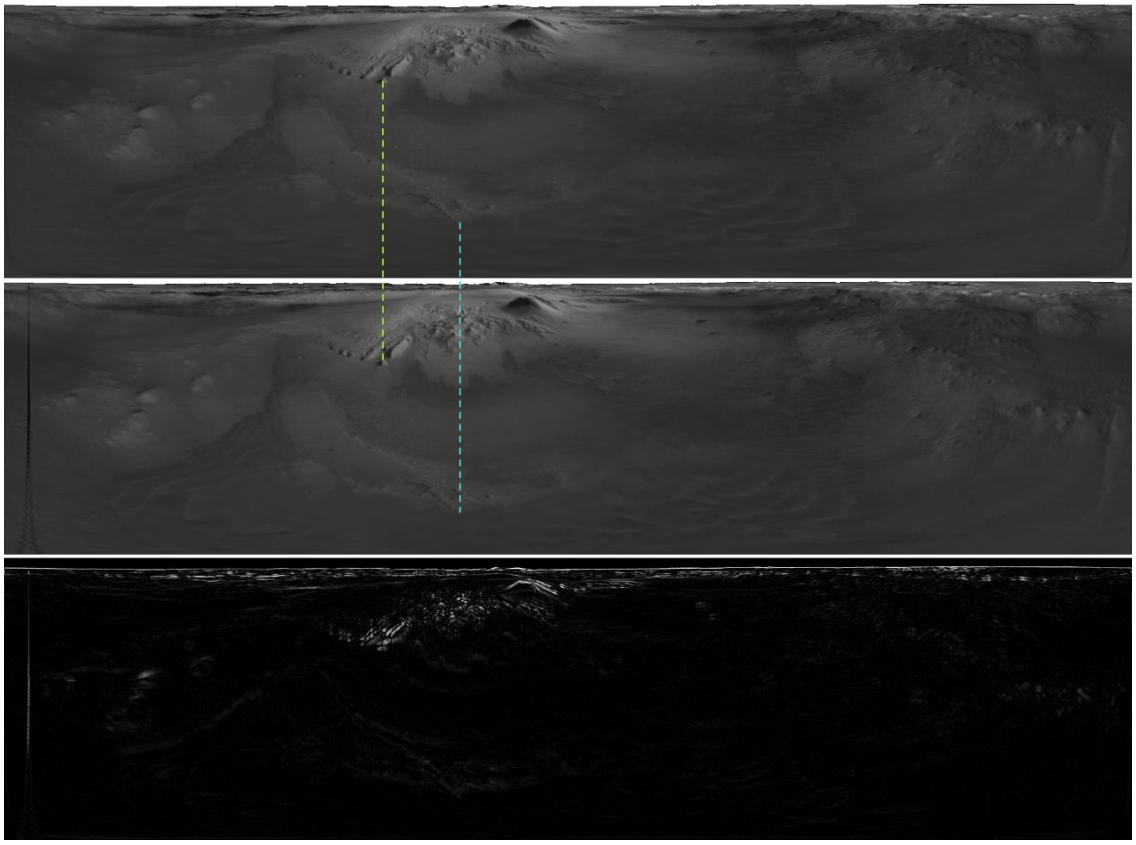

Figure 5.28: The images of Figure 5.26 after rectification (top and centre). The dashed lines indicate the same surface features as in Figure 5.26, and demonstrate their alignment. The bottom panel shows a difference map of the two images.

When it is desirable for images from multiple channels to be combined, capturing those images simultaneously and rectifying as shown here has a key advantage. Whilst the orientation and spin motion of the camera during the capture of the images may differ from the ideal or expected case, introducing errors such as those presented in section 5.5.1, it is known that each image was captured under the same conditions, and from the same location. The relationship between images captured simultaneously through different channels is a function of the camera's optics, which can be well characterised as part of the instrument construction and commissioning. In contrast, rectifying images which were captured at different times during a descent is potentially complicated by the motion of the penetrator, and the change in viewing perspective.

5.5.2.2 Temporally Separated Images

For a real mission, it may well be that a penetrator descent camera would not image the surface from the ideal case, introduced in section 5.2, of a constant, vertical descent path and spin axis. Two likely deviations from this case are a descent with a changing nadir point (i.e. the penetrator drifts horizontally whilst it descends), and a spin axis with a changing direction (e.g. nutation). This could be by design (the fuel required to deploy a penetrator with a purely vertical descent trajectory may be deemed unjustified) or because of external factors, such as atmospheric winds, though it should be noted that penetrator survivability studies suggest that the spin axis and descent trajectory should not deviate >5° from vertical at impact (e.g. Folkner et al. (2006); Lorenz (2010); Seiff (1993); Soderblom et al. (2007)). If unaccounted for, even small deviations would impact the accuracy of analysis performed on the descent images. Both types of deviation from the idealised descent have the same effect, that two images captured at different times will not share a common scan spin axis.

Two images which were captured with different spin axes cannot be rectified as described in 5.5.2.1, but can still be rectified by projecting them to common azimuth coordinates. Consider two images, each captured over 360° of penetrator rotation at two different locations along its descent trajectory. The descent trajectory has both horizontal and vertical components, and is not parallel to the penetrator's spin axis during either of the image captures. An axis can be defined by a straight line joining the locations of each image capture, as shown in Figure 5.29.

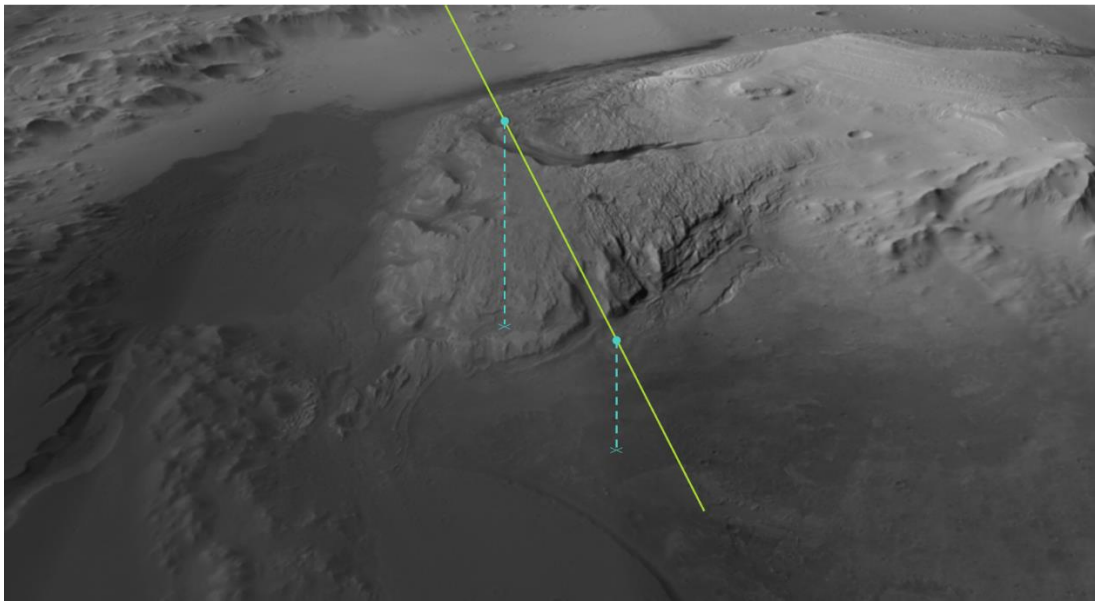

Figure 5.29: Illustration of a non-vertical descent to a surface. Solid circles indicate two imaging locations, with corresponding nadirs shown by the crosses and dashed lines. The solid green line indicates the axis defined by the two imaging locations.

Two images, simulated as having been captured during a non-vertical descent, at locations 5.21°S, 137.30°E and 5.22°S, 137.33°E at altitudes of 9 km and 6 km respectively, during a descent to Mars, are shown in Figure 5.30.

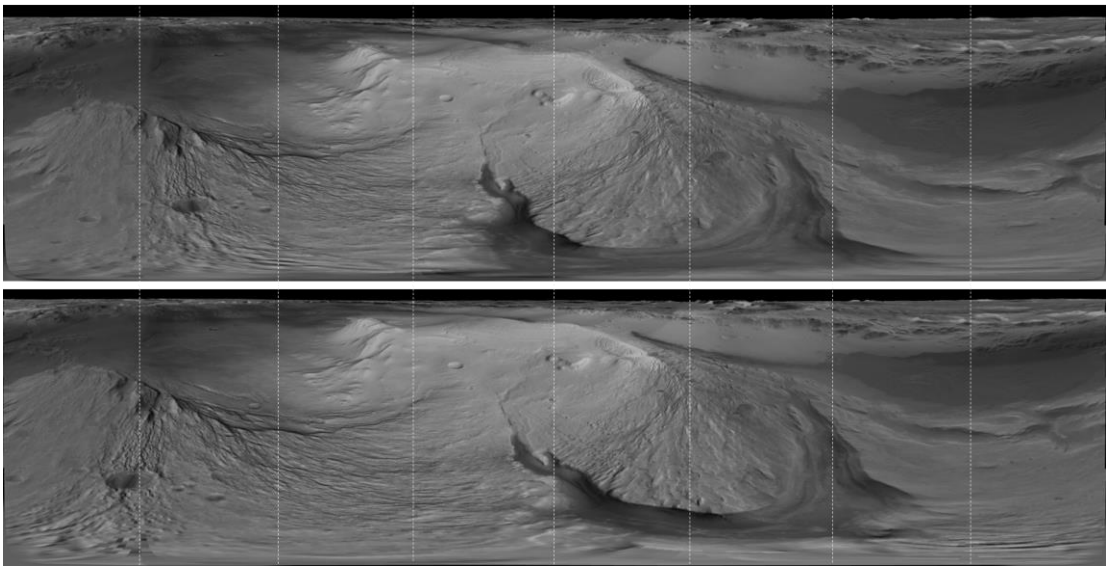

Figure 5.30: Two simulated images, captured during a non-vertical descent (datasets 5.1 and 5.4). Vertical lines are for comparison of surface features' horizontal positions in the two images.

Providing the orientation and position of the penetrator can be suitably determined or estimated for the times of each image capture, the connecting axis can be used to project the images to common azimuth coordinates. Estimation of these properties would potentially be possible, after the descent, using measurements from sensors onboard the penetrator such as magnetometers, Sun sensors and gyroscopes, as well as measurements, such as Doppler tracking, taken from Earth or an accompanying spacecraft (e.g. Aboudan et al. (2008); Folkner et al. (2006)). This method relies on the assumption that the penetrator's position and spin axis do not change significantly over the course of a single image capture (i.e. the time period of any nutation or precession is significantly longer than the spin period).

The connecting axis can be used to define two spherical coordinate systems, centred on each imaging location. The two coordinate systems share as their z axes (the axis around which azimuth angles are measured) the connecting axis. Therefore, a point in space will have the same azimuth angle in each coordinate system. Rectification of an image is achieved by transforming its x-axis to azimuth, as measured around the connecting axis, and transforming its y-axis to off-axis angle, where the off-axis angle is the angle measured from the connecting axis, to the point of interest, around the imaging location. The geometry described here is illustrated in Figure 5.31. Crucial to the rectifying process is the fact that a point's azimuth, when measured around the connecting axis, is the same from both image locations. Off-axis angle is location specific.

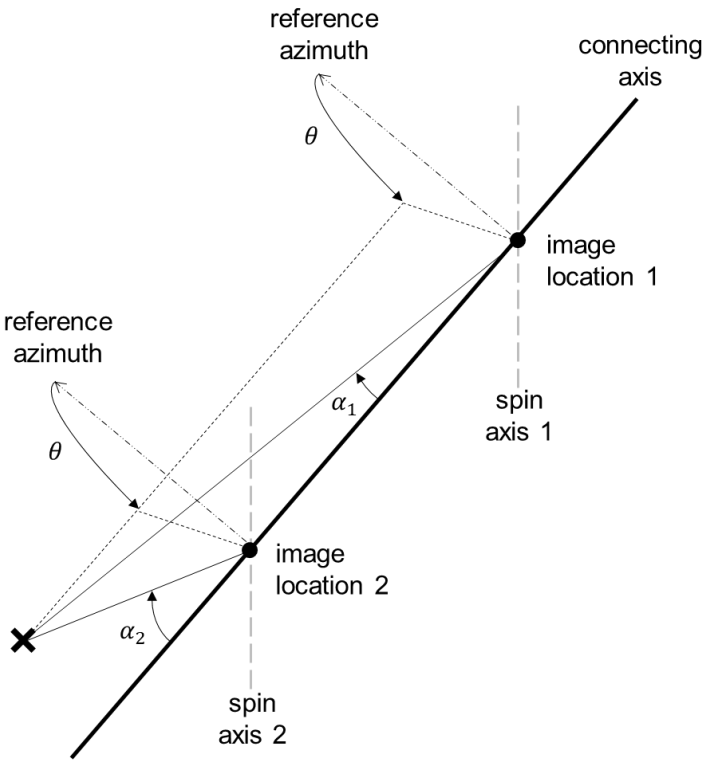

Figure 5.31: The geometry of two images and their rectification using their connecting axis. Images are captured at locations indicated by the solid black circles, around spin axes indicated by the thick grey dashed lines. The azimuth (θ) and off-axis (α_i) angles of a point in space (×) as measured from each imaging location, using spherical coordinate systems defined by the connecting axis, are shown.

The rectified versions of Figure 5.30's images, transformed as described above, are shown in Figure 5.32. Note that any given surface feature, as intended, appears at the same horizontal coordinate in both images. The distortion of the transformation vertically compresses and stretches portions of the original images, altering, significantly in places, the spatial resolution of the data.

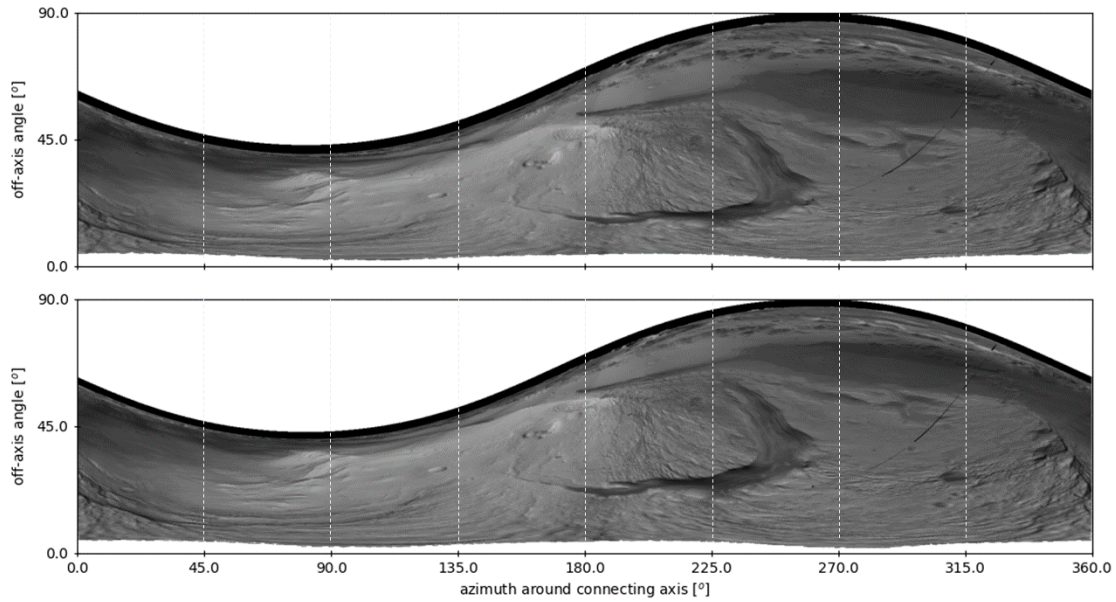

Figure 5.32: The images of Figure 5.30 after having been rectified, using their connecting axis, according to the description in the main text.

5.5.2.3 Further Considerations

The two rectification processes described here address the two main sources of image misalignment: the FOV's non-uniform imaging geometry, and the change in imaging position and pose throughout descent. For both examples, an 'idealised' image acquisition scenario has been assumed, in which every image can be treated as being captured from a single point, the penetrator spin axis is fixed for the duration of each image capture, and the accurate determination of the penetrator position and spin axis were possible before carrying out rectification. These assumptions are justified, as this chapter is presenting a description and analysis of a general imaging concept in its very early stages, and a study of the complex and varied effects of multiple error sources is beyond its scope.

However, it is important to note that a real implementation of this imaging technique would of course deviate in some way from the ideal case. Likely mechanisms for this include a varying spin axis orientation with a time period comparable to or smaller than an image capture duration, and uncertainty in the spin axis and penetrator location at the times of image capture. This would impact the efficacy of the rectification processes described here. Achieving perfect horizontal alignment (where any point in the scene shares the very same column in both rectified images) would be unlikely. Instead the images may be expected to align to within some horizontal error range, whose magnitude would depend on the accuracy with which the imaging geometries are known.

An alternative approach to constraining the descent camera's imaging geometries, the relationship between its images and the penetrator's location and attitude, which is not discussed in detail here, is bundle adjustment (Furukawa & Ponce, 2009; Läbe & Förstner, 2006; McLauchlan & Jaenicke, 2002). Bundle adjustment is a photogrammetric technique for measuring, from a set of multiple images of the same scene, the 3D structure of the scene, and simultaneously computing camera positions at the time of each image capture. Its use with spacecraft imagery, including descent images, is well established (Remetean, et al., 2016; Scholten, et al., 2019; Soderblom, et al., 2007). A bundle adjusted algorithm tailored to the nature of the descent camera may be effective in cases where significant geometric errors are present within images, and the simple rectification processes of sections 5.5.2.1 and 5.5.2.2 may not be sufficient. The implementation of bundle adjustment with the penetrator descent camera's images is beyond the scope of this thesis, but could be a valuable future study.

5.5.3 Co-registration

A key benefit of the descent camera concept's use of scan imaging (alongside achieving large surface coverage), is the opportunity to capture images through multiple optical filters in spite of the rapid scene motion and the short duration of operation. Deriving composite products such as colour images or polarisation maps requires the combination of images from multiple channels, a prerequisite of which is their co-registration. This section will demonstrate a simulation of that process, and illustrate the descent camera's fundamental concept of multi-channel imaging, using the analogue drone data first discussed in section 5.5.1.

The original video data was captured in colour by the drone mounted camera, and comprised the standardised red, green and blue channels. For the geometric correction discussion of section 5.5.1, greyscale images derived from only the green channel were used. In the work presented here, all three channels were used separately, to simulate the capture of three separate images through different spectral filters. Thus, three push-broom images were independently constructed (in the same way as described in section 5.5.1), one for each colour channel. The red, green and blue push-broom images were constructed from detector columns 1984, 2048 and 2112 respectively (Figure 5.33). This represents very well a typical dataset that the penetrator camera could return.

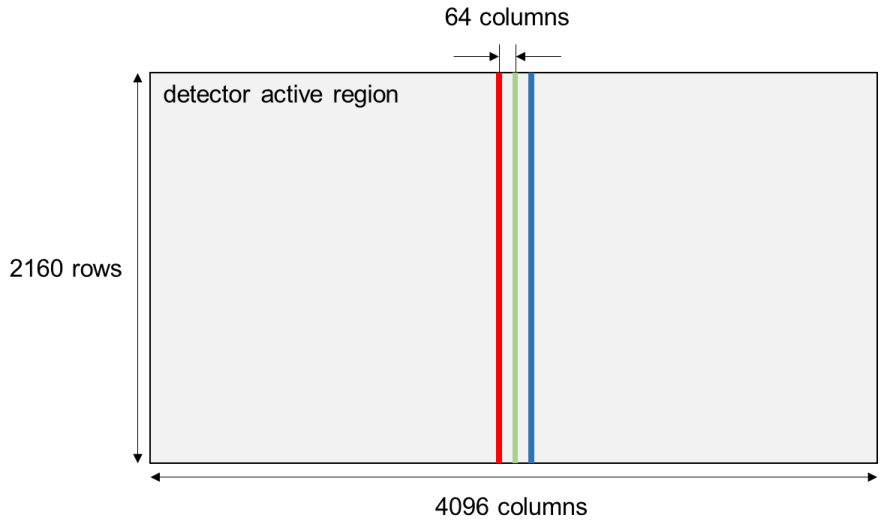

Figure 5.33: Illustration of the detector columns used to construct three push-broom images for multispectral imaging.

Images which are captured simultaneously through different filters see the surface from the same viewpoint, and therefore view any point in the scene with the same azimuth and off-nadir angle (assuming a vertical spin axis). After geometric correction and rectification as outlined in 5.5.1 and 5.5.2 respectively, co-registration is a matter of transforming the images to common off-nadir coordinates. The off-nadir coordinates viewed by the camera's pixels are, much like the viewed azimuths (Figure 5.27), a function of the optics' imaging geometry, whose a priori characterisation and calibration would be necessary. As an example, for the ideal equidistant 90° x 90° FOV of the descent camera, the nature of the variation in pixel off-nadir angles is plotted in Figure 5.34.

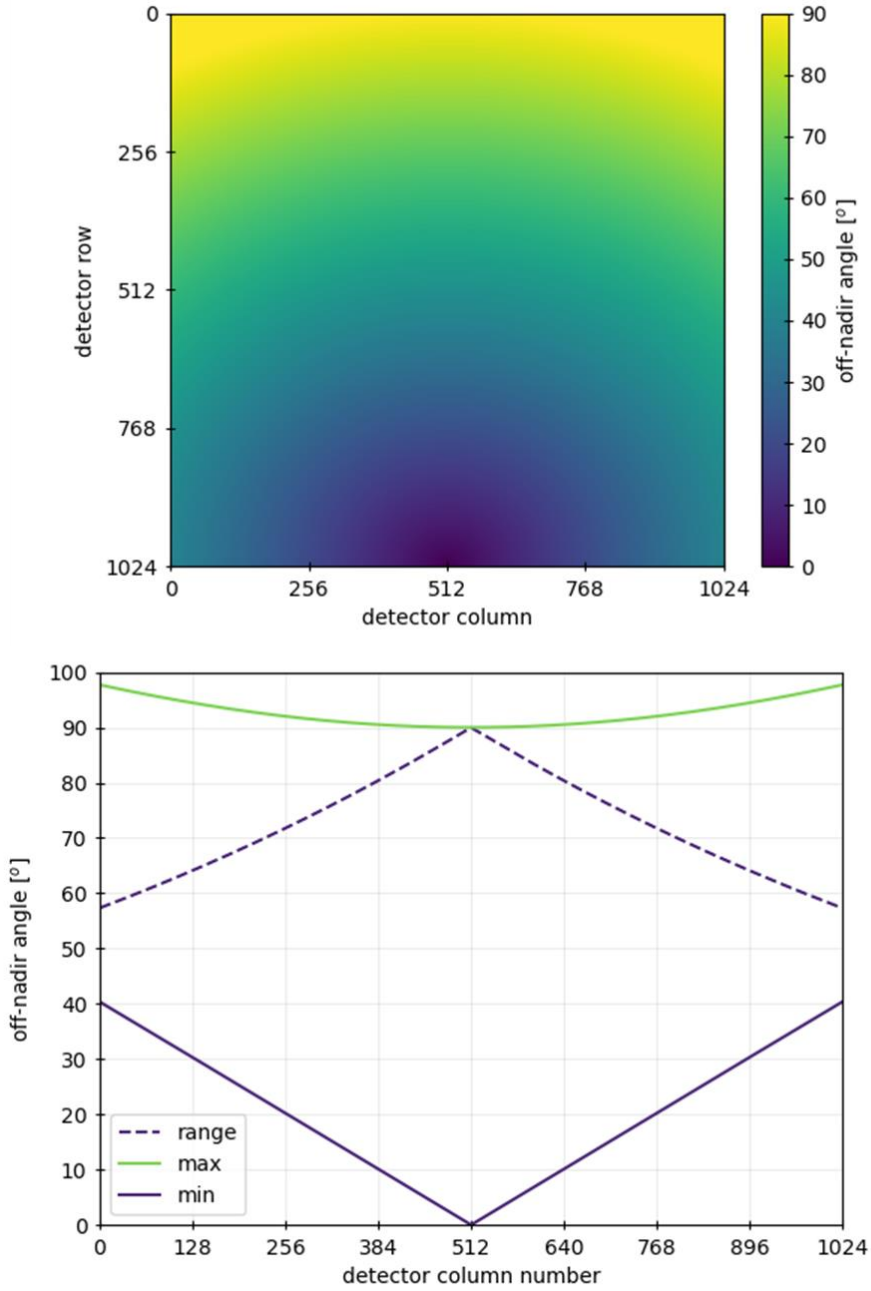

Figure 5.34: (top) the off-nadir angle viewed by each of the descent camera's pixels. (bottom) The minimum and maximum off-nadir angle, and their difference, viewed by each of the camera's detector columns.

Figure 5.34 indicates the significant variation across the columns of the detector in off-nadir angles viewed. The central column, aligned with the principal point, covers the largest range of off-nadir angles, and therefore has the lowest off-nadir resolution. To avoid downgrading image resolution when transforming them to common off-nadir

coordinates, images should be up sampled to match the image of best off-nadir resolution, rather than down sampling the highest resolution images.

The three individual drone-derived push-broom images were geometrically corrected, rectified, and transformed to common off-nadir coordinates (based on the ~50° FOV of the drone camera's optics, not the descent camera optics of Figure 5.34). A multispectral RGB image was then created comprising as its red, green and blue channels these red, green and blue push-broom images respectively. The resulting multispectral image is shown in Figure 5.35 (a). Also shown is a zoomed portion of the product (b) alongside the same zoomed portion of a colour image produced in the same way, but with the omission of the rectification and off-nadir alignment steps (c). The comparison of (b) and (c) shows the significant misalignment present when no rectification is performed. Image (d) shows a difference map between the derived multispectral image (a), and a colour image derived from column 2048 of the raw frames, retaining the original RGB channels. The latter contains perfect alignment of the three colour channels, providing a reference. Qualitatively, the difference map shows good agreement between the derived and reference multispectral images. Some differences in the red and blue channels can be seen, occurring at edges of features. These, and the colour artefacts in (b), are indicative of some remaining misalignment between the three colour channels. These are to be expected, as the optics of the drone camera were not characterised, and were instead modelled with ideal equidistant behaviour. Nonetheless the principal of the instrument's multichannel imaging is well demonstrated.

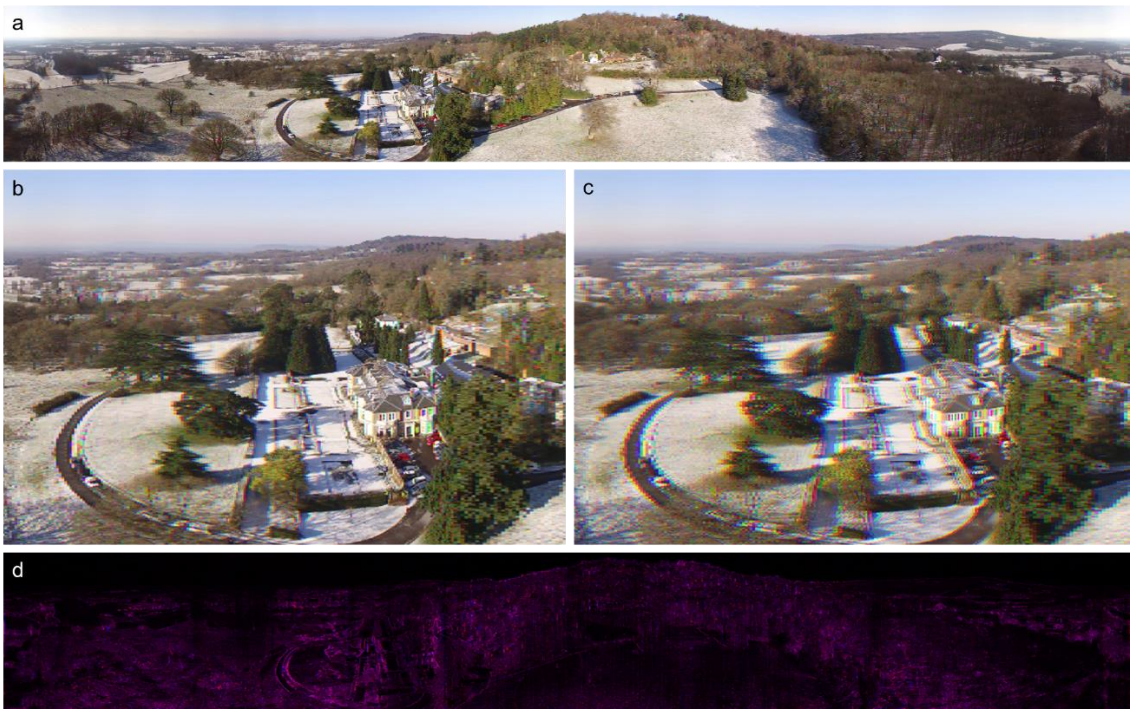

Figure 5.35: (a) Multispectral image produced from three separate channels, as described in the main text. (b) Zoomed section of (a). (c) Zoomed section of multispectral image produced without performing rectification or alignment, covering same region as (b). (d) Difference map (with colours stretched for clarity) of (a) and a colour image derived from column 2048 of the raw video frames, retaining the original RGB channels.

Whilst the data used to simulate the co-registration and combination of images here deals with combining colour filters to produce a multispectral image, the same process would apply to other multi-channel products. For example, the combination of images captured through different polarising filters to produce a map of polarisation.

Because the three channels that were combined to produce the multispectral push-broom were captured simultaneously, their registration was simplified by the fact they all viewed the surface from the same location. When combining images captured at different locations, the process of co-registration is not so simple. The difference in position of a point in the scene between two images becomes a function of the point's position, i.e. the surface shape. Accordingly, co-registration of satellite imagery often utilises, where available, prior knowledge of the topography (in the form of a surface shape model) to align images onto common coordinates (e.g. Leprince et al. (2007)). If the surface's shape is not known, this approach is not available. Instead, correspondence between images must to be measured from the images themselves. If the viewing geometry of

each image is known, this not only facilitates the co-registration of images, but also the measurement of the surface's topography. Thus, the process of measuring penetrator descent image correspondence will be covered in the following section, in the context of inferring surface topography.

5.6 Descent Stereophotogrammetry

In this section of chapter 5, a study of the penetrator camera's application to measuring surface topography is presented. The work and material discussed here is also published in Brydon et al. (2021).

Images of planetary surfaces can reveal many properties: reflectivity, composition, geological activity, age. But a single image gives a flattened representation of the surface, and the possible analysis can be somewhat elevated by determining the surface's topography (i.e. its three dimensional shape). From this, additional geological insight is gained.

Commonly, surface topography is measured stereoscopically with orbital imagery in the visible spectrum. 3D surface structure is recovered by comparing two (or more) images with overlapping surface footprints and distinct known viewing geometries. Such stereoscopic datasets are commonly acquired by imaging the same region of surface during two slightly different orbits (e.g. Henriksen et al. (2017)), or with a multi-view camera (e.g. Jaumann et al. (2007)). In addition to stereoscopy, surface topography of solar system bodies has also been measured by methods such as orbital LIDAR (Smith, et al., 2001) and photoclinometry (Kirk, et al., 2003).

High resolution global-coverage topographic mapping has been carried out for the Earth (e.g. Vernimmen et al. (2020)), the Moon (Scholten, et al., 2012) and Mars (Tao, et al., 2018), whilst numerous other solar system bodies including Phobos (Willner, et al., 2014), Europa (Steinbrugge, et al., 2020) and Titan (Soderblom, et al., 2007) have also been mapped, but with significantly lower coverage and/or resolution. Additionally, local-coverage, very high resolution topography measurements have been derived from rover imagery at Mars (Barnes, et al., 2018).

Although less prolific than orbital mapping, descent images offer an alternative perspective from which to measure surface topography. The changing imaging altitude provides a vertical stereo baseline from which 3D geometry can be reconstructed. The topography of the Huygens probe's landing site on Titan was obtained from stereo images captured during the last 40 km of its descent (Soderblom, et al., 2007). Localised

lunar topography was measured with centimetre accuracy using the descent images of the Chang'e 3 and 4 landers (Li, et al., 2015; Jia, et al., 2018). Descent images captured by Rosetta's Philae lander were used to derive a DTM of its landing site on comet 67P/Churyumov-Gerasimenko's surface (Mottola, et al., 2015).

Likewise, the vertical descent of a penetrator offers the opportunity to measure surface topography from its descent camera's images, potentially assisting with contextualising the penetrator's in-situ measurements (taken after impact) and constraining the location of the penetrator's landing site (Liu, et al., 2019). Given that similar wide angle descent/ascent images can also be collected from other spacecraft such as soft landers, aerial vehicles and balloon probes and used for stereo imaging, there is ample motivation to study the method's capabilities.

5.6.1 Method

Measurements of the topography of solar system bodies are commonly stored, presented and used in the form of a digital terrain model (DTM). These two-dimensional image representations of surfaces are intuitive to inspect by eye, can be readily interpreted by 3D analysis and visualisation software, and permit rigorous quantitative analysis of the represented surface. DTMs are a common end product of topographic measurements, and a crucial intermediate step between acquiring a stereo image pair and producing a DTM, is the triangulation of points on the surface, resulting in a point cloud. Whilst a point cloud contains the discrete 3D topography measurements derived from a stereo pair, it differs from a DTM in that it does not give a model of a single connected surface. The accuracy of a DTM is largely a function of the accuracy of the point cloud from which it is constructed, but is also dependent on the method of construction. The work here will focus on the former of these, by investigating the accuracy with which surface elevation can be measured from stereo descent images.

The accuracy of surface stereophotogrammetry is a function of many factors (including the camera properties, stereo geometry (Morgan, et al., 2010) and lighting conditions (Kirk, et al., 2016)), and cannot be fully assessed or estimated without knowledge of the surface it represents. Traditionally, the quality of a point cloud or DTM will therefore be investigated by comparing it to check points: measurements of the surface obtained by another means (Li, 1988). For stereo images with high spatial resolution and large surface coverage, achieving a sufficiently large and distributed set of check points can be a barrier to fully assessing measurement accuracy (Li, 1991).

When the objective is not to map a specific surface, but to investigate an instrument's ability to measure topography (as is the case here), simulated images of a pre-existing

surface model can be used. By utilising a pre-existing DTM as a ground truth, images of a surface can be simulated and subsequently used to measure the ground truth topography. The measured topography can be directly compared with the ground truth in order to assess its accuracy and quality.

A benefit of this approach is that the ground truth need not be a highly accurate representation of its real surface, as the objective is to reproduce the ground truth, not the surface it represents. An important property of the ground truth however is that it contains features that are representative of the surfaces on which the topographic measurement method is intended to be used (Kirk, et al., 2016).

A challenge of simulating imagery from digital surface models is that they exhibit low contrast at scales approaching their sampling length (e.g. on the scale of a single DTM pixel), whilst real surfaces contain variation in their appearance at all scales. To mitigate this, the surface model should have a higher spatial density than can be resolved by the simulated imaging system (see also chapter 7, section 7.3). When simulating high-resolution imagers, achieving this can be computationally costly.

5.6.1.1 Ground Truth Data

Satisfying the above criteria for a ground truth data set for the penetrator descent camera is not necessarily straight forward. The wide FOR of the camera, coupled with its high altitude images, means that a ground truth with extensive coverage is required. Further to this, the camera's wide FOV and significant change in altitude over the course of a descent means that its GSD varies by orders of magnitude (see section 5.3.1), and a ground truth with a high spatial resolution is required. Datasets which meet this combination of requirements are not common: generally, large coverage DTMs have a coarse resolution, and high resolution DTMs have a small extent. This investigation benefitted greatly therefore from the availability of a DTM mosaic and accompanying ORI mosaic covering the entirety of Mars' Gale crater, the landing site of the Mars Science Laboratory (Wray, 2013).

The dataset (dataset 5.1 in Table 5.1) comprises both CTX (Malin, et al., 2007) and HRSC (Jaumann, et al., 2007) imagery, and its production is described in Persaud et al. (2019). It was selected due to the DTM's large coverage (spanning a region of approximately 170 by 200 km), its high spatial resolution of 18 m per pixel, and the variety of terrain types, surface features and elevation variation present in the data. The accompanying greyscale ORI covers the same extent as the DTM at 6 m per pixel horizontal spacing and provides additional visible surface detail. The DTM and ORI which together constitute the ground truth data set are shown in Figure 5.36.

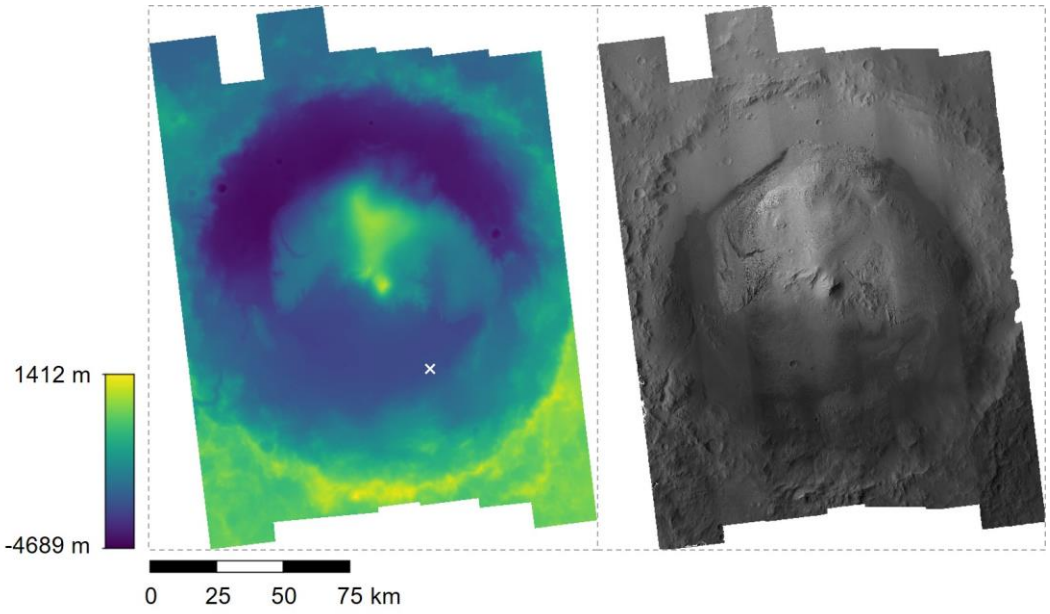

Figure 5.36: The ground truth data, from Persaud et al. (2019), used for simulation of descent images (dataset 5.1). (left) DTM of Gale Crater with 18 m spatial resolution and elevation extent of -4689 m (dark purple) to 1412 m (light yellow). (right) 6 m per pixel ORI of Gale Crater. Both images span 170 km horizontally and 214 km vertically (northward). The white cross indicates the location of this study's simulated descent.

In spite of its large extent, the ground truth dataset does not extend all the way to the visible horizon when simulating high altitude images. DTMs and ORIs covering the entirety of Mars' surface are available (e.g. Fergason et al. (2018)), and could have been used to supplement dataset 5.1 and extend the area of simulated surface. Indeed this was done for many of the simulated images used in the study of the penetrator camera (see Figure 5.21 for example). However, for the investigation of stereophotogrammetry, the decision was taken not to do this, in order to avoid the merging of datasets of different grid spacing and vertical accuracy, on which stereo matching may have performed differently.

5.6.1.2 Image Simulation

Images were simulated as being captured under ideal conditions (as described in section 5.3), i.e. a vertical descent with constant vertical spin axis orientation and constant spin rate. Deviations from this in real data could, to some extent, be characterised and corrected for with data from complementary sensors (e.g. magnetometer, star-tracker, Sun sensor) and potentially the images themselves (e.g. Di et al. (2013)). Images were simulated as being captured through the central column of the camera's detector, and azimuthally rectified as described in section 5.5.2. The simulated descent has a nadir

location of 5.886°S, 138.170°E, purposely offset from the crater centre to avoid alignment with the radial symmetry in the topography.

The camera was simulated to have a coarse angular resolution of 0.18° x 0.18° per pixel (corresponding to an active detector region of 500 pixels spanning the full vertical FOV) in order to keep its pixels' footprints larger than the ground truth DTM's horizontal grid spacing. Results are normalised relative to camera resolution, such that they can readily be applied to a range of imaging systems of different angular resolution.

The focus of this study is the dependency of stereo measurement accuracy on baseline and viewing geometry, whilst the impacts of surface illumination are beyond its scope. For this reason, surface brightnesses, reflections and shadows were no photometrically simulated, but instead mapped directly from the high resolution ground truth ORI, which contains calibrated albedo and illumination variation of the Martian surface (Edmundson, et al., 2012). During a temporally short descent/ascent, surface lighting conditions will remain relatively unchanged, and the most significant photometric affect would likely be the change in phase angle across a stereo pair.

Because of the limited physical extent of the ground truth, images were restricted to being simulated at or below altitudes of 13 km (above which it was judged that too much of the camera's FOV was viewing the region of no data beyond the DTM's boundaries). To minimise the impact of the camera's spatial resolution exceeding the ground truth DTM grid spacing, no images were simulated below 5 km altitude. Stereo pairs were simulated at altitudes throughout this range, and 33 different stereo pairs were utilised in order to investigate the effect of both imaging altitude and stereo baseline on the achievable elevation measurement accuracies. For any stereo pair, the geometry is described by its imaging height (the altitude of the lower image) and baseline (the difference in altitude between the higher and lower image). The height and baseline-to-height ratios of the 33 simulated stereo pairs used in this study are plotted in Figure 5.37.

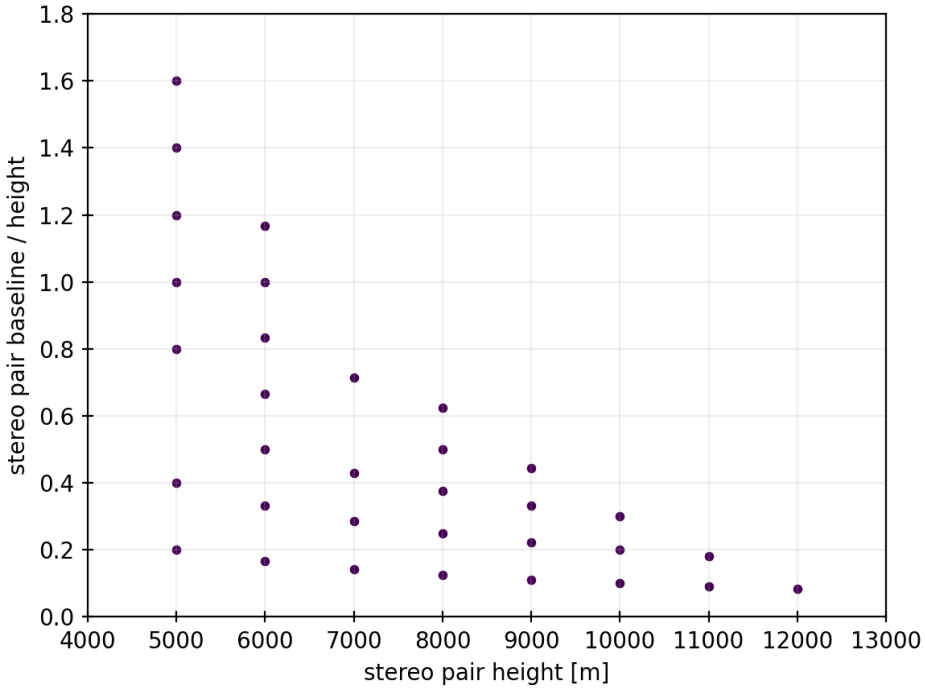

Figure 5.37: The heights and baselines of the 33 simulated stereo pairs used in this study.

5.6.1.3 Topography Measurement

For a given stereo pair, surface points were triangulated by applying stereo matching to the pair's two images in order to produce a single surface point cloud. Stereo matching was performed with the StereoBinarySGBM class of the OpenCV library (Bradski, 2000), which implements a modified Hirschmuller semi-global matching algorithm (Hirschmuller, 2008) to find correspondence between two rectified images. This algorithm is not tailored to planetary data, but was used for the simplicity of its implementation in this study, and its quick execution time. Choice of stereo matching algorithm does affect the quality of topographic measurement, as does its optimisation for the appearance of planetary surfaces, but the study of this is beyond the scope of this work (Heipke, et al., 2007). Semi-global matching does not enforce a smooth surface, and accommodates occlusions, both of which are beneficial characteristics for the viewing geometry and surface topography used in this study. However, the matching algorithm does not account for significant changes in the scale and shape of image features, both of which occur in the penetrator descent camera data due to its changing viewing perspective. Therefore, an error model is used (section 5.6.1.5) to assist with

identifying the effect of these scale and shape changes on the accuracy of stereo matching.

The matching algorithm attempts to find full correspondence, i.e. for every pixel in one image, it aims to find the pixel in the other image which views the same point in the scene. For each pair of corresponding pixels, the algorithm returns its disparity, D, (the difference in the pixels' positions in the two images). The output is a disparity map, which displays the disparity measured for each pixel as a greyscale colour (Figure 5.38 (b)).

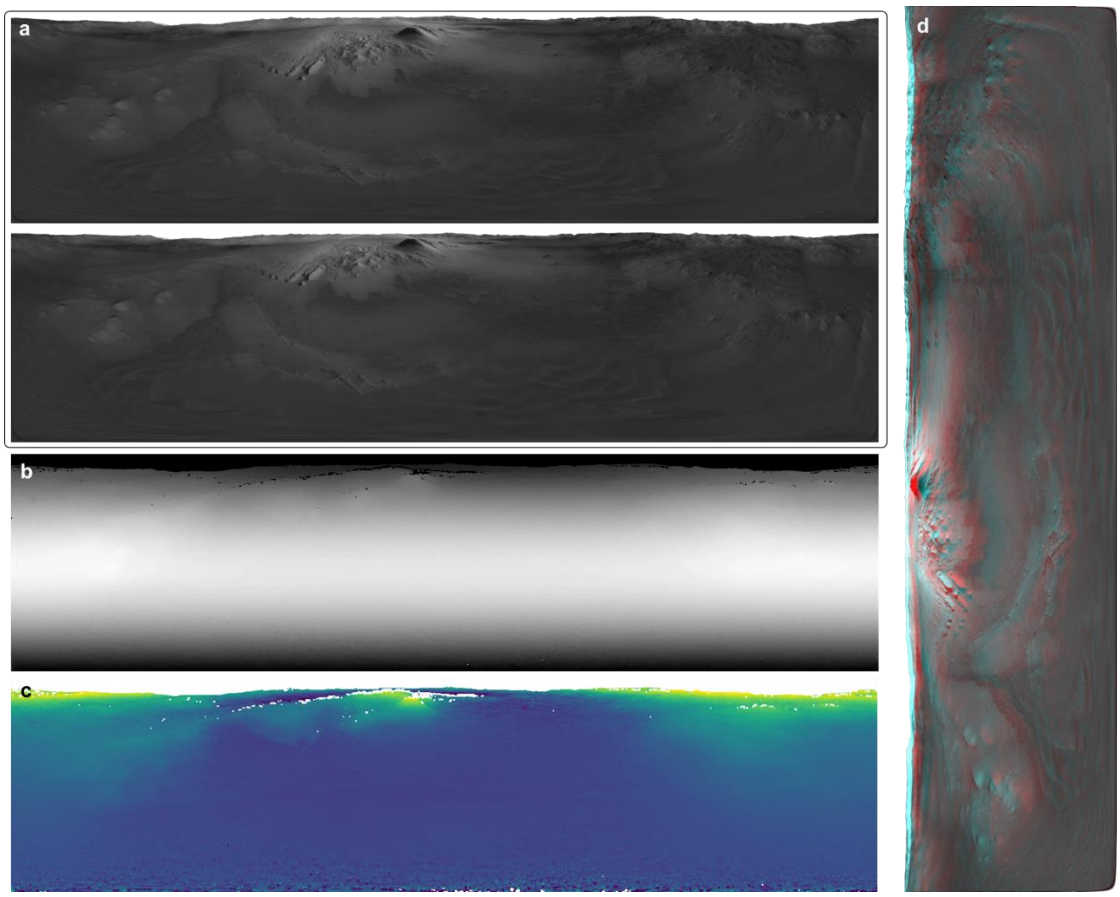

Figure 5.38: (a) Stereo pair of simulated images with imaging altitude 4 km and baseline 2 km captured during a descent to Gale crater (dataset 5.1). (b) The stereo pair's disparity map. Black (white) indicates low (high) disparity. (c) Surface elevation map derived from the disparity map as described in the main text. Dark (light) colours indicate low (high) terrain. (d) Red-blue anaglyph produced from the stereo pair. The anaglyph must be viewed in its portrait orientation, in order to align its vertical baseline with the eyes' horizontal baseline.

Recalling the imaging geometry of the penetrator camera (described in section 5.3.1), the off-nadir angle of the line of sight of a pixel in the central detector column is given by

$$a = \frac{\pi}{2}\left(1 - \frac{n_{row}}{N-1}\right) \tag{5.16}$$

where N is the total number of pixels spanning the vertical FOV and $n_{row} \in \{0 \dots N-1\}$ is the number of the pixel's row, with a value of zero corresponding to viewing horizontally (i.e. the top of the image) and a value of $N-1$ viewing the nadir.

Taking a pixel in the lower altitude image of a stereo pair (image 1), the off-nadir angle of its line of sight is given by equation 5.16, and the off-nadir angle of the line of sight of the corresponding pixel in the higher altitude image (image 2) is given by

$$a_2 = a_1 - \frac{\pi\, D}{2N} \tag{5.17}$$

The baseline length, b of a stereo image pair and the pixel LOS angles a_1 and a_2 constrain the location of a viewed point relative to the imaging locations, as in Figure 5.39.

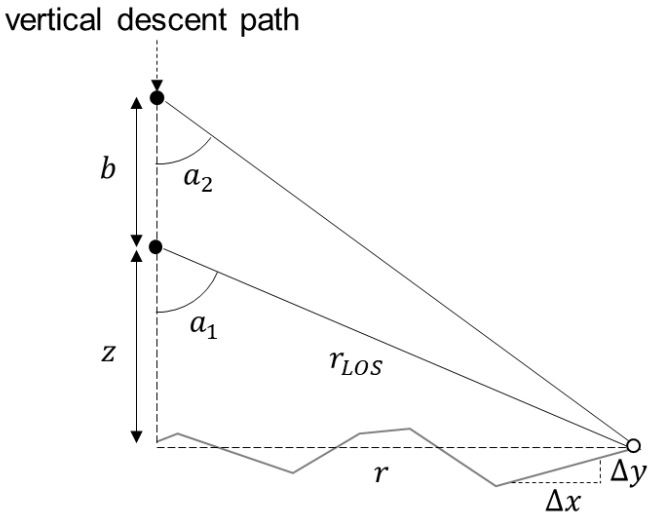

Figure 5.39: Stereo viewing geometry by which the physical positions of points (empty circle) on the surface (grey line) are triangulated from two imaging locations (solid circles).

The point's vertical displacement from the lower imaging location (whose altitude defines the stereo pair's altitude), is given by

$$z = r_{LOS} \cos a_1 \tag{5.18}$$

whilst its radius, measured perpendicular to the descent axis, is given by

$$r = r_{LOS} \sin a_1 \tag{5.19}$$

where the line of sight distance, r_{LOS} is given by

$$r_{LOS} = b\frac{\sin a_2}{\sin(a_1 - a_2)} \tag{5.20}$$

From z, r and viewing azimuth, the location of each pixel's viewed point is constrained, and readily transformed back into the planetocentric coordinates of the ground truth data. Performing this for every matched pixel results in a point cloud representing the measured surface, which can be directly compared with the ground truth DTM. Some examples of these point clouds are shown in Figure 5.46.

For a real penetrator descent, the camera's height will decrease over the course of capturing a single image, and image altitude will be a function of azimuth. If acceleration is occurring, as would happen for a penetrator in free-fall, the baseline between two images will not be constant and will also be a function of azimuth. Given that the images in a stereo pair are rectified by azimuth, this effect can be accounted for in equations 5.18-5.20 by making both z and b functions of azimuth, and treating each column of the stereo images independently. Given that this does not require functionally different analysis, the effect of vertical motion during image capture was not simulated in this study.

5.6.1.4 Point Cloud Elevation Accuracy

For the purpose of quantitatively assessing the quality of the descent camera's topography measurement, elevation accuracy was estimated by subtracting the derived elevation values from the ground truth's elevation values at corresponding longitude and latitude, to produce difference clouds (an example is shown in Figure 5.40). Because of the simulated camera's significant change in viewing geometry over its wide FOV, it is to be expected that measurement accuracy will vary across the FOV. Given that the camera's imaging geometry has radial symmetry, and its GSD is a function of radius (see section 5.3.2), the elevation accuracy was measured as a function of radius from nadir. This was achieved by segmenting the difference cloud into radial bins (illustrated in Figure 5.40). The elevation root mean square error (RMSE) was calculated for each annulus individually.

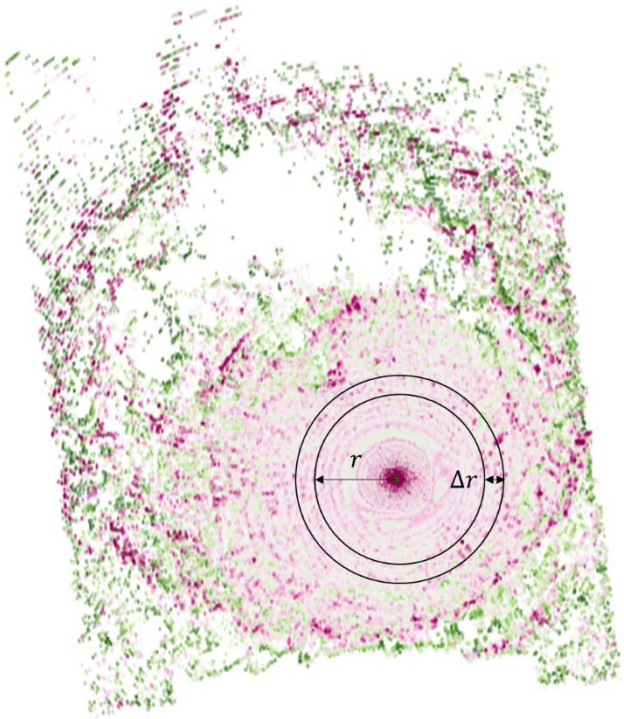

Figure 5.40: Map of difference (green indicates positive, pink negative) between measured elevations and ground truth DTM. An annular bin, as was used for elevation error assessment, is illustrated.

To generalise the results, rather than presenting values in units tied to the physical dimensions of the ground truth data set (i.e. metres), they will be given relative to the imaging geometry. Baselines will be normalised to imaging height (the height of the stereo pair's lower image), elevation errors will be normalised to the camera's GSD, and viewing geometry will be expressed by the ratio of imaging height to ground radius.

5.6.1.5 Modelling Stereo Matching Accuracy

There are many different software packages and algorithms (e.g. SOCET SET (Walker, 2007); Ames Stereo Pipeline (Moratto, et al., 2010)) available and in use for creating DTMs from stereo images. Because of the unusual imaging geometry of the penetrator descent camera images, and because stereophotogrammetry was not the sole focus of this thesis, accessing and utilising such software was beyond the scope of this study. Instead, perhaps unusually, bespoke but relatively simple software, which leverages a general purpose disparity algorithm, was used for the triangulation of points' positions. Given that a large contributing factor to the accuracy of topography measurements is the performance of the stereo matching algorithm, an error model is here presented which assists with contextualising the results of this study, and making comparison with other stereo matching software.

As discussed in section 5.6.1.3, Figure 5.39 and equations 5.18-5.20 describe the stereo viewing geometry by which surface points are triangulated.

Performing error analysis yields

$$\sigma(z)^2 = \sigma(r_{LOS})^2 \left(\frac{\partial z}{\partial r_{LOS}}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial z}{\partial a_1}\right)^2 \tag{5.21}$$

$$\sigma(r)^2 = \sigma(r_{LOS})^2 \left(\frac{\partial r}{\partial r_{LOS}}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial \mathrm{r}}{\partial a_1}\right)^2 \tag{5.22}$$

where

$$\frac{\partial z}{\partial r_{LOS}} = \cos a_1 \tag{5.23}$$

$$\frac{\partial z}{\partial a_1} = -r_{LOS} \sin a_1 \tag{5.24}$$

$$\frac{\partial r}{\partial r_{LOS}} = \sin a_1 \tag{5.25}$$

$$\frac{\partial r}{\partial a_1} = r_{LOS} \cos a_1 \tag{5.26}$$

and

$$\sigma(r_{LOS})^2 = \sigma(b)^2 \left(\frac{\partial r_{LOS}}{\partial b}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial r_{LOS}}{\partial a_1}\right)^2 + \sigma(a_2)^2 \left(\frac{\partial r_{LOS}}{\partial a_2}\right)^2 \tag{5.27}$$

where

$$\frac{\partial r_{LOS}}{\partial b} = \frac{\sin a_2}{\sin(a_1 - a_2)} \tag{5.28}$$

$$\frac{\partial r_{LOS}}{\partial a_1} = -b \frac{\sin a_2 \cos(a_1 - a_2)}{sin^2(a_1 - a_2)} \tag{5.29}$$

$$\frac{\partial r_{LOS}}{\partial a_2} = b \frac{\sin a_1}{sin^2(a_1 - a_2)} \tag{5.30}$$

This study assumes exact knowledge of the stereo geometry, meaning $\sigma(b) = 0$. With the camera's equidistant imaging geometry, $\sigma(a_1)$ and $\sigma(a_2)$ are taken to be constant over the entire FOV. The above equations could equally be applied to different imaging geometries by including an appropriate dependence on a in each of these uncertainties.

In deriving the elevation e of a point on the surface, error is introduced by the vertical uncertainty in z, and the horizontal uncertainty in r, over which the actual elevation of the surface will change based on the slope of the terrain $s = \Delta y/\Delta x$ (see Figure 5.39). The variance of e is therefore given by

$$\sigma(e)^2 = \sigma(z)^2 + {s_{RMS}}^2\, \sigma(x)^2 \tag{5.31}$$

where s_{RMS} is the ground truth's RMS slope. Figure 5.41 plots the RMS slope of the ground truth DTM, measured over a range of horizontal baselines. Also plotted is an histogram of the relative frequency of radial GSDs with which the surface is imaged in the set of 33 stereo pairs used in this study. The majority of pixels' radial GSDs are below 120 m, and therefore s_{RMS} is taken to be 0.19 at all times.

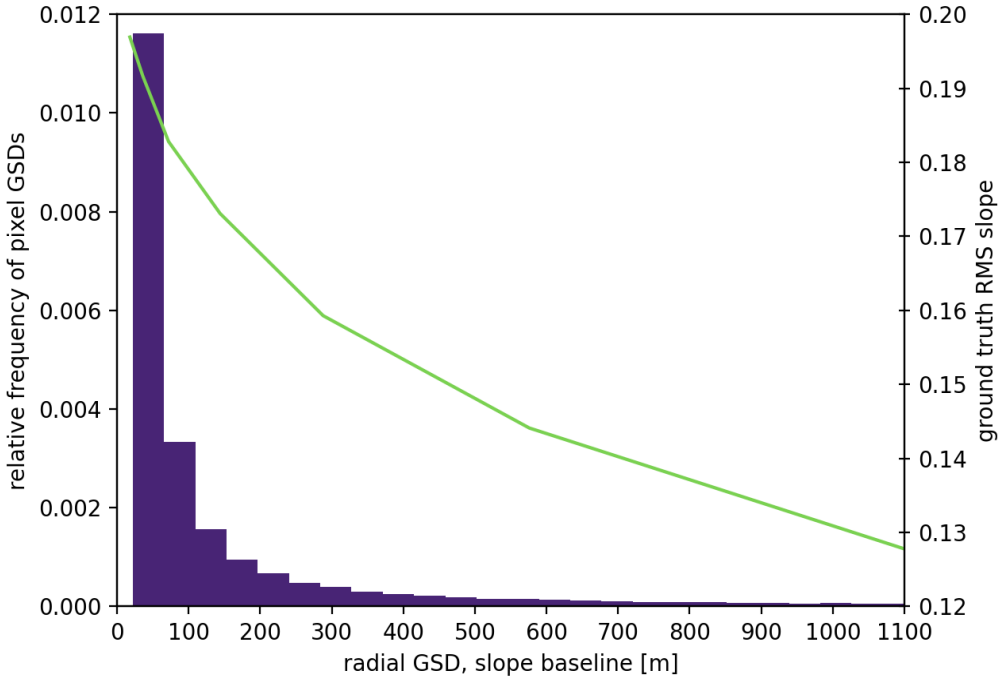

Figure 5.41: Histogram (purple) of relative frequency of radial GSDs of the pixels in the 33 simulated stereo pairs, plotted alongside the RMS slope (green) of the ground truth DTM as a function of baseline length over which slope is measured.

The following three plots illustrate the nature of elevation error, as modelled by equations 5.21-5.30. Figure 5.42 shows the behaviour of elevation error as a function of ground radius, for three different baseline geometries, when imaging from a fixed height. The curves are calculated based on the measured slope of 0.19, and an error of 0.2x the camera's angular resolution in each of a_1 and a_2. Elevation error is normalised to the GSD of the lower camera's nadir-viewing pixel.

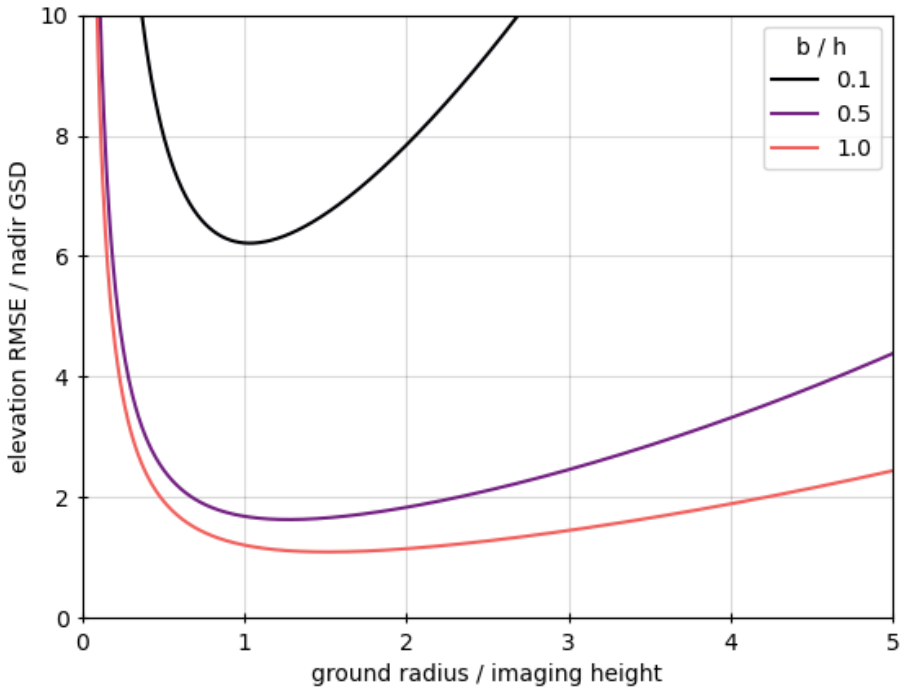

Figure 5.42: Modelled elevation RMSEs as a function of ground radius, when imaging from a fixed height. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution. Elevation error is normalised to the GSD of the lower camera's nadir pixel.

Figure 5.42 demonstrates the significant effect of the camera's wide range of viewing perspectives. Surface near the nadir is close to the descent axis, so exhibits minimal parallax between the two images of a stereo pair, leading to high error in elevation measurement. Large ground radii also exhibit growing elevation error, and a minimum in measurement error occurs around a ground-radius-to-height ratio of 1-3, depending on baseline. Larger baselines achieve greater elevation measurement accuracy, and exhibit a larger range of ground radii over which elevation measurement is optimal.

Figure 5.43 shows a complementary plot of elevation RMSE versus imaging altitude, when observing a fixed point on the surface of given ground radius. Slope and angular

error values match those of Figure 5.42. Elevation error is normalised to the GSD of the camera's nadir viewing pixel were the camera at an altitude equal to the ground radius of the observed point.

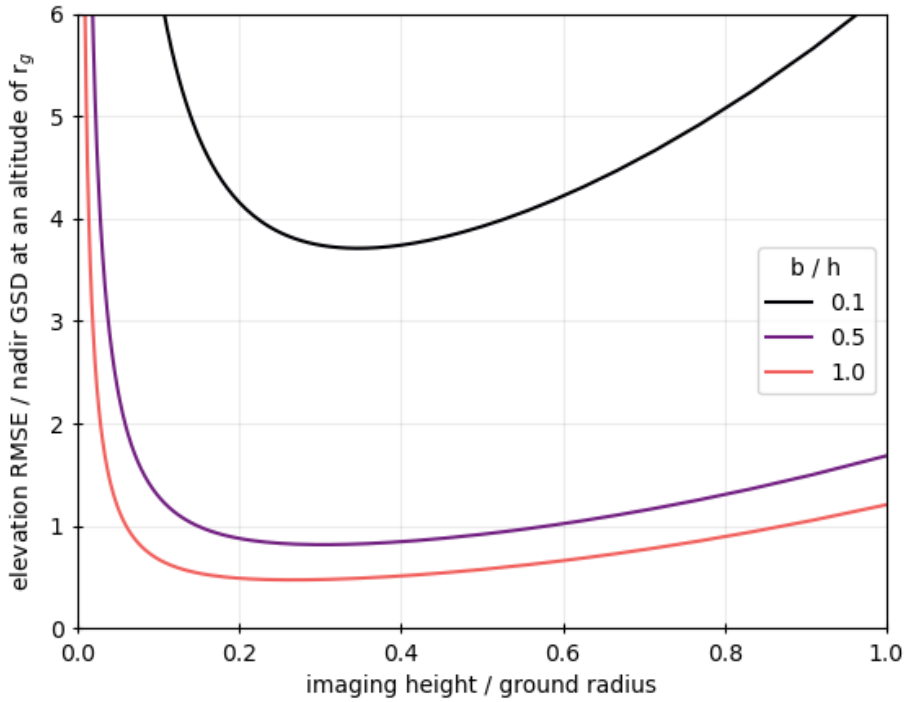

Figure 5.43: Modelled elevation RMSEs as a function of imaging altitude, when viewing a point on the surface with fixed ground radius. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution.

Just as there is an optimum ground radius to measure from a given height, Figure 5.43 shows that there is an optimum height from which to measure a given ground radius. For the plotted baselines, optimum height-to-ground-radius ratio occurs between 0.25 and 0.3. Note that the imposition of an optimum imaging altitude by the stereo viewing geometry is comparable to that dictated by the radial GSD, covered in section 5.3.2.4, but arises by a different mechanism and occurs at a different altitude-to-ground-radius ratio (~0.3 versus 1). It should also be noted that the minima in elevation measurements of Figure 5.43 only occur for non-zero values of s_{RMS}, otherwise the elevation RMSE decreases monotonically with decreasing imaging height (comparable to azimuthal GSD, see Figure 5.12).

Finally, Figure 5.44 plots the elevation RMSE as a function of stereo baseline length, when imaging a fixed point on the surface from a fixed altitude. Four different stereo

geometries, labelled by their ratios of the ground radius to imaging altitude, are plotted. Modelled slope and angular errors remain unchanged from Figure 5.42.

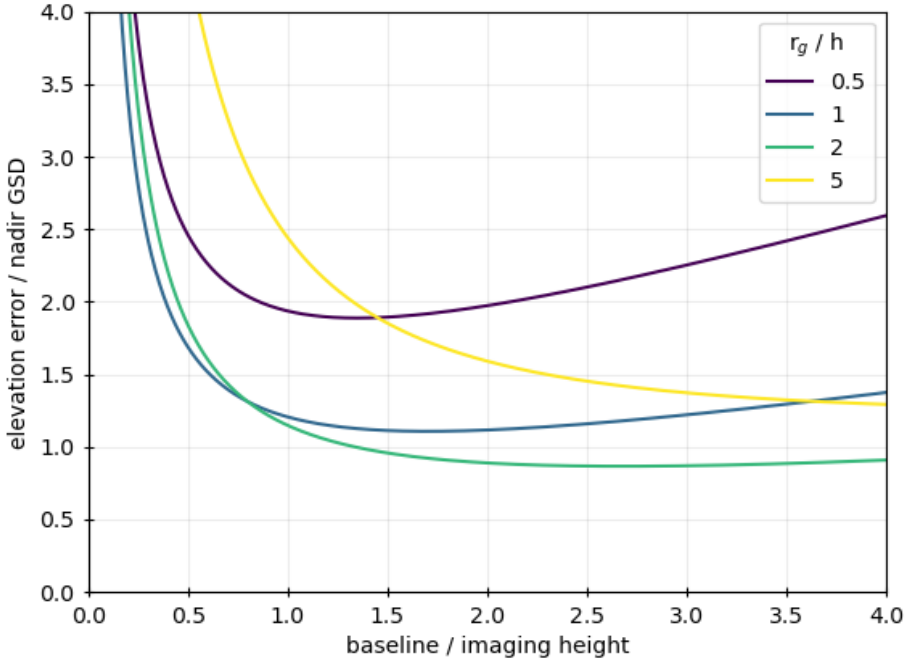

Figure 5.44: Modelled elevation RMSE as a function of stereo baseline. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution.

Modelled curves of the types plotted in Figure 5.42 and Figure 5.43 are plotted alongside the measured results presented in section 5.6.2 in order to assist with their interpretation. First however, the model is here used to estimate the uncertainties $\sigma(a_1)$ and $\sigma(a_2)$ associated with the stereo matching algorithm when applying it to simulated images of the ground truth. Images of a flatted ground truth (achieved by mapping the entire surface to a single elevation) were simulated, and its topography then stereoscopically measured by the method described in section 5.6.1.3 (flattening the ground truth removes the uncertainty in the slope's contribution). The measured error curves were then used to estimate $\sigma(a_1)$ and $\sigma(a_2)$ (assuming the two to be equal), by plotting them alongside model error curves, and varying the value of $\sigma(a_1)$ and $\sigma(a_2)$ to achieve best agreement (judged by eye) with the measured curves. Through this approach, it was found that model and measured curves best agreed when $\sigma(a_1)$ and $\sigma(a_2)$ were modelled as having a value of 0.11x the pixel angular resolution. Three example comparisons of measured and modelled curves for a flat surface are plotted in Figure 5.45.

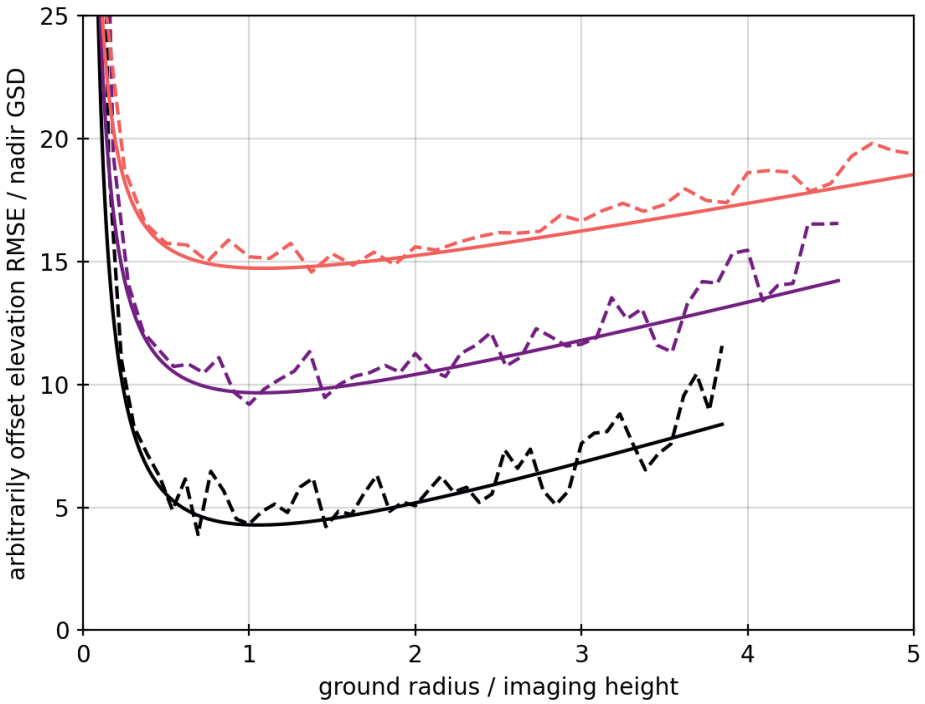

Figure 5.45: Example modelled elevation RMSEs (solid) and measured elevation RMSEs (dashed) for a flat surface, for three different stereo pairs of simulated images (offset vertically for clarity). Modelled angular error is 0.11x the pixel angular resolution for all three curves.

5.6.2 Results

Figure 5.46 shows a selection of elevation point clouds generated from simulated stereo pairs in this study. The descent nadir location is indicated by a yellow dot in the bottom right quadrant of each. Point cloud spatial density decreases with distance from nadir, due to the camera's decreasing spatial resolution, but large areas of no data (white) also exist (particularly north west of nadir) in regions where high elevation surface obscures more distant, lower elevation surface from the camera's view in at least one of a stereo pair's images.

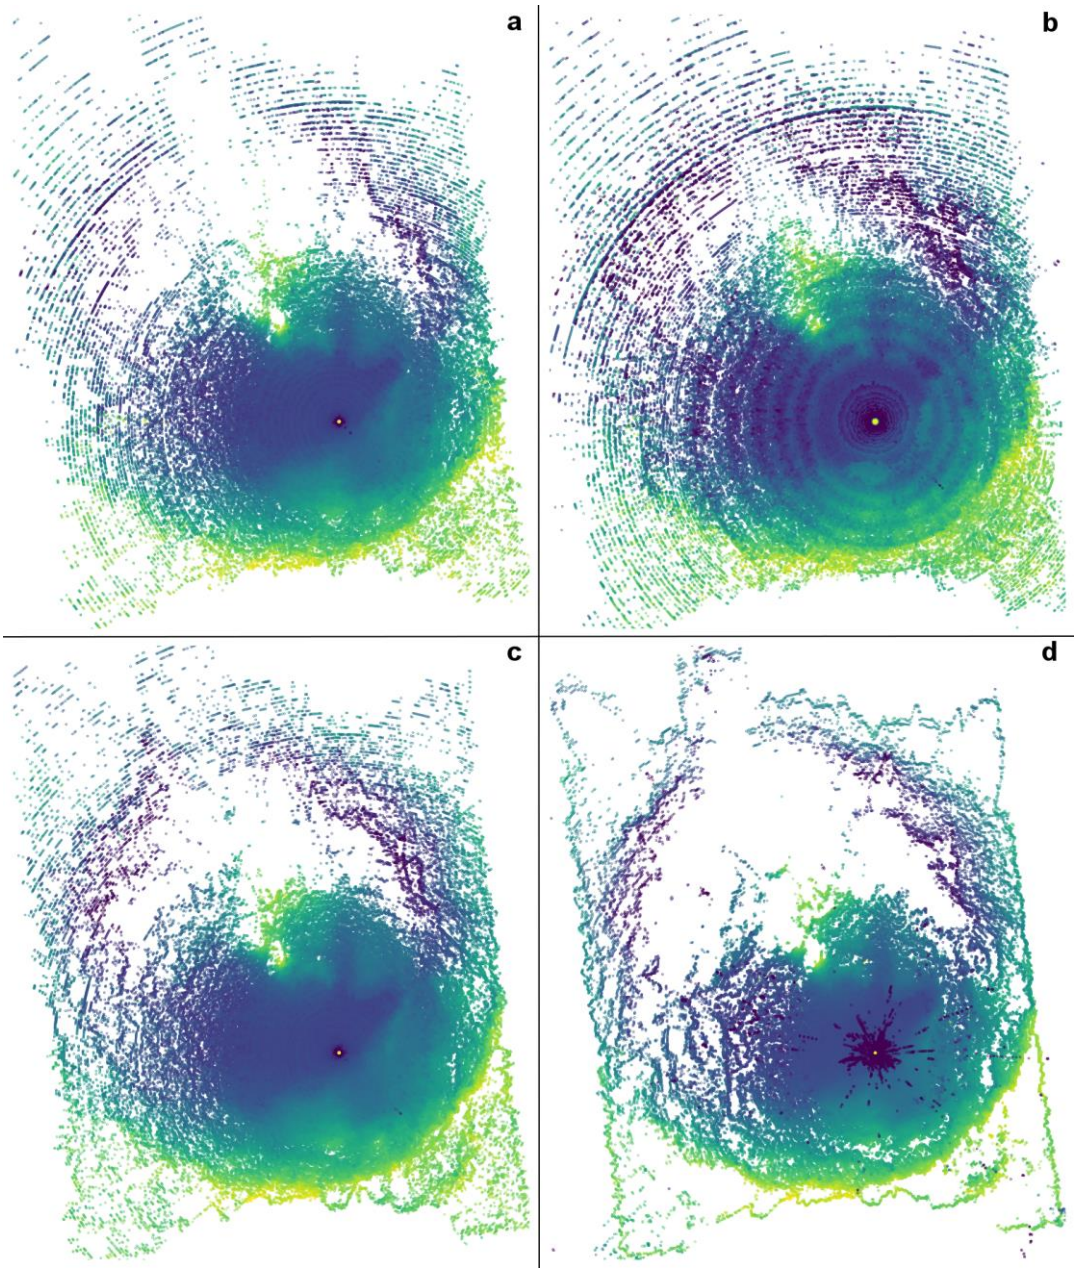

Figure 5.46: Four example elevation point clouds measured in this study. They were obtained from images with baseline/altitude of (in the ground truth's physical units) 1/5, 1/12, 2/7, 6/5 km for a, b, c and d respectively. All are displayed with the same elevation scale with dark (light) shading indicating low (high) terrain. White indicates no data.

There are clear visible differences between the four example point clouds of Figure 5.46. The high altitude images of (b) reduce the impact of occlusions (blind spots), leading to the smallest region of no data of all four products. However, its small baseline leads to large elevation errors, noticeably manifesting as concentric rings. These result, at least in part, from the chosen stereo matching algorithm, and different algorithms may mitigate

this. With its large baseline, (d) does not exhibit these rings, but the region of error around nadir is significant, and the low altitude results in widespread occlusions.

5.6.2.1 Error Versus Ground Radius

Figure 5.47 displays the results of measuring the elevation RMS error in 33 different topography point clouds, each derived from a different stereo pair of simulated images covering a range of baseline-to-imaging-height ratios of 0.07-1.14. Solid coloured lines show the measured error, whilst the grey dashed lines show the modelled error (as in section 5.6.1.5) for each stereo geometry, assuming a surface slope RMS of 0.19 and angular error of 0.020° (0.11x the angular resolution of the pixels). For clarity, the data are grouped into six separate plots, based on their baseline-to-imaging-height ratios.

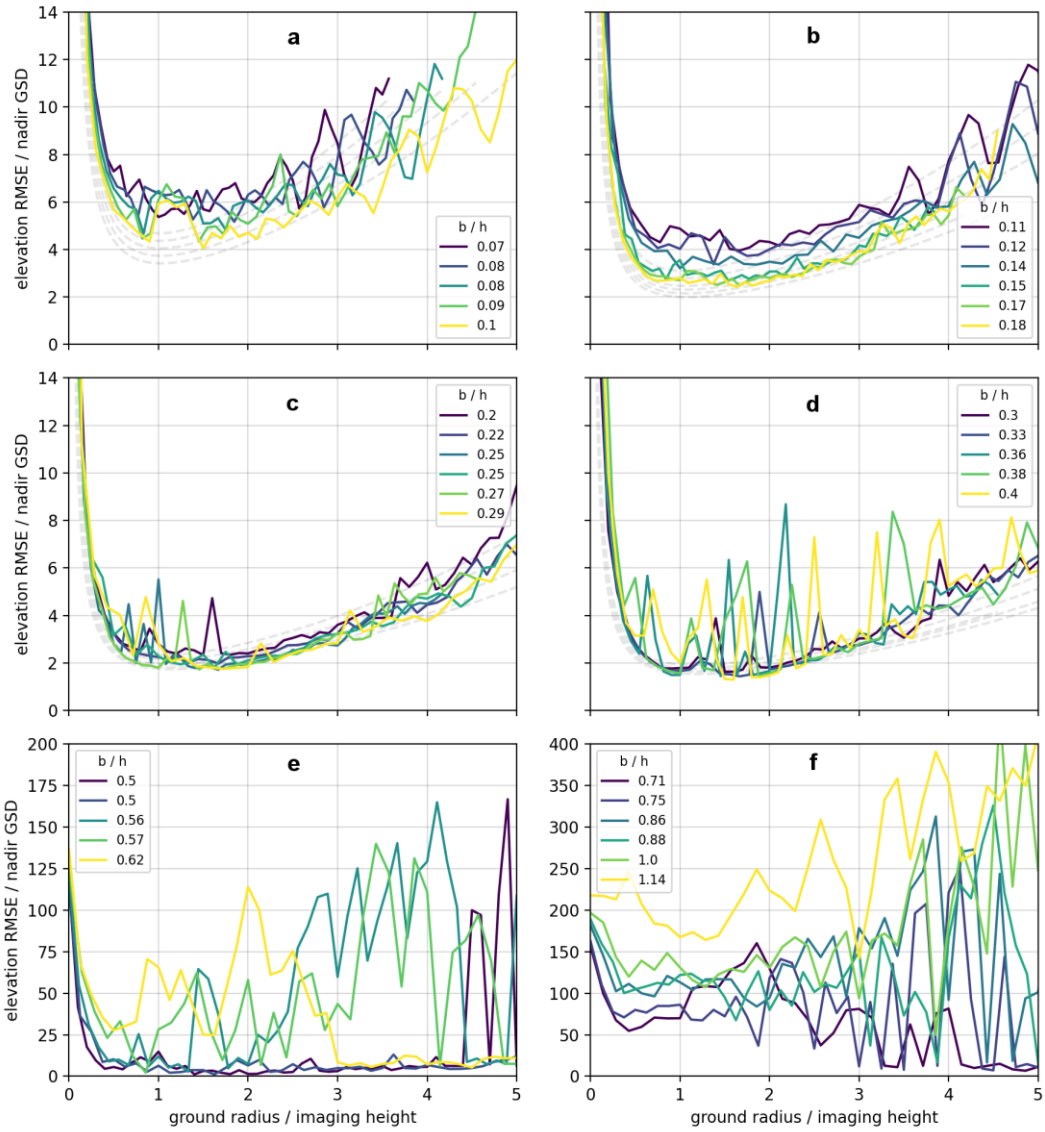

Figure 5.47: Measured elevation RMS error as a function of radius from nadir for 33 stereo pairs (solid coloured lines), segmented into six separate plots based on their different b/h values. Dashed grey lines are the modelled error for each stereo geometry, assuming a surface RMS slope of 0.19 and angular error of 0.11, 0.11, 0.14 and 0.16x pixel angular resolution for plots a, b, c and d respectively.

The overall shapes of plots (a), (b) and (c) of Figure 5.47 closely match their model error curves, whilst the measured curves of (d) deflect upwards from the model at a ground-radius-to-height ratio of ~3, in addition to also exhibiting many significant sharp, localised deviations at a variety of radii. The upward deflection of the curves in plot (d) suggests a decrease in stereo matching accuracy at the larger ground radii, potentially due to the more oblique viewing.

The increasing baseline (0.07–0.29) through plots (a) to (c) is accompanied by a reducing elevation RMSE. Elevation error is yet smaller for the larger baselines (0.3–0.4) of plot (d), but sharp spikes in error are also more frequent and greater in amplitude. In (e) and (f) elevation errors no longer resemble the model behaviour, exhibit an increase with increasing baseline, and are significantly larger than the errors of plots (a) to (d). This is likely due to the significant change in the surface's appearance (scale and orientation of surface features) between two images of a stereo pair, as a result of the baseline size, leading to a significant reduction in stereo matching efficacy. The sharp spikes in the plotted curves, most striking in plot (d), are likely due to isolated regions of surface whose appearance strongly negatively impacts the efficacy of stereo matching. This may be due to the fundamental appearance of the surface in an image (e.g. exhibiting low contrast), or the change of its appearance between the two images of a stereo pair (e.g. a significant change in shape and/or size). Similar error features were also seen in some of the measured curves of flat-surface images, indicating that surface topography is not solely responsible. For the curves of plots (a) to (d), minimum error occurs around a ground-radius-to-height ratio of 1.5, and error remains less than double its minimum value (disregarding spikes) within the ground-radius-to-height ratio range of 0.5–3.

5.6.2.2 Error Versus Imaging Height

Figure 5.48 plots the measured elevation RMS errors of all data within the topographic point clouds derived from the same 33 stereo pairs as presented in Figure 5.47. Error is plotted as a function of imaging height, when observing a fixed ground radius. Height is normalised to this ground radius, whilst elevation RMSE is normalised to the camera’s nadir GSD were it at an altitude equal to the ground radius. The data are colour coded by their baseline-to-imaging-height ratio (in groups corresponding to the six plots of Figure 5.47).

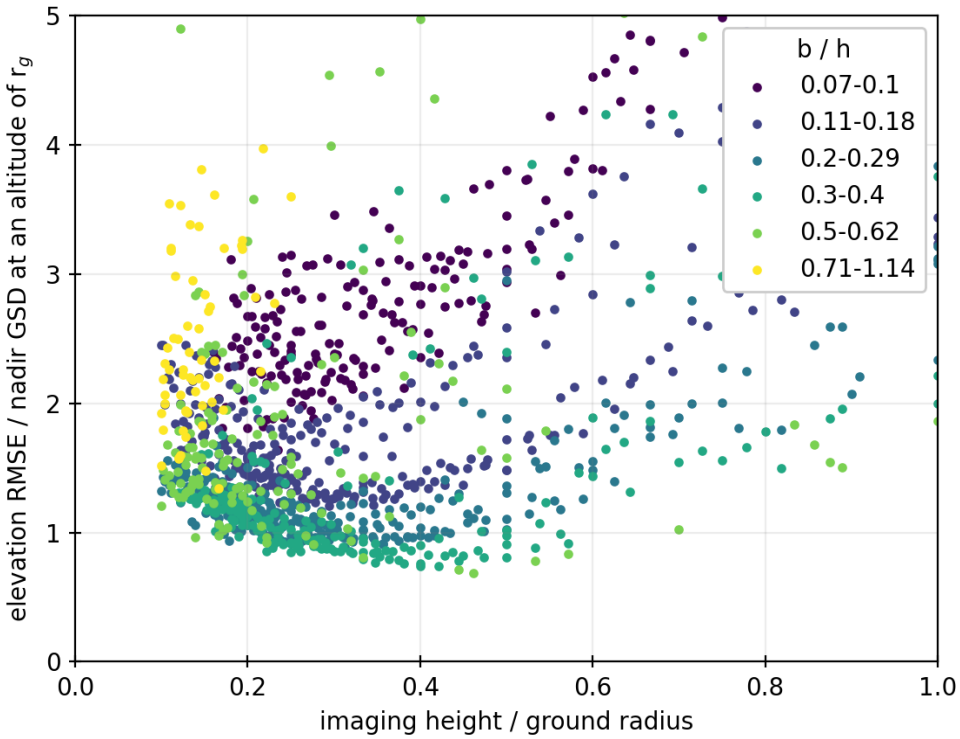

Figure 5.48: Measured elevation RMS error as a function of imaging height for a fixed ground radius, for 33 stereo pairs (same as are plotted in Figure 5.47).

There is a large spread in the full dataset plotted in Figure 5.48, and the shape of the modelled curves presented in section 5.6.1.5 is not immediately recognisable within. Note however that the groups with lowest elevation RMSE (groups with b/h ranges 0.11–0.18, 0.2–0.29, 0.3–0.4) do exhibit a minimum error between height-to-ground-radius ratios of 0.2 and 0.5, with error climbing steeply either side of this range, similarly to the modelled curves. The segmentation of the six groups is generally clear, except for at the very lowest imaging heights, and indicates lowest elevation RMSEs occurring for a baseline-to-imaging-height ratio between 0.2 and 0.4.

5.6.2.3 Surface Visual Assessment

To give a complete impression of the nature of the descent camera's topography measurement, this section provides some very brief qualitative discussion on the derived point clouds' representations of the surface, compared to that of the ground truth. For this, an example point cloud has been visualised from a variety of views, where each point in the cloud is given the colour of its corresponding pixel in the lower of the two stereo images from which the point cloud is derived. Three such visualisations are shown in Figure 5.49.

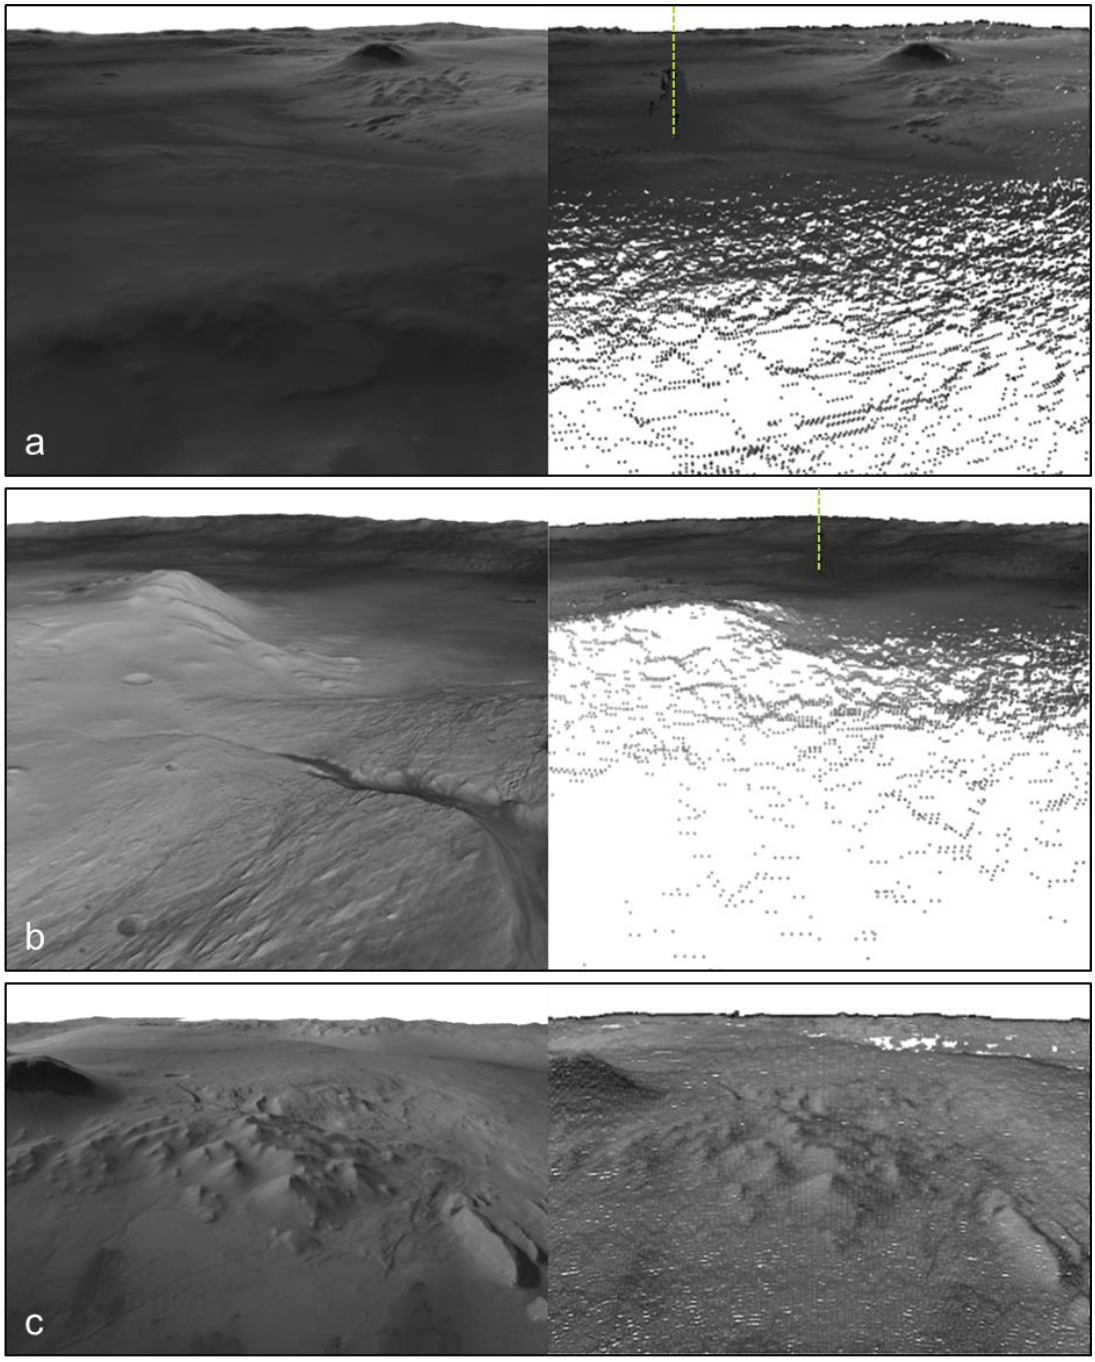

Figure 5.49: Three views of the Gale crater ground truth (left) and the same views of an example point cloud derived from a simulated descent stereo pair (right). (a) A view of the crater and landing site from its south east rim. (b) The peak of Mount Sharp, and distant landing site, viewed from the north west. (c) A view from the penetrator's descent, looking approximately northward. Green dashed lines in (a) and (b) indicate the penetrator's descent and landing site. White indicates no data.

The decreasing spatial resolution of the data with increasing ground radius is clear to see in Figure 5.49 (a) and (b). Also clear is the connection between elevated terrain and

occlusions, leading to regions of no data. Panel (b) shows a large area of no data on the far side of the peak of Mount Sharp from the descent location. A similar blind spot exists in panel (c) behind a ridge of the crater's central mound. Panels (a) and (b) show a significant artefact in the region immediately surrounding nadir (green dashed line), where elevation error grows rapidly. Both the topography and appearance of this region of surface is lost. Figure 5.49 (particularly panels (a) and (c)) does however demonstrate that significant detail on both the topography and appearance of the surface is captured within the point clouds.

5.6.3 Descent Stereophotogrammetry: Key Conclusions

The geometry of stereo observations has long been known to impact the accuracy of topographic measurements (e.g. Hallert (1960), Johnsson (1960)), so it is no surprise that it is also an important factor in descent stereophotogrammetry. This study, which is also published in Brydon et al. (2021), shows that accuracy of elevation measurement is a strong function viewing angle, thus carrying significantly over the surface of the observed body, and resulting in optimal stereo geometries. From the measured elevation RMSEs presented above, an optimum ground-radius-to-height ratio of 1-3 is estimated, whilst the optimum height-to-ground-radius ratio is found to be ~0.2-0.5 (similar values are predicted by the error model of section 5.6.1.4). Additionally, the measured elevation RMSEs indicate an optimum baseline-to-height ration of 0.2-0.4.

Notably, the model curves of section 5.6.1.4 indicate significantly larger optimum baseline-to-height ratios of at least 1 for near surface, and as much as 3-4 for distant surface. In the measured curves, measurement error grows rapidly as baseline increases beyond ~0.5 and ceases to follow the model behaviour. The reason for this contrast with the modelled curves is that the error model does not encapsulate the dependence of angular uncertainty ($\sigma(a_1), \sigma(a_2)$) on the surface appearance within the images of a stereo pair. Larger baselines result in a more significant change in surface appearance, whereby the shape and size of a given surface feature can differ markedly between images, increasing the likelihood of inaccurate or failed image matching. This highlights the importance of testing imaging techniques on representative data, and the value of utilising planetary datasets to simulate camera products. The suitability and performance of different stereo matching algorithms, which may allow the use of larger baselines and yield smaller elevation errors, would be a valuable topic of further study.

The findings of this study apply not only to the penetrator descent camera concept of this chapter, but also the general concept of wide-angle descent stereophotogrammetry, and they represent the first comprehensive quantitative analysis of such a technique. In

addition to this, they demonstrate the applicability of high fidelity DTMs and surface images to simulated studies of imaging instrumentation. The methods used in this study could be readily utilised and adapted for the study of other specific descent stereophotogrammetry setups.

Valuable areas for further research on this topic include the comparison and optimisation of different stereo matching algorithms for descent image pairs, the impact of photometric effects, and the recovery of topography when camera positions and/or orientations are uncertain. The nature of the descent topographic measurements presented here share some characteristics, particularly their low density and incomplete coverage at oblique viewing angles, with those obtained from rover imagery (e.g. Tao et al. (2016) and their Figure 11). Some areas of active research associated with rover stereo imagery, in particular its fusion with orbital data and use for localisation, would also potentially be applicable to descent imagery, and this too would be a valuable area of further research (Tao, et al., 2016).

5.7 Discussion

This chapter has described and investigated a descent camera concept for use on a planetary penetrator. Scanning motion of the penetrator, and wide-angle optics, permit large surface coverage around the site of the penetrator's impact. The imaging geometry results in significant variation in the viewing conditions over the camera's FOV, and the opportunity to capture images from multiple altitudes provides a range of perspectives. Section 5.3 discusses the effect this has on imaging resolution and scene motion. The camera's radial GSD is the limiting factor in pixel scale, and dictates an ideal relationship between imaging altitude and ground radius (1:1) if best possible imaging resolution is to be achieved (section 5.3.2.4).

Section 5.6 details a study of performing stereophotogrammetry with the camera's images, using their vertical offset as a stereo baseline to infer surface shape. For this, simulated images of a high resolution Martian dataset, first introduced in section 5.4, were used. Recovery of surface topography surrounding the descent location was shown to be feasible. Comparison of measured topography with the original surface model indicated a range of optimum stereo geometries, which, like the camera's GSD characteristics, can be used to inform image planning.

To demonstrate this, let us consider an example penetrator descent to Europa's surface. A vertical free fall, starting from a stationary position 30 km above the surface, will be assumed. For maximum surface coverage and context, let the first image be taken at

30 km altitude. Considering first an objective of imaging the entire visible surface with a radial GSD no greater than 150% the minimum possible value for any given ground radius, Figure 5.13 indicates an acceptable height-to-ground-radius ratio range of 0.4-2.5. Using equation 5.14, images should, at a minimum, be captured at altitudes of 30 km, 4.8 km, 0.77 km, 0.12 km and, if possible, 0.02 km (by this stage of the descent, there would be significant vertical motion of the penetrator during an image capture). The radial GSDs with which each of these images captures the surface are plotted in Figure 5.50.

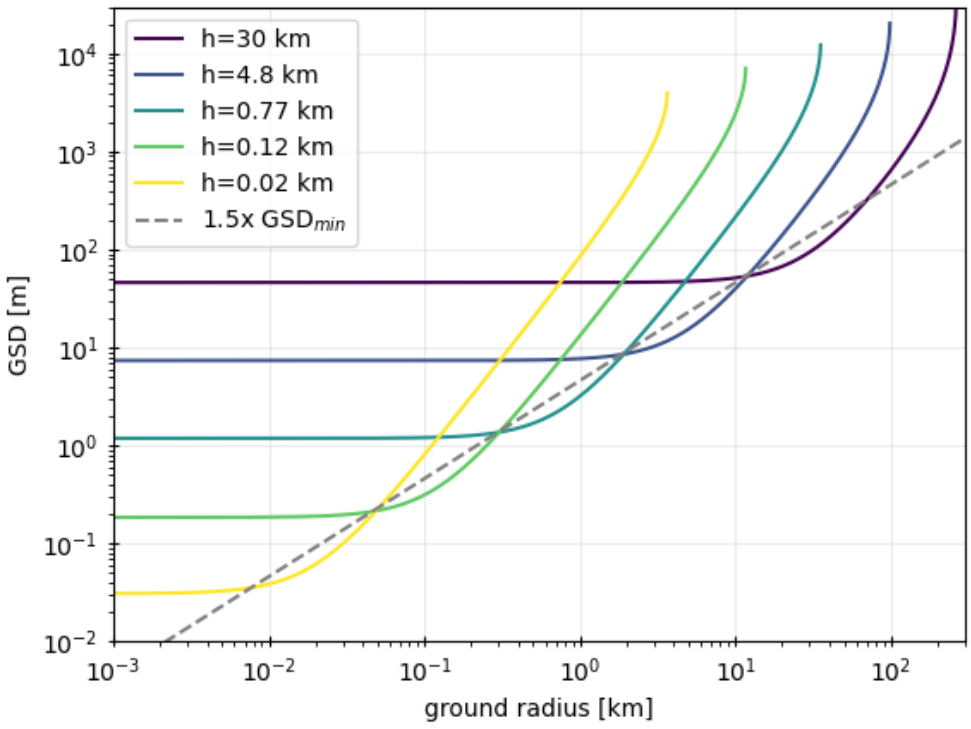

Figure 5.50: Radial GSD of camera with angular resolution 0.088° per pixel, from 5 imaging heights during descent to Europa, as a function of ground radius. The dashed line plots 1.5x the minimum possible GSD at each ground radius.

The set of images plotted in Figure 5.50 allows all surface within ≈100 km of the landing site to be imaged with radial GSDs no greater than 150% the minimum possible value (grey dashed line). Surface within 50 km is imaged at 100 m per pixel or less, whilst surface within 2 km is measured with GSDs less than 10 m. At their smallest, pixels cover approximately a metre down to tens of centimetres, depending on the fidelity of the final images. This is achieved with just 5 images, but it requires an intensive period of image collection toward the end of the descent. The last 1 km of descent, during which three

images should be captured, lasts only 3 seconds, less than 2% of the total descent duration. The times and heights at which these images could be acquired (as well as the number that could be acquired) would likely be dictated by what imaging cadence the camera were capable of, rather than the altitudes of optimal imaging geometry. In addition to this, vertical motion blur starts to become significant below 1 km. Therefore, achieving optimal imaging positions would probably not be feasible below 1-2 km. Any images below this would still have low absolute GSDs in the near-nadir region compared to higher altitude images.

The resolution of the descent images may not be the only consideration when planning their acquisition. Topography elevation measurement is optimally performed at a height-to-ground-radius ratio of 0.2-0.6, with a stereo baseline of 0.3x imaging height. Using equation 5.14 as before, this can be achieved with the set of images listed in Table 5.2.

Table 5.2: Altitudes of stereo image pairs for optimal surface elevation measurement.

Lower Image Altitude [km]	Higher Image Altitude [km]
23	30
7.7	10
2.6	3.4
0.9	1.2
0.3	0.4

Table 5.2 shows that, once again, optimal image altitudes are biased toward the end of the descent, and accurately achieving more than three stereo pairs of optimal imaging geometry would likely be challenging. Accounting only for the first three stereo pairs of Table 5.2, it is estimated that a region of surface extending from approximately 1 to 50 km could be well measured, with approximate elevation accuracy as is plotted in Figure 5.51.

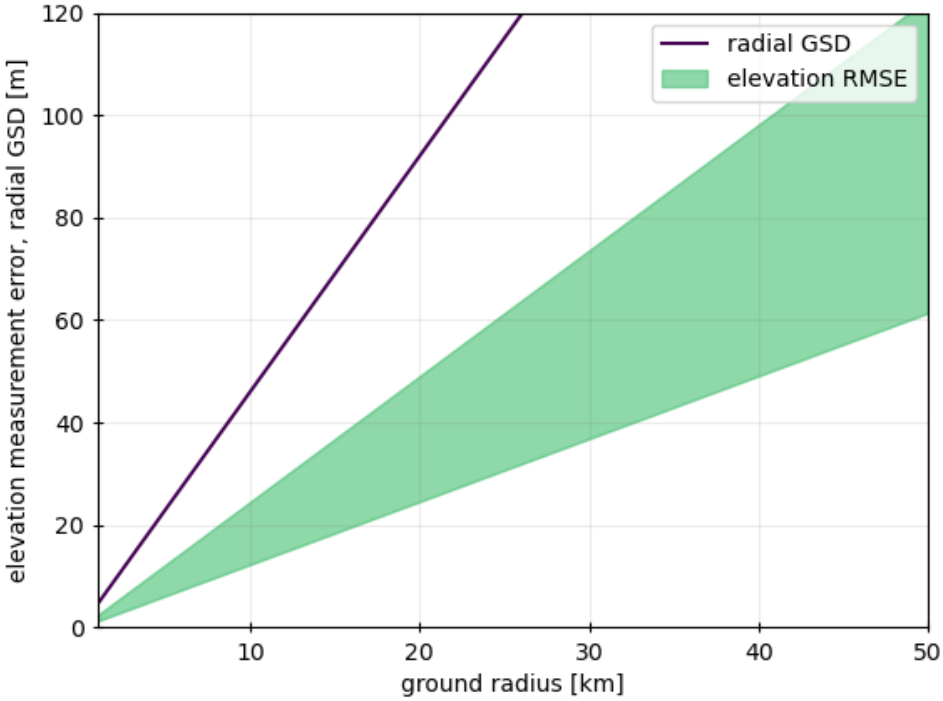

Figure 5.51: Approximate range of achievable accuracies of surface elevation measurement, as a function of surface ground radius. Also plotted is the approximate radial GSD of the stereo images.

Figure 5.51 indicates elevation accuracies of 100 m or better for surface within 40 km of the landing site, and as low as 10 m within 5 km. Steinbrügge et al. (2020) estimate their localised Europa DTMs, derived from stereo Galileo images, to have a comparable vertical precision.

As mentioned at the beginning of this chapter, a major benefit of penetrators is their potential to access surfaces and sub-surfaces at relatively low cost, which may otherwise require high cost missions or be inaccessible all together. Equally, the value of including a camera aboard a penetrator comes from the opportunity to image a surface which hasn't yet been imaged, or improve/supplement previous image data by achieving improved observations. As a demonstration, and in order to put the capabilities of the camera into context, Figure 5.52 shows the penetrator descent camera's radial GSD as a function of the penetrator's altitude, during its descent to Europa. Four lines are plotted, each showing the radial GSD of a different fixed point on the surface. Also marked on the plot, for comparison, are the approximate GSDs at which major solar system bodies have already been imaged (vertical lines indicate the range of GSDs with which a body has been observed). The plotted GSD curves can be considered the largest GSD within

their associated radius of the nadir (i.e. the line corresponding to 100 km from nadir indicates the largest GSD at which the surface is imaged anywhere within 100 km of nadir). The radii of 100 km, 10 km, 1 km and 0.1 km correspond to approximate areas on the surface of 30,000 km^2, 300 km^2, 3 km^2 and 0.03 km^2 respectively.

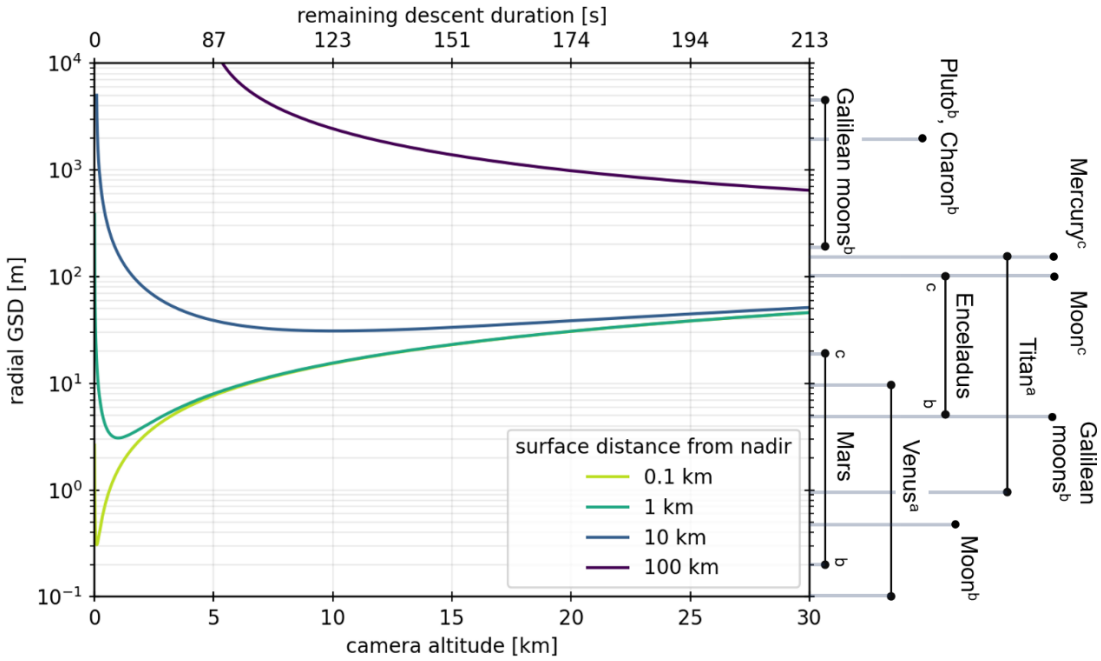

[a] Coverage localised to probe descent/landing site.
[b] Limited coverage.
[c] Global coverage.

Figure 5.52: Across-track GSD with which four fixed points on the surface (defined by their radii from nadir) are imaged versus camera altitude (bottom axis) and corresponding time until impact (top axis). The approximate GSDs of previous observations of a variety of solar system bodies are indicated for comparison (vertical lines indicate ranges of GSDs).

It is evident from Figure 5.52 that penetrator descent imaging has the scope to provide improved GSDs at a wide variety of solar system bodies. On Pluto and Charon, where image sets of significant surface coverage (>25%) have GSDs on the order of 1 km, a descent camera could provide improved resolution over the whole of its FOR (Moore, et al., 2016; Stern, et al., 2015). Whilst the Galilean moons (Carr, et al., 1995; Clark, et al., 1998; Greeley, et al., 1998; McEwen, et al., 2000) and Enceladus (Roatsch, et al., 2008) have been imaged at GSDs comparable to the best of the descent camera (~1-10 m), their coverage is limited, and a penetrator could well be deployed outside of this coverage. Larger (global in the case of Enceladus) surface coverage of these bodies exists with longer GSDs (~100-1000 m), which exceed the GSDs achieved by the

penetrator camera for much of its descent. Titan (Karkoschka, et al., 2007) and Venus (Ksanfomality, 2012), whose visible surface images were obtained by landers, have been imaged with similar ranges of GSDs to those expected from the descent camera. However, these images are very localised, and at these bodies, penetrator descent imaging would very likely provide the opportunity to observe unseen surface.

For Mars (Jaumann, et al., 2007; Malin, et al., 2007; McEwen, et al., 2007) and the Moon (Li, et al., 2010; Robinson, et al., 2010), the most thoroughly imaged solar system bodies, a penetrator descent camera may well not outperform the highest resolution existing images, but these have limited coverage. Because Mars and the Moon have also been globally mapped at high resolutions, improved GSDs would likely only be achieved during the late stages of a descent, for the near surface.

The analysis of this chapter need not apply solely to a penetrator descent, nor be restricted to scanning imagers. Similar imaging geometry can be achieved with fisheye or hemispherical optics and framing cameras on spacecraft performing a descent or ascent (e.g. balloon probes, powered descent stages and aerial vehicles). The research presented here provides a foundation from which more specific scenarios can be studied (e.g. a certain solar system body, more constrained camera design).

6 EnVisS: The Entire Visible Sky Camera for the Comet Interceptor Mission

In this chapter, a rotationally scanning camera is investigated for a very different application to that of chapter 5: imaging the large scale structures of a comet during a spacecraft flyby. The imaging concept – rotational push-frame with a wide field of view – is similar to the penetrator descent camera of the previous chapter, but is designed to image extended features across the entire sky, rather than map a planetary surface. The work that was carried out for and is presented in this chapter focusses on a real instrument, under development by an instrument team comprising members at several institutes across Europe. The instrument, called EnVisS (Entire Visible Sky) was devised for and exists as a part of the scientific payload for the European Space Agency's Comet Interceptor mission. The Comet Interceptor mission and the EnVisS imager are introduced in section 6.1. In section 6.2, a more detailed description of EnVisS is given, including calculations of its imaging geometry and operation. Sections 6.3 and 6.4 describe a comprehensive study of the expected signals and noise in EnVisS images, and possible methods to optimise them. Finally, EnVisS' imaging polarimetry is introduced, described and analysed in section 6.5 before the chapter's findings are summarised in section 6.6.

6.1 The Comet Interceptor Mission and the EnVisS Camera

6.1.1 Comet Interceptor

In June 2019 Comet Interceptor was selected by the European Space Agency (ESA) as its first F-class mission. The F, standing for fast, reflects the fact that the mission will be developed on a relatively short timescale of under ten years from concept to launch. It is planned to launch in 2029 alongside the Ariel spacecraft, sharing the same launch vehicle (Snodgrass & Jones, 2019).

The overarching goal of the Comet Interceptor mission is to perform a flyby of a long-period comet (LPC, by definition a comet with an orbital period greater than 200 years), that preferably is also a Dynamically New Comet (DNC). DNCs are a subset of LPCs which have never yet passed through the inner solar system (Levison, 1996). This is in contrast to short period comets such as 1P/Halley and 67P/Churyumov-Gerasimenko, both of which have been visited by highly successful spacecraft missions (discussed in chapter 4, section 4.3). Short period comets and long period comets which have already passed through the inner solar system are altered from their original state by solar heating. Some nucleus material ejected by outbursts of subliming gas as a comet's temperature increases slowly returns, due to gravity, to the surface, resulting in an altered structure and distribution (Keller, et al., 2017).

The comparatively unaltered state of DNCs, having not gone through this process, therefore provides an opportunity to see the comet in its original form, dating back to the early stages of the solar system. This is something that has never before been done. The reason for this is that DNCs approaching the inner solar system can only be detected at relatively close range, typically a few astronomical units (AU), providing too short a warning time to design and launch a mission to visit them before they have already passed through and departed the inner solar system once more. The Comet Interceptor mission overcomes this by launching to space before its target comet has been identified. It will wait at Lagrange point L2, whilst Earth based telescopes search for approaching DNCs whose orbits are suitable for intercepting. Once a suitable target is identified and selected, the Comet Interceptor spacecraft will depart from L2 on course to perform a close flyby of the comet. The search for and identification of a suitable target will rely largely on the currently-under-construction Vera Rubin Observatory and its Legacy Survey of Space and Time (LSST, formerly standing for Large Synoptic Survey Telescope).

The Comet Interceptor mission will comprise three separate spacecraft: a main spacecraft (A) and two sub spacecraft, or probes (B1 and B2). Spacecraft A will carry B1 and B2 until the flyby of the target body commences, when B1 and B2 will each separate and follow their own flyby trajectories. This will allow multipoint measurements of the comet and its environment, permitting spatial and temporal variations to be characterised in a way not possible with a single spacecraft.

The 3-axis stabilised main spacecraft A will pass the comet at a relatively large, safe distance, where the dust environment will not pose a serious risk to its survival. Its science payload will include dust and plasma analysers, an infrared (IR) imager and a

visible camera which will image the nucleus at high resolution. Spacecraft B1, which is being built by the Japanese Space Agency JAXA (whereas spacecraft A and B2 are both provided by ESA), is a small 3-axis stabilised spacecraft carrying a suite of plasma instruments and several imagers. It will approach the nucleus closer than spacecraft A. Spacecraft B2 is spin stabilised, and will pass closest to the nucleus of all three craft. Much like the Giotto spacecraft (see chapter 4, section 4.3.2), its spin axis will be aligned with its velocity relative to the comet, and a forward facing dust shield will protect it from impacts. Because it is intended to pass close to the nucleus, its chances of survival are the lowest of all three spacecraft. It will collect data on its inbound journey and transmit them to spacecraft A, and if it survives beyond its closest approach to the comet, will have the opportunity to make and transmit further observations on its outbound journey. Spacecraft B2 will carry field and particle analysers, a visible framing camera, and the instrument which is the subject of the work presented throughout this chapter: EnVisS. A summary of the three spacecraft's intended instrument payloads at time of writing are given in Table 6.1.

Table 6.1: Instruments of the Comet Interceptor mission.

Spacecraft	Instrument		Details
A	MIRMIS	Modular Infrared Molecules & Ices Sensor	Hyper and multispectral imager addressing ices, emission, surface composition and temperature.
	CoCa	Comet Camera	Mission's primary science camera. High resolution imaging of nucleus.
	MANiaC	Mass analyser for Neutrals in a Coma	Mass spectrometer.
	DFP	Dust Field and Plasma	Magnetic field, electric field and plasma parameters.
B2	OPIC	Optical Imager for Comets	Nucleus, inner coma and dust jets.
	EnVisS	Entire Visible Sky	Wide angle scanning camera. Multispectral and polarimetric sky mapping.
B1	HI	Hydrogen Imager	Lyman-alpha emission imager.
	PS	Plasma Suite	Measurement of magnetic field and plasma velocity distribution.
	NAC and WAC	Narrow and Wide Angle Cameras	High latitude nucleus imaging to complement CoCa.

6.1.2 EnVisS

The EnVisS (Entire Visible Sky) imager is a wide angle scanning camera, which will capture images from aboard spacecraft B2. At the time of writing, the instrument is undergoing development, and some aspects of its design are not fully constrained, and will likely change. Therefore, in this thesis, details of the instrument will often be described as provisional, indicating that they are reasonably representative of the final design and have been used in the analysis presented here, but may differ in the final instrument.

EnVisS is in many ways similar to the descent camera concept studied in chapter 5. It will image through wide-angle optics in order to achieve a large across-track FOV, and the spin of the spacecraft on which it is mounted will scan that FOV across the sky to build up large images. Images will be captured by an array detector, on which narrow optical filters will be mounted, allowing multichannel imaging. There will be both spectral and polarising filters. Utilising fisheye optics, the camera will achieve a ~180° across-track FOV. The camera will be mounted on the side of spacecraft B2, and for the work of this thesis, two possible configurations are considered:

- The principal axis of EnVisS is perpendicular to the spacecraft spin axis, such that the camera's 180° FOV extends between the parallel-to-spin-axis and antiparallel-to-spin-axis directions (Figure 6.1 (a)).
- The principal axis of EnVisS is tilted toward the antiparallel-to-spin-axis direction, resulting in a blind spot in the forward facing direction and an oversampled region in the rearward facing direction (Figure 6.1 (b)).

The latter case allows the entirety of the EnVisS instrument to be protected by the spacecraft's dust shield, reducing the likelihood that the instrument is destroyed by a dust impact during the flyby. However, the former case benefits from complete coverage of the sky, as well as its unmitigated view of the comet's inner coma throughout the whole of spacecraft B2's inbound trajectory.

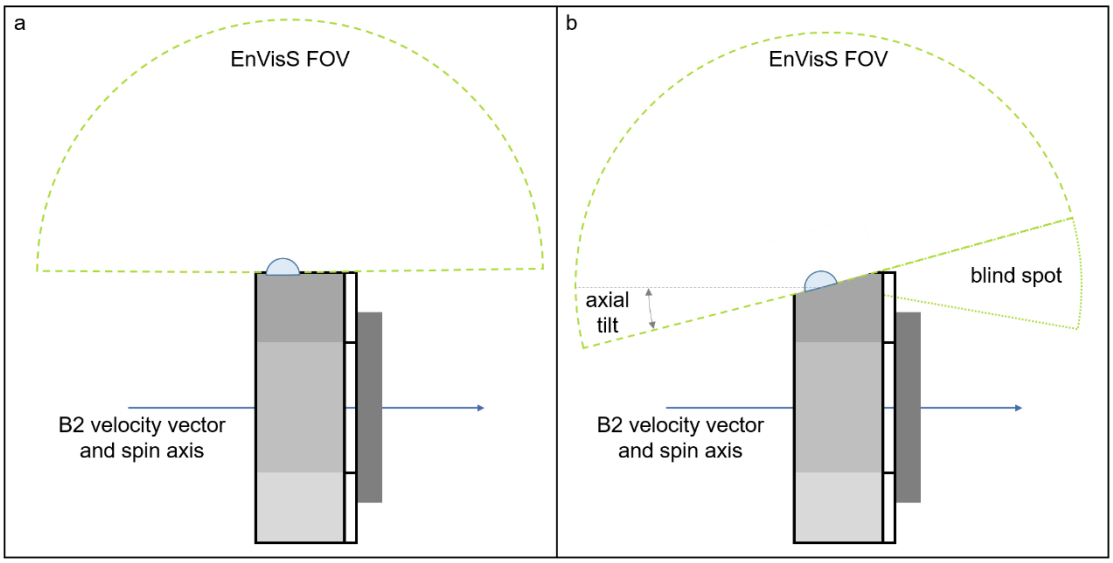

Figure 6.1: A profile illustration of the two considered configurations for mounting EnVisS on spacecraft B2. The positioning of EnVisS behind the dust shield in (b) provides protection from dust impacts, at the cost of a blind spot in its coverage.

The high-level science goals originally envisioned for EnVisS are listed in Table 6.2.

Table 6.2: EnVisS high-level science goals.

Science Goal	Measurement
Characterise the dust's physical nature and its spatial variation.	Polarimetric sky mapping at multiple locations within the coma.
Observe ion tail features and their three-dimensional structure.	Ion tail imaging during approach and traversal of coma.
Determine structure of dust jets	All-sky images of continuum at multiple locations within coma.
Determine structure of gas jets	Neutral gas filter images at multiple locations.

6.2 EnVisS: Camera Description

Whilst the concept of the EnVisS camera was briefly introduced in section 6.1.2, here a more formal and thorough description is given, including analysis of its imaging geometry and operation. At the time of submission of this thesis, the camera is very much still under development, with the ultimate end target of integration onto the Comet Interceptor B2 spacecraft in 2026, and launch in 2029. Whilst the instrument's design is largely constrained, some aspects will inevitably change between the conclusion of this thesis and the launch of Comet Interceptor. Indeed, the design has already evolved several

times from its original concept. Some of these changes were influenced by the work of this thesis, whilst many were driven by other technical and programmatic factors. The calculations of instrument fields of view, fields of regard, resolution, exposure durations and data form presented throughout this section were carried out for this thesis, and played an integral part in the development of the EnVisS instrument. The EnVisS optical design, without which these calculations could not be made, was the work of EnVisS team members at CNR-IFN and INAF-OaPD.

At the time that this thesis was submitted, the structure of the EnVisS team and the responsibilities of the institutes involved were changing. Table 6.3 outlines the members and structure of the EnVisS team as it was throughout the period that this thesis' work was carried out. During this time, the instrument was led by principal investigator Geraint Jones at MSSL and co-principal investigator Vania Da Deppo at CNR-IFN.

Table 6.3: EnVisS instrument team member institutes.

Institute	Responsibilities
UCL MSSL, UK	Lead institute; focal plane assembly; detector and on-detector filters.
CNR-IFN, Italy	Optics, mounting and baffle.
INAF-OaPD, Italy	
IAA-CSIC, Spain	DPU and power supply.
Aalto, Finland	Data handling, communication between EnVisS and OPIC instrument.

6.2.1 Imaging Geometry

EnVisS will utilise the push-frame technique to acquire its images. As discussed in section 2.2.2.2, this allows large-coverage multichannel imaging with a short cadence – particularly important given the high velocity and relatively short duration of Comet Interceptor's flyby. A fisheye lens will provide the necessary FOV to achieve all-sky coverage. This section describes the resulting imaging geometry of the camera.

6.2.1.1 EnVisS' Optics

EnVisS will image through wide-angle optics, covering ~180° across the FOV's largest extent, with a narrower ~45° coverage perpendicular to this. Such a large field of view requires a fisheye lens, whose projection geometry differs from the far more common central perspective geometry used for narrow FOV cameras. Imaging geometries and mapping functions are introduced in section 2.1. The details of the design and

performance of EnVisS' optical design are described by Pernechele et al. (2021). The optics are designed to use an equidistant imaging geometry. All analysis, modelling and simulation of EnVisS carried out for and presented in this thesis assumes an ideal equidistant imaging geometry.

EnVisS will be mounted such that the 180° axis of its FOV is in plane with the spacecraft's spin axis (as described in section 6.1). If the camera is mounted in the tilted configuration, it will be aligned with a provisional 10° axial tilt (see Figure 6.1 (b)).

6.2.1.2 EnVisS' Focal Plane

The focal plane design and layout of EnVisS is crucial to its mode of image capture. In order to acquire images covering the whole sky, EnVisS will use the full 180° FOV of its fisheye optics. Such optics project a full hemisphere onto a circle on the camera's image plane. The optics are designed such that this circle's diameter will span the active region of the detector. A diagram of the EnVisS focal plane is shown in Figure 6.2. The detector will be a CMV4000 CMOS, within a 3D Plus camera head, and will have an active region of 2048 by 2048 pixels. Optical filters (both spectral and polarimetric) will be mounted directly over the detector. Each filter's footprint on the detector will be long and narrow, spanning all the detector's rows but only a few columns. Filters will be placed adjacent to each other, and centred on the optics' principal point. The exact number and dimensions of filters are not yet finalised. Provisionally, they will each have an along-track FOV of ≈5° at its narrowest (corresponding to a width of 50 pixels on the detector). The regions of the detector which are not covered by a filter will not be used for imaging, permitting some unused portions of the fisheye optics' primary lens to be cut away to save mass.

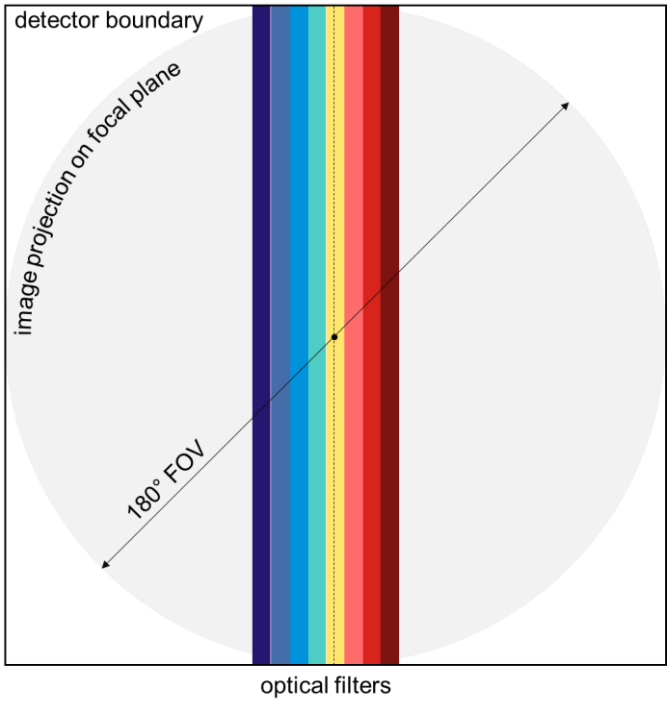

Figure 6.2: The EnVisS focal plane. The central black circle indicates the principal point of the imaging geometry. The eight vertical stripes illustrate the placement of eight example optical filters. The dashed line indicates both the vertical axis of the detector and the orientation of the 180° across-track FOV. The detector boundary describes 2048 columns and 2048 rows of pixels.

A coordinate system, called normalised detector coordinates, will be used throughout this chapter to describe positions on the focal plane. This is an alternative to defining a position by its detector column and row numbers, and is adopted to keep analysis applicable to imaging systems of different resolutions. For this thesis, this was first used before EnVisS' detector (and number of pixels) had been selected, but even with the detector now constrained, it remains valuable for considering the capture of binned images, and for generalising analysis to other similar cameras.

For a square detector with N_{pix} pixels per side (i.e. the detector comprises an array of $N_{pix} \times N_{pix}$ pixels), and a point on that detector with column and row number n_c and n_r respectively, the normalised detector coordinate of that point is given by

$$\vec{\tilde{p}} = \begin{bmatrix} \tilde{n}_c \\ \tilde{n}_r \end{bmatrix} = \frac{1}{N_{pix}} \begin{bmatrix} n_c \\ n_r \end{bmatrix} \tag{6.1}$$

Figure 6.3 illustrates the relationship between detector coordinates and normalised detector coordinates.

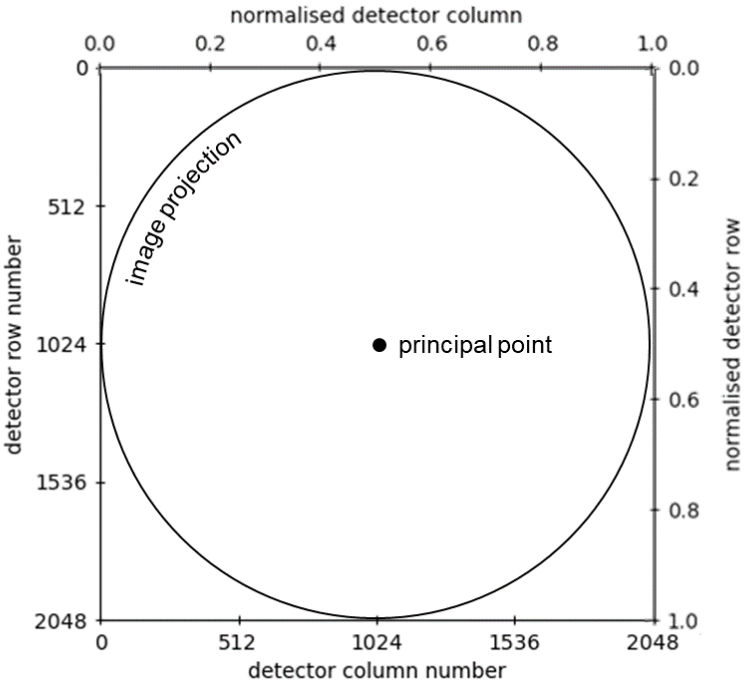

Figure 6.3: Detector coordinates and normalised detector coordinates for the EnVisS focal plane ($N_{pix} = 2048$).

6.2.1.3 Sky Coordinates

As with any scanning imaging system, the use of the push-frame technique means that EnVisS' viewing geometry will not only be a function of its optics, but also the instrument's scanning motion. Before describing the imaging geometry of the instrument further, it is useful therefore to first define a simple coordinate system that describes both the scan motion and EnVisS' view of the sky.

EnVisS will scan the sky by continuous rotation around its spacecraft's spin axis, which will be aligned with the spacecraft-comet relative velocity vector. Let that spin axis define the polar axis of a spherical coordinate system (henceforth referred to as B2-centric coordinates), with north pointing in the negative velocity direction (i.e. away from the comet). The origin of the coordinate system is the spacecraft's centre of mass (which can also be assumed to be the position of EnVisS). An arbitrary axis perpendicular to the polar axis defines a direction of 0° azimuth, such that the azimuth of the camera's principal axis steadily changes as the spacecraft spins. At the time of writing, the spin period of B2 is not finalised, but is constrained to a range of 4-15 s (with the shorter end of this range expected to be favoured to maximise spacecraft stability). Also undefined is the spin direction of EnVisS (i.e. clockwise or anticlockwise), but the design and

operation of the instrument is the same in each case (this work assumes a positive rate of change, i.e. anti-clockwise rotation when viewed from a position of < 90° polar angle).

EnVisS will be mounted on the spacecraft such that the projection of its central detector column (the column which coincides with the principal point; dashed line of Figure 6.2) at any instant forms a line of constant azimuth. Each of the camera's pixels views a direction with a certain polar angle (which is constant as the spacecraft spins) and azimuth (which changes as the spacecraft spins). The principal axis of the camera will be used to define its pointing azimuth, θ_{enviss}. This coordinate system is shown in Figure 6.4.

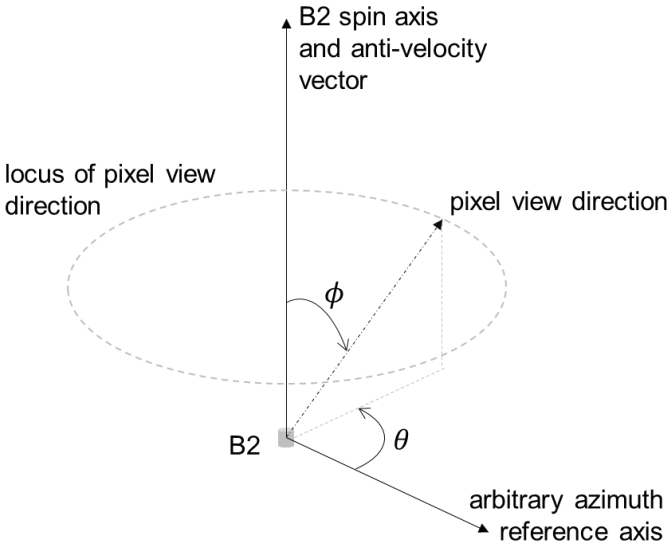

Figure 6.4: Illustration of the B2-centric coordinate system used in this thesis to describe EnVisS' imaging geometry. θ and ϕ are azimuth and polar angle respectively.

Figure 6.5 shows the region of sky, represented by this coordinate system, that is projected by EnVisS' optics onto its focal plane when $\theta_{enviss} = 0°$ for both untilted (panels a and b) and tilted (panels c and d) mounting cases. Because the B2-centric coordinate system is aligned with the spacecraft's spin axis, the polar angles seen by each pixel are independent of spacecraft orientation. By contrast, the azimuths viewed by the camera's pixels are a function of the spacecraft's orientation, and this provides the mechanism by which the camera scans a large FOR.

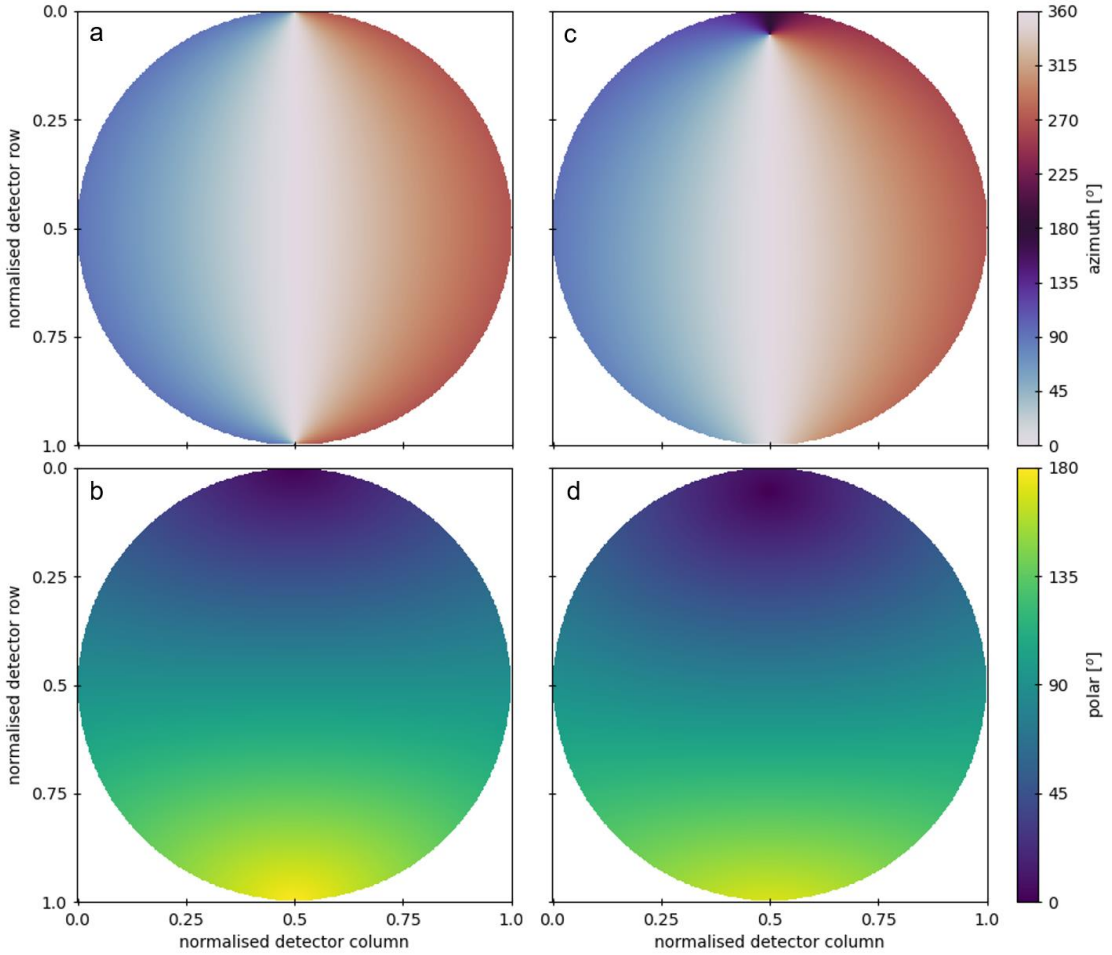

Figure 6.5: Azimuth (a and c) and polar (b and d) B2-centric coordinates seen by EnVisS with $\theta_{enviss} = 0°$ for untilted (a and b) and tilted (c and d) mounting. The azimuths seen vary as the spacecraft spins, whereas polar values are constant. The top of the focal plane views rearward and the bottom forward of the spacecraft.

The difference between the untilted and tilted mounting is most noticeable in Figure 6.5 (c), where the full 360° range of azimuths is captured instantaneously near the trailing spin axis, and also in Figure 6.5 (d), which shows that no polar angles greater than 170° are visible to the tilted mounting. For the majority of the camera's FOV, the effect on the imaging geometry by the camera tilt is minimal. One notable feature of the imaging geometry is that, for a given column on the detector, its across-track FOV depends strongly on its location on the detector. A central column observes the largest across-track FOV, whilst that viewed by a column near the edge is significantly smaller. Bearing in mind that EnVisS will image through narrow filters which run parallel to the detector's columns, the filters' FOVs will therefore differ from each other, and will depend on their positioning on the detector (Figure 6.6).

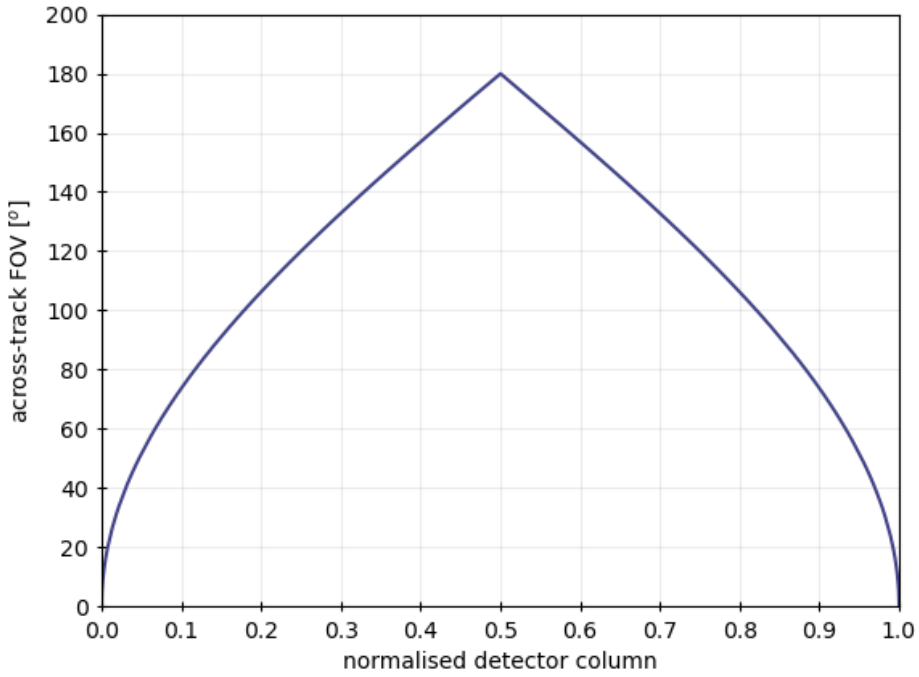

Figure 6.6: EnVisS across-track FOV as a function of normalised detector column.

For the tilted mounting of EnVisS, the across-track FOVs of Figure 6.6 do not equate to the across-track FORs (the maximum range of polar angles observed) achieved by the detector's columns. For example, the centre column of the detector samples a full 180° strip of the sky, but over the course of a full spacecraft spin, it does not image the whole sky, because two rearward viewing portions of the FOV (above and below the spin axis) cover the same polar angles. The geometry of this is illustrated in Figure 6.7, and the across-track FOR is plotted as a function of normalised detector column in Figure 6.8.

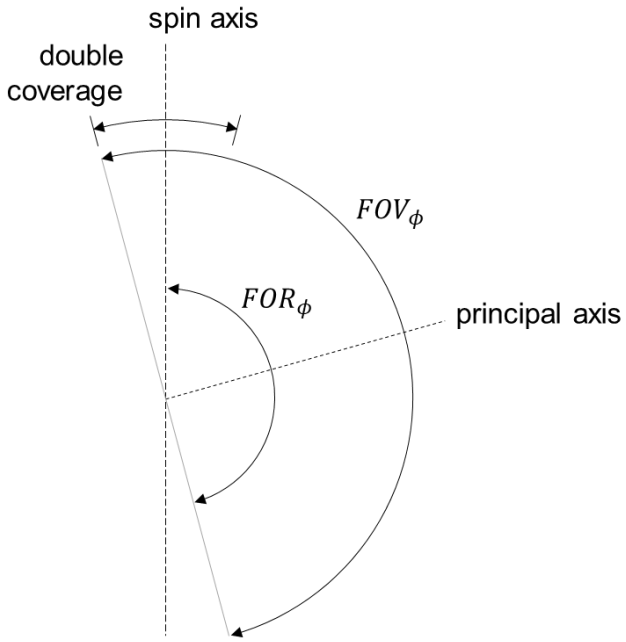

Figure 6.7: The difference between EnVisS' across-track FOV and across-track FOR for the tilted mounting case.

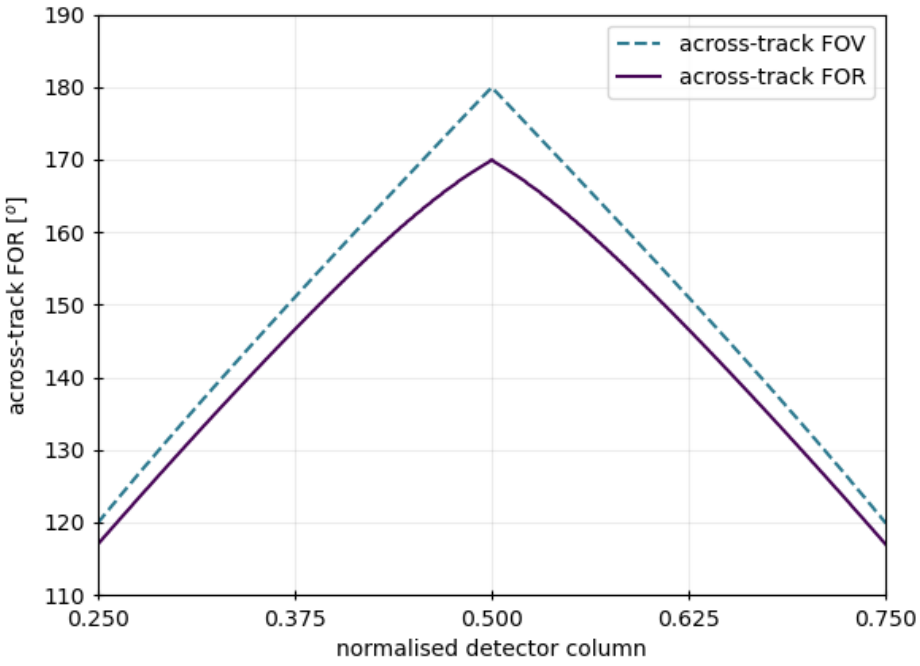

Figure 6.8: EnVisS across-track FOR as function of normalised detector column for the tilted mounting case. The dashed line shows the instrument's across-track FOV for comparison.

The 10° axial tilt of EnVisS in the tilted mounting case results in a maximum across-track FOR of 170°, achieved at the centre of the FOV. This is in contrast to the 180° across-track FOR of the untilted mounting, and results in a blind spot of 20° diameter, centred on the B2 spin axis. This blind spot obscures the nucleus and inner coma from the camera for a significant portion of B2's approach to the comet, and dictates the minimum visible impact parameter. The geometry of this is illustrated in Figure 6.9 (top), and the size of the blind spot at the nucleus and the minimum visible impact parameter are plotted in Figure 6.9 (bottom) as a function of B2's distance to closest approach. At the beginning of the flyby, even very large structures (10^6 km) will be out of EnVisS' view. The nucleus is likely to become visible ≈2000 km (≈40 s, depending on flyby velocity) from closest approach.

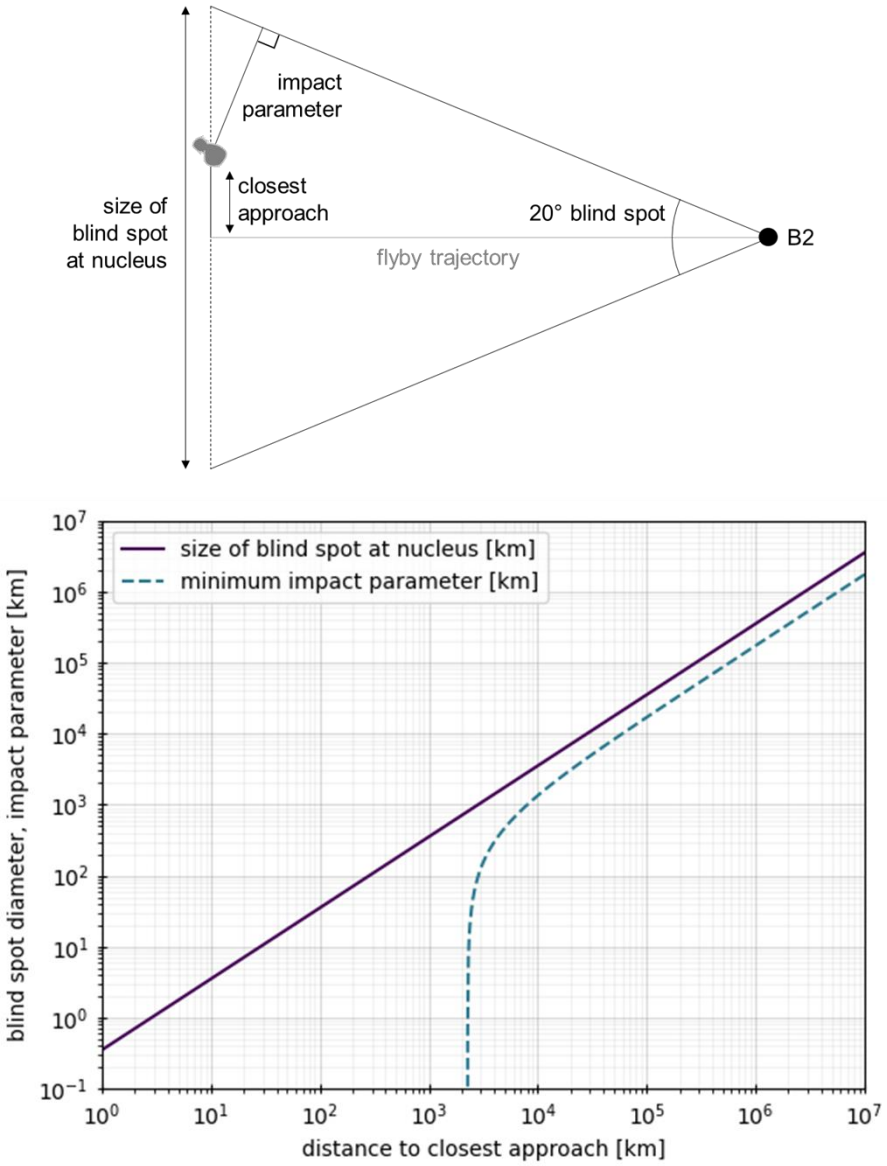

Figure 6.9: (top) Illustration of the geometry of EnVisS' blind spot when mounted with the tilted configuration. (bottom) Plot of the blind spot's size at the nucleus and the minimum visible impact parameter as a function of B2's distance to closest approach (assuming a closest approach of 400 km).

Plotting B2-centric coordinates on Cartesian axes (polar versus azimuth) leads to an equirectangular projection of the whole sky, with a standard parallel latitude of 0° (i.e. the plate carrée projection) (Botley, 1951). This representation of the sky will be used throughout this chapter to visualise EnVisS' imaging geometry and data. Figure 6.10 displays the EnVisS focal plane in this equirectangular all-sky projection. As the spacecraft spins, the projection of the focal plane on the sky maintains the same size,

shape and polar coordinates, whilst its azimuth coordinates uniformly increase/decrease in line with the spacecraft's spin rate and direction.

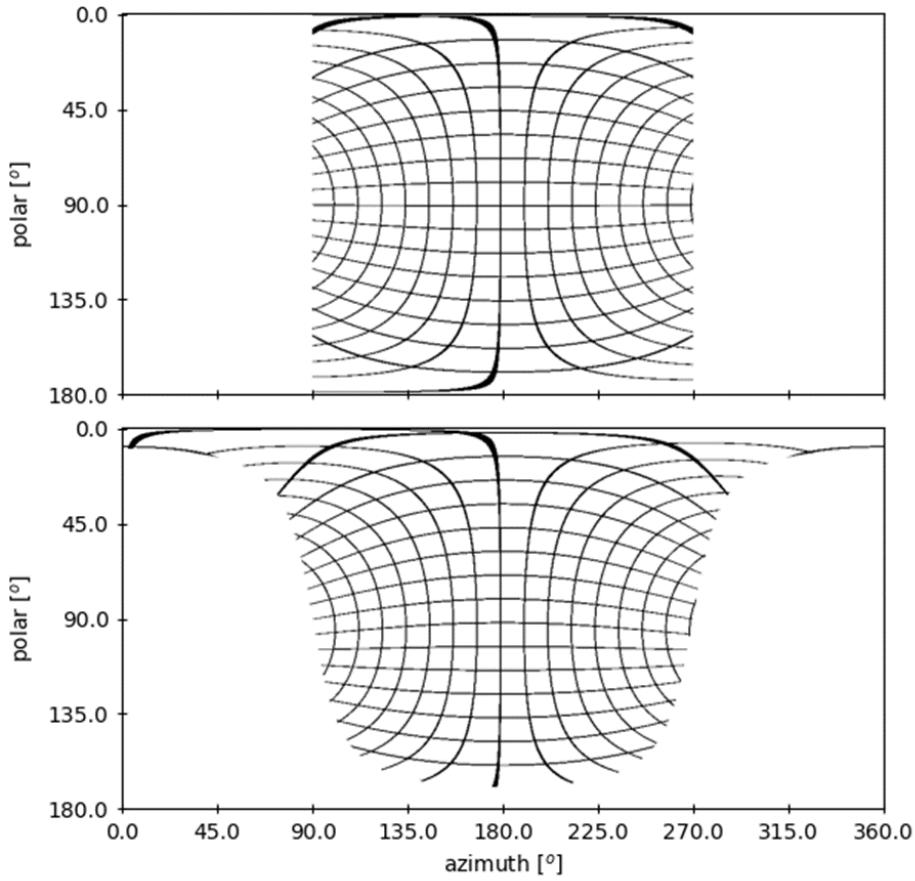

Figure 6.10: Illustration of how EnVisS' focal plane projects to B2-centric sky coordinates, with θ_{enviss} = 180°, for untilted (top) and tilted (bottom) mounting. The black lines divide the detector into 16 columns and 16 rows. The white space outside the black mesh is the portion of sky unseen by the FOV. Detector column number increases from right to left, row number from top to bottom.

With the untilted mounting, the projection of the EnVisS focal plane is symmetric about the 90° polar angle line, and spans 180° of azimuth at all polar angles. The tilted configuration spans the full 360° azimuth range at low polar angles, meaning this region of the sky is always in view, no matter the value of θ_{enviss} (c.f. Figure 6.5 (c)).

For completeness, Figure 6.11 shows the footprints of EnVisS' filters (assuming the provisional focal plane arrangement of Figure 6.2) projected onto the sky for both mounting configurations. It can be seen from Figure 6.11 how the across-track FOR of filters decreases as they move away from the detector's central column.

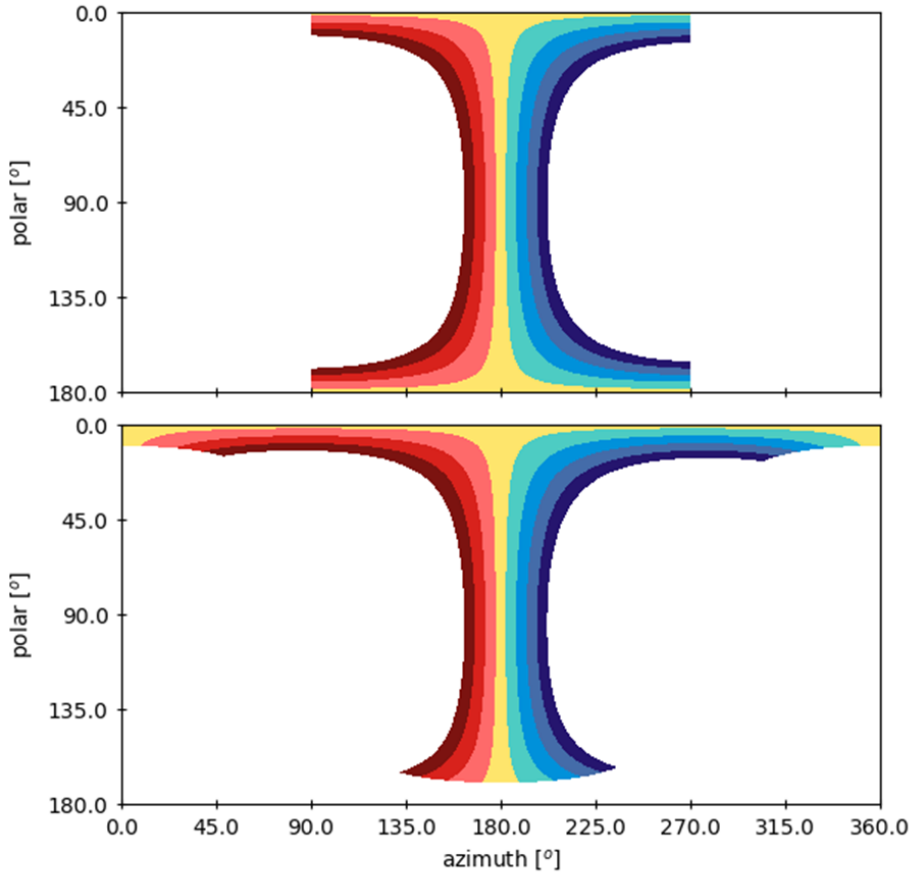

Figure 6.11: Projection of EnVisS' filters onto the sky, represented in B2-centric coordinates, with θ_{enviss} = 180°, for untilted (top) and tilted (bottom) mounting. Colours correspond to Figure 6.2.

6.2.2 Imaging Mechanics

For EnVisS, the nature of its images cannot be fully described just by the camera's imaging geometry, but must also consider the instrument's imaging mechanics (i.e. the operation of the camera: its exposure durations, imaging frequency, detector operation). These aspects will be addressed here, and their impact on the nature of EnVisS' data will be described. The analysis throughout this section is very similar for both the untilted and tilted mounting configurations of EnVisS. For simplicity, just the tilted case will be presented, and it can be taken for granted that values (e.g. exposure times) are functionally the same for the untilted case.

6.2.2.1 Exposure Time

When imaging, EnVisS will repeatedly capture framelets which, over the course of a 360° spacecraft rotation, build up to provide a full sky image. Each framelet is a narrow image,

every pixel of which is captured simultaneously. For clarity, the exposure time is the duration for which the detector collects signal in its capture of a single framelet, and is distinct from the term integration, which throughout this chapter refers to the camera's collection of a set of multiple framelets over the course of a 360° rotation, and their combination into an all-sky image.

Given that cometary comae and structures are relatively faint, the signals being collected by EnVisS will be well below the detector's saturation level, and exposure times will need to be as large as possible (this is addressed in detail in section 6.3). EnVisS will be moving relative to the scene it is imaging, due both to the spin and the flyby velocity of the B2 spacecraft. The upper limit of its exposure time is therefore dictated by the effect of motion induced pixel smear.

To determine EnVisS' exposure time, it is here assumed that motion blur degrades the angular resolution of an EnVisS image when pixel smear on the detector exceeds a length greater than the pixel pitch (i.e. the width of one pixel) during a single exposure. Thus, the maximum exposure time without motion blur is given by

$$\tau = \frac{1}{N_{pix}\left|\frac{d\vec{\tilde{p}}}{dt}\right|} \tag{6.2}$$

where τ is the exposure time, N_{pix} is the number of pixels spanning a side of the square detector and $d\vec{\tilde{p}}/dt$ is the rate of change of the scene's position in normalised detector coordinates.

Assuming initially that the spacecraft is stationary relative to the comet, and its only motion is its spin stabilisation, the pixel motion can be expressed in terms of the spacecraft spin rate and the derivative of normalised detector coordinates with respect to spacecraft azimuth:

$$\frac{d\vec{\tilde{p}}}{dt} = \frac{d\vec{\tilde{p}}}{d\theta_{enviss}} \cdot \frac{d\theta_{enviss}}{dt} = \frac{d\vec{\tilde{p}}}{d\theta_{enviss}} \cdot \frac{2\pi}{T_{spin}} \tag{6.3}$$

where T_{spin} is the period of the spacecraft's rotation. Equations 6.2 and 6.3 combine to give the spin-limited exposure time (τ_{spin}) in terms of N_{pix} and T_{spin}:

$$\tau_{spin} = \frac{T_{spin}}{2\pi\, N_{pix}\left|\frac{d\vec{\tilde{p}}}{d\theta_{enviss}}\right|} \tag{6.4}$$

Moving N_{pix} and T_{spin} to the left hand side of equation 6.4 gives a 'normalised' exposure time, (herein called the exposure factor, $\tilde{\tau}$), which is independent of these two parameters, and a function only of the optics' imaging geometry:

$$\tilde{\tau}_{spin} = \frac{N_{pix}}{T_{spin}} \tau_{spin} = \frac{1}{2\pi \left| \frac{d\vec{\tilde{p}}}{d\theta_{enviss}} \right|} \tag{6.5}$$

The spin-limited exposure factor of equation 6.5 over the whole of the EnVisS focal plane is plotted in Figure 6.12. An absolute exposure time is found by multiplying exposure factor by T_{spin}, and dividing by N_{pix} (equation 6.5). Use here of the exposure factor rather than exposure time permits assessment and comparison of different spin rate and resolution combinations, the former of which is, at the time of writing, constrained to a possible range of 4-15 s, and the latter of which will vary due to the use of pixel binning (covered in section 6.4.1).

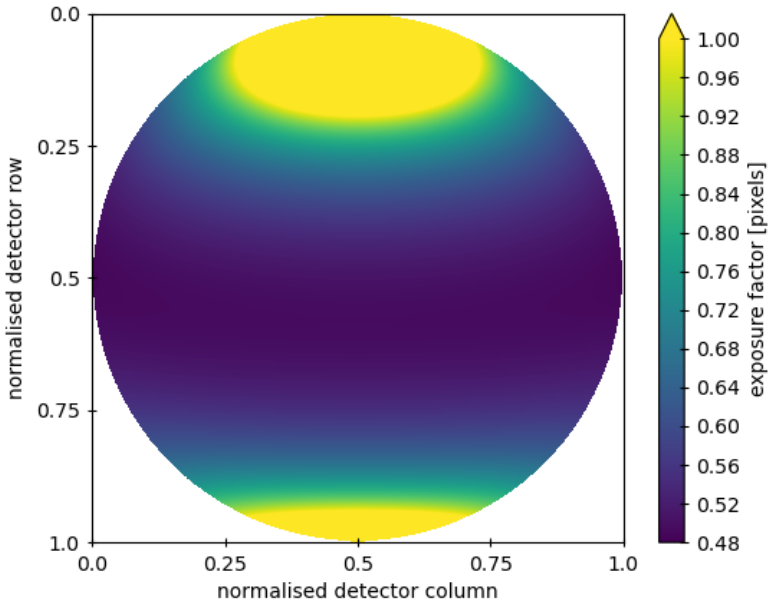

Figure 6.12: Spin-limited exposure factor (equation 6.5) for a stationary, spinning EnVisS. Multiplying by T_{spin} and dividing by N_{pix} gives the corresponding spin-limited exposure time.

The instrument's minimum spin-limited exposure factor occurs in the central portion of the detector. Over a fairly wide horizontal band covering the majority of the image the value differs little from this minimum. At the vertical extremes of the focal plane, near to the projections of the spin axis' poles on the detector, scene motion is slower and thus

exposure factors are larger. Exposure time should be based on the minimum value of exposure factor if motion blur is to be totally avoided. Increasing the exposure time beyond this would introduce motion blur first at the equator of the image, which would migrate towards the poles as exposure length grew. Table 6.4 gives some example exposure times, based on Figure 6.12, for three different spacecraft spin periods and detector resolutions.

Table 6.4: Absolute spin-limited exposure times for an example set of T_{spin} and N_{pix} values.

T_{spin} [s]	N_{pix} [pixels]	τ_{spin} [ms]
4	512	3.9
	1024	1.9
	2048	0.97
8	512	7.8
	1024	3.9
	2048	1.9
12	512	11
	1024	5.8
	2048	2.9

Separate from the spin of B2, the spacecraft's velocity relative to the target comet also causes scene motion across the focal plane. Whilst the spin causes a constant pattern of motion across the detector, the motion induced by the flyby speed will have its strongest effect for closer objects, and so is dependent on the structure of the scene. Broadly, it will become increasingly significant as the spacecraft nears closest approach, and an important question is whether and to what extent it will contribute to motion blur and impact the possible exposure times as the flyby progresses.

For the analysis of this effect, let us now consider EnVisS on a non-rotating spacecraft which is travelling with constant velocity relative to the target comet. Let l be the spacecraft's distance along it's flyby trajectory from some arbitrary start point. In an analogous manner to equation 6.3, the motion of a point in real space across the detector can be expressed in terms of the spacecraft velocity v and the derivative of normalised detector coordinates with respect to l:

$$\frac{d\vec{p}}{dt} = \frac{d\vec{p}}{dl} \cdot \frac{dl}{dt} = \frac{d\vec{p}}{dl}\, v \tag{6.6}$$

It is helpful to rewrite equation 6.6 in terms of $d\vec{\tilde{p}}/d\phi$, the rate of change of pixel coordinates with respect to polar angle. Because ϕ is measured in a coordinate system centred on EnVisS, the constant spacecraft velocity, which is parallel to the coordinate system's polar axis, results in a changing polar angle, $d\phi/dl$ of the observed point in space (recall Figure 6.4).

$$\frac{d\vec{\tilde{p}}}{dt} = \frac{d\vec{\tilde{p}}}{d\phi} \cdot \frac{d\phi}{dl}\, v \tag{6.7}$$

In equation 6.7, $d\vec{\tilde{p}}/d\phi$ is a function of the optics' imaging geometry, and varies over the camera's field of view, whilst $d\phi/dl$ is dependent on the scene being observed. Consider a point in the scene being viewed by one of the camera's pixels, with a distance R from the spacecraft, as is shown below in Figure 6.13. The polar angle of the point in the scene is related to its perpendicular ($R_\perp$) and parallel (l') distances from EnVisS by:

$$\phi = \tan^{-1}\frac{R_\perp}{l'} \tag{6.8}$$

The distances l and l' are measured along the same axis, but from different reference points. Therefore $dl = dl'$, and thus it follows from equation 6.8 that

$$\frac{d\phi}{dl} = \frac{d\phi}{dl'} = \frac{-R_\perp}{{R_\perp}^2 + l'^2} \tag{6.9}$$

From Figure 6.13 it is clear that

$$R_\perp = \sqrt{R^2 - l'^2} \tag{6.10}$$

and

$$l' = R\cos\phi \tag{6.11}$$

Substituting equations 6.10 and 6.11 into equation 6.9 yields

$$\frac{d\phi}{dl} = \frac{-\sqrt{1-\cos^2\phi}}{R} \tag{6.12}$$

Equations 6.2, 6.7 and 6.12 therefore combine to give the velocity-limited exposure time:

$$\tau_{vel} = \frac{-R}{N_{pix}\, v\sqrt{1-\cos^2\phi}\left|\frac{d\vec{\tilde{p}}}{d\phi}\right|} \tag{6.13}$$

from which the velocity-limited exposure factor $\tilde{\tau}_{vel}$ follows:

$$\tilde{\tau}_{vel} = \frac{N_{pix}\, v}{R}\tau_{vel} = \frac{-1}{\sqrt{1-\cos^2\phi}\left|\frac{d\vec{\tilde{p}}}{d\phi}\right|} \tag{6.14}$$

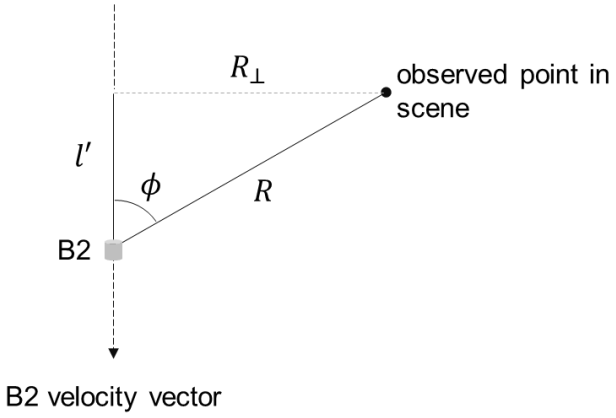

Figure 6.13: Illustration of geometry used in equations 6.8-6.12.

The velocity-limited exposure factor is plotted over the whole focal plane of EnVisS in Figure 6.14. An absolute exposure time for a particular system is obtained by dividing the exposure factor by both N_{pix} and the ratio v/R.

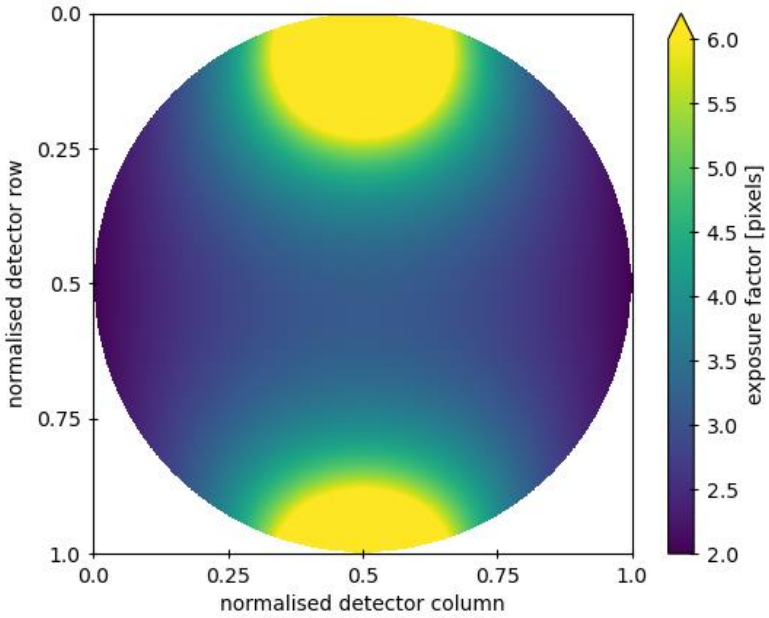

Figure 6.14: Velocity-limited exposure factor (see equation 6.14) for a non-spinning EnVisS with constant flyby velocity. Dividing by both N_{pix} and v/R ratio gives the corresponding velocity limited exposure time.

Figure 6.14 shows that, as with the spinning case, blur-free velocity-limited exposure times are largest near the poles of the spin axis. This is because this is where points in the scene have the smallest angular velocity as viewed by EnVisS. Minimum exposure times occur at the left and right edges of the image, in the unused portions of the FOV. Table 6.5 presents some actual velocity-limited exposure times for a range of N_{pix} and v/R values. The flyby velocity, v will not be known until the target comet is detected, long after the finalisation of EnVisS' design. Based on the expected delta-v capabilities of Comet Interceptor, the distribution of LPC trajectories, and target selection requirements, it is expected to be between 10 kms^{-1} and 80 kms^{-1} (Snodgrass & Jones, 2019). The distance, R from EnVisS to the scene it is imaging will depend on the shape of the scene (i.e. the 3D structure and size of the comet) and can only be approximated. Indeed, a further complication is that EnVisS will be viewing extended, optically thin structures, and each pixel's signal will not derive from a single point in space but an extended volume. However, here v/R ratio will be approximated by taking R to be the distance between B2 and the comet nucleus. At the beginning of EnVisS' operation, the comet will be millions of km from B2, whilst its closest approach to the nucleus will be on the order of 400 km. Therefore, v/R ratio is likely to have a minimum value as low as 1x10^{-5} s^{-1} and a maximum $\leq$ 0.2 s^{-1}.

Table 6.5: Absolute velocity-limited exposure times for an example set of v/R ratios and N_{pix} values.

v/R [s^{-1}]	N_{pix}[pixels]	τ_{vel} [ms]
1	512	5.9
	1024	2.9
	2048	1.5
0.1	512	59
	1024	29
	2048	15
0.01	512	590
	1024	290
	2048	150

The spin-limited exposure times listed in Table 6.4 are significantly lower (up to two orders of magnitude) than the velocity-limited exposure times of Table 6.5, even for the extreme case of $v/R = 1$ s^{-1}, which exceeds the highest v/R ratio that EnVisS is expected to observe. The spacecraft's spin motion is clearly the main limiting factor in

EnVisS exposure times, but the linear combination of the two sources of scene motion is strictly what dictates the maximum possible exposure time:

$$\frac{d\vec{\tilde{p}}}{dt} = \frac{d\vec{\tilde{p}}}{d\theta_{enviss}} \cdot \frac{d\theta_{enviss}}{dt} + \frac{d\vec{\tilde{p}}}{d\phi} \cdot \frac{d\phi}{dt} \tag{6.15}$$

Given the significant change in viewing geometry that will occur over the course of B2's flyby, the relative contributions of spacecraft spin and flyby velocity to the rate of scene motion across the focal plane will vary as the flyby progresses. The maximum exposure time (calculated with equations 6.2 and 6.15) as a function of v/R ratio and B2 spin period is plotted in Figure 6.15.

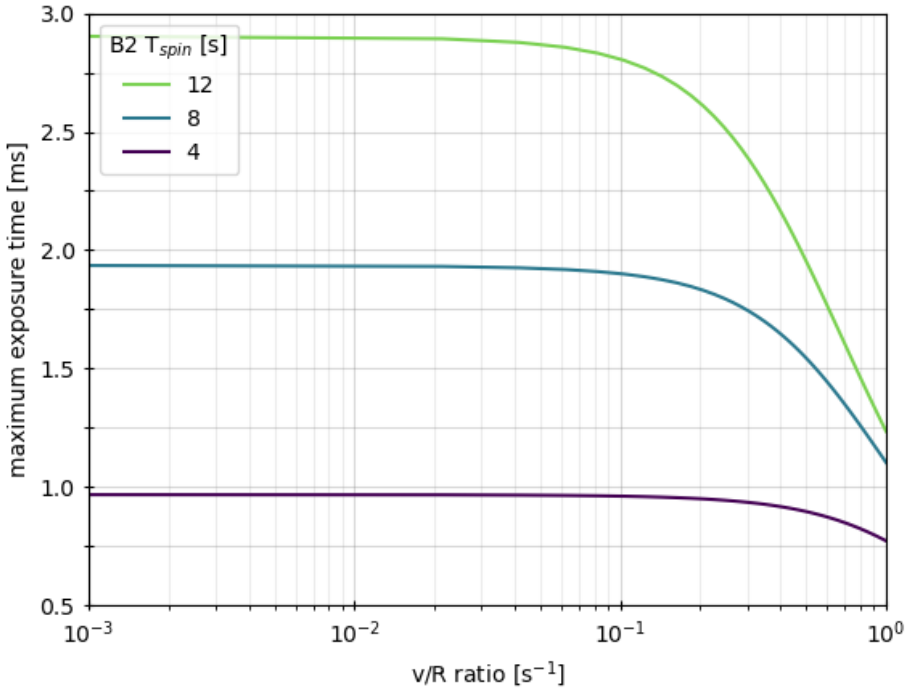

Figure 6.15: EnVisS' maximum blur-free exposure time (with $N_{pix} = 2048$) for three possible spin periods.

It can be seen from Figure 6.15 that the effect of the spacecraft's flyby motion on exposure time only becomes significant at $v/R \gtrsim 0.1$ s^{-1}. With a maximum expected likely v/R ratio of ~0.2 s^{-1}, this indicates that by using an exposure time based on τ_{spin} (equation 6.4), significant pixel smear due to B2's flyby motion is unlikely to arise in any of EnVisS' images.

6.2.2.2 Exposure Frequency

EnVisS will produce all-sky scans by repeatedly capturing framelets – small images covering a narrow strip of the sky – as the spacecraft spins. The combination of these framelets into a single product yields the all-sky image. The imaging geometry of these framelets dictates the frequency with which they must be captured in order to cover the entire visible sky in a single 360° rotation of the spacecraft. As is discussed in section 6.2.1, the exact imaging geometry of a framelet depends on which columns of the detector are used to capture it. However, the principle by which the framelet's imaging geometry relates to the minimum imaging frequency is the same for all filters.

As Figure 6.11 illustrates, the along-track extents (the range of azimuths instantaneously viewed) of EnVisS' filters vary significantly as functions of the viewed polar angle. Consider a single of EnVisS' filters imaging the sky. If the framelets it captures are to combine to give full coverage of the visible sky, there must be no gaps between the footprints of two consecutive framelets. To achieve this, the spacecraft must rotate no farther than the filter's minimum along-track extent between consecutive exposures. Figure 6.16 illustrates this. The angle by which a filter's footprint moves between the capture of two consecutive framelets is dictated by the spacecraft spin period and the sampling frequency. The sampling frequency must be sufficiently high to eliminate the green region of Figure 6.16.

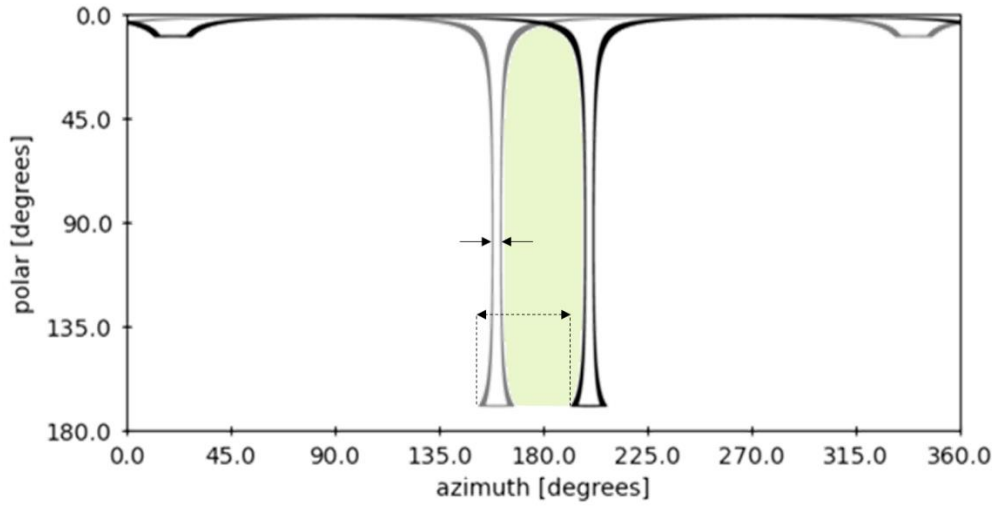

Figure 6.16: Two footprints (grey and black solid outlines respectively) of the same filter at two different times. The dashed arrow shows the angle by which the spacecraft rotated between the capture of the two footprints. The solid arrows indicate the minimum along-track extent of the filter. The separation of the two footprints is greater than the minimum along-track extent, meaning that a portion of sky between the two footprints (green) is not imaged, despite being within the filter's FOR.

With a spin period of T_{spin} and a filter footprint whose minimum along-track extent is Θ, the minimum framelet sampling frequency required in order to image 100% of the filter's FOR is

$$f_0 = \frac{2\pi}{\Theta\, T_{spin}} \tag{6.16}$$

As an example, for a provisional filter width of 5° in the along-track direction (at its narrowest point) and a spacecraft spin period of 4 s, a minimum framelet sampling frequency of 18 Hz is required. This minimum sampling frequency must be below the maximum readout frequency of the camera's detector and data handling electronics. These are themselves functions of Θ. Widening the along-track extent of a filter means imaging with a larger number of detector columns. This increases the time taken to read out the framelet and thus reduces the maximum imaging frequency the camera can achieve. Likewise, imaging with multiple filters simultaneously reduces the maximum imaging frequency for the same reason. The exact relationship between the number of columns being used for imaging and the maximum detector readout rate is dependent on the specific hardware, and at the time of writing, a characterisation of the camera hardware's performance in that regard is yet to be carried out. However, to maximise the

possible sampling frequency, it is already established that EnVisS will only perform imaging with a single filter at any one time.

Note that exposure time is dictated by pixel angular resolution and independent of filter width (every captured framelet should be blur free, regardless of when and how often they're captured), whilst sampling frequency is dictated by filter width, and independent of pixel angular resolution (framelet footprints should conjoin, regardless of their pixels' angular resolution).

6.2.3 Simulated Data: Form and Acquisition

A single framelet, captured from one of EnVisS' filters, represents a very small portion of the whole sky. Figure 6.17 (a) shows an example framelet, captured through a 50 pixel wide filter centred on the principal point of the optics. Integrating an all-sky scan requires repeatedly capturing these framelets over the course of a complete 360° spacecraft rotation. Figure 6.17 (b and c) shows two concatenated sets of 82 framelets (the minimum required to achieve full coverage with the 50 pixel wide filter) captured during an integration. These framelet sets show the same scene, but illustrate the difference in the arrangement of the data when spin direction is reversed. Although it seems trivial, it is important that the camera's data processing software is setup for the correct polarity (e.g. there are no sign errors) otherwise data will be irreversibly corrupted (e.g. Hansen et al. (2017)).

It should be noted that the simulated images of Figure 6.17 provide an idealised illustration of EnVisS' view from within a comet's coma, with the inclusion of cometary tails, dust emission features and stars in order to provide detail to demonstrate the form of the data. It's important to note that the absolute and relative brightnesses of the image constituents are not physically representative (the Sun and its effects are also omitted, but solar signal will lead to saturation of some pixels). For radiometrically-accurate simulated images, see the discussion of signal and noise in EnVisS data in section 6.4.

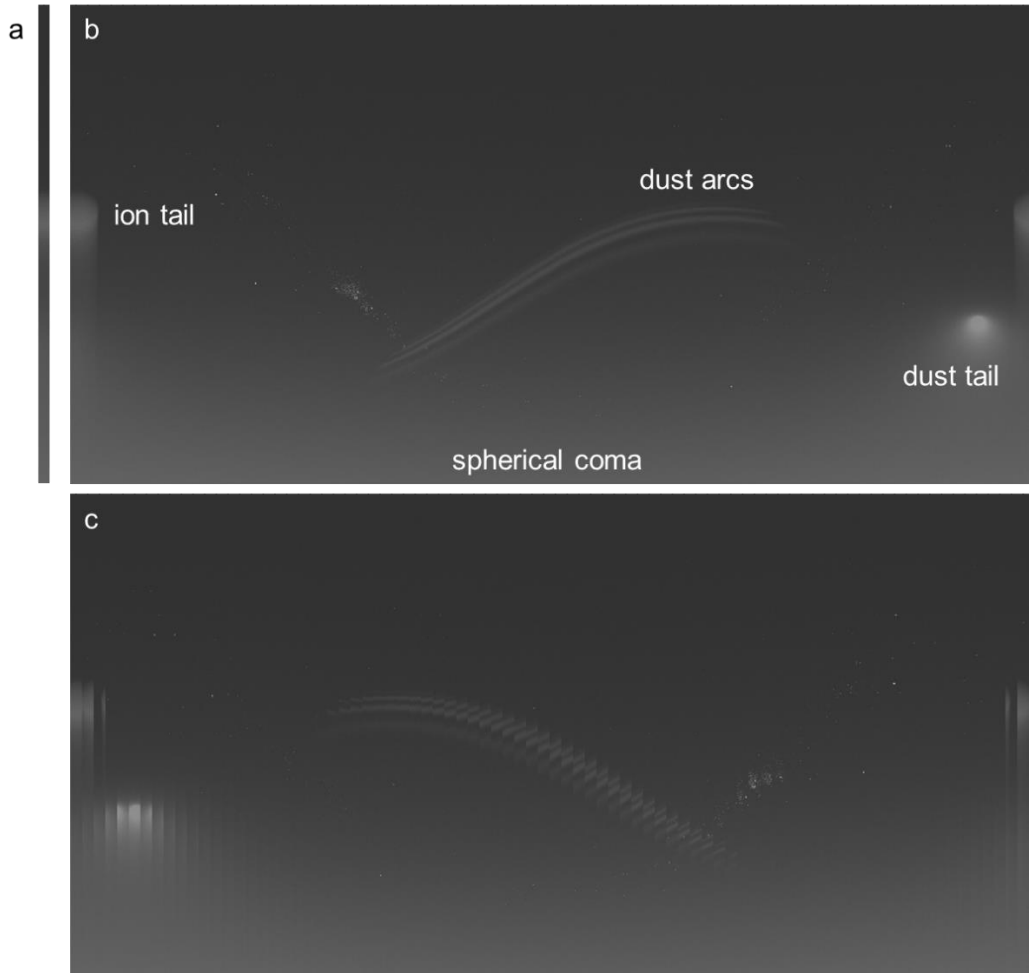

Figure 6.17: (a) a single framelet captured with EnVisS' central filter. (b) a full set of 82 concatenated framelets from a single clockwise 360° rotation. (c) the same as (b), but for an anti-clockwise rotation. The major cometary features represented within the simulated images are labelled in panel b.

With an unlimited telemetry budget, the raw framelets could be transmitted in their entirety to Earth, where their processing would include their stitching together to produce all-sky images. With EnVisS' limited telemetry allocation, this is not possible. The alternative is to construct the all-sky images onboard by mapping the framelets to a single projection of the whole sky, and discarding the original framelets.

Figure 6.18 shows a cylindrical all-sky projection, in B2-centric coordinates, of Figure 6.17's example data. The top panel displays a single reprojected framelet, whilst the bottom shows the reprojection of the full set of framelets, achieving maximum possible coverage of the sky. Note that these data represent the tilted mounting configuration of EnVisS, and there is therefore a region of no data at polar angles >170° due to the camera's blind spot. The reprojection process is the same for the untilted mounting, and

the data would be very similar except for the lack of blind spot, resulting in the nucleus and inner coma of the comet also being visible.

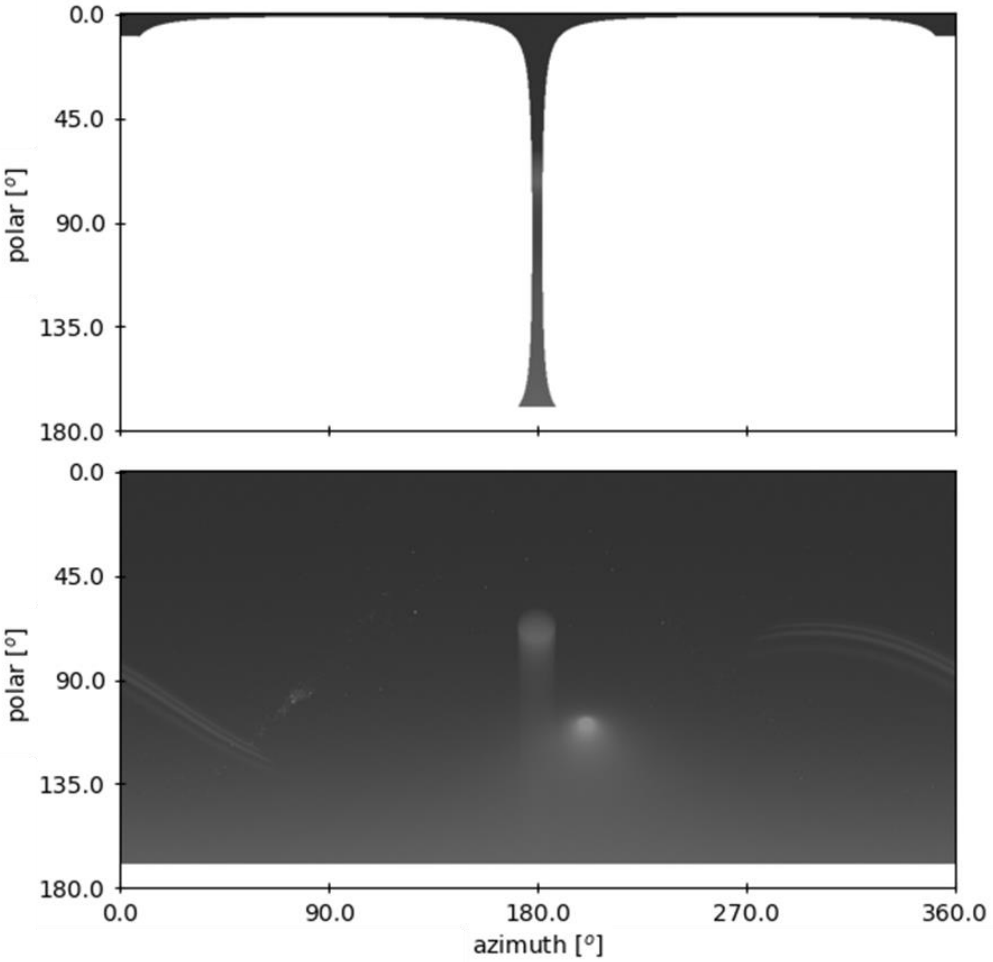

Figure 6.18: cylindrical all-sky projection (B2-centric coordinates) of example EnVisS data (for tilted EnVisS mounting). Top panel displays a single reprojected framelet (from Figure 6.17 (a)), bottom panel displays reprojection of full framelet set (from Figure 6.17 (b) and (c)).

The mapping of raw framelet pixels to pixels in all-sky images is not one-to-one, and is not uniform, as is evident from Figure 6.18 (top). Pixels at the B2-centric equator map to single pixels in the cylindrical projection, but pixels at and near the spin axis map to a significantly larger number, due to the azimuthal stretching of the framelet. This is shown in Figure 6.19 (a), which plots the number of cylindrical projection pixels to which raw framelet pixels map, for a central portion of EnVisS' detector. Additionally, many pixels in the cylindrical all-sky projection are written to more than once, because the region of sky they represent is imaged by multiple framelets. This is the case for a significant portion of the cylindrical projection, as shown in Figure 6.19 (b).

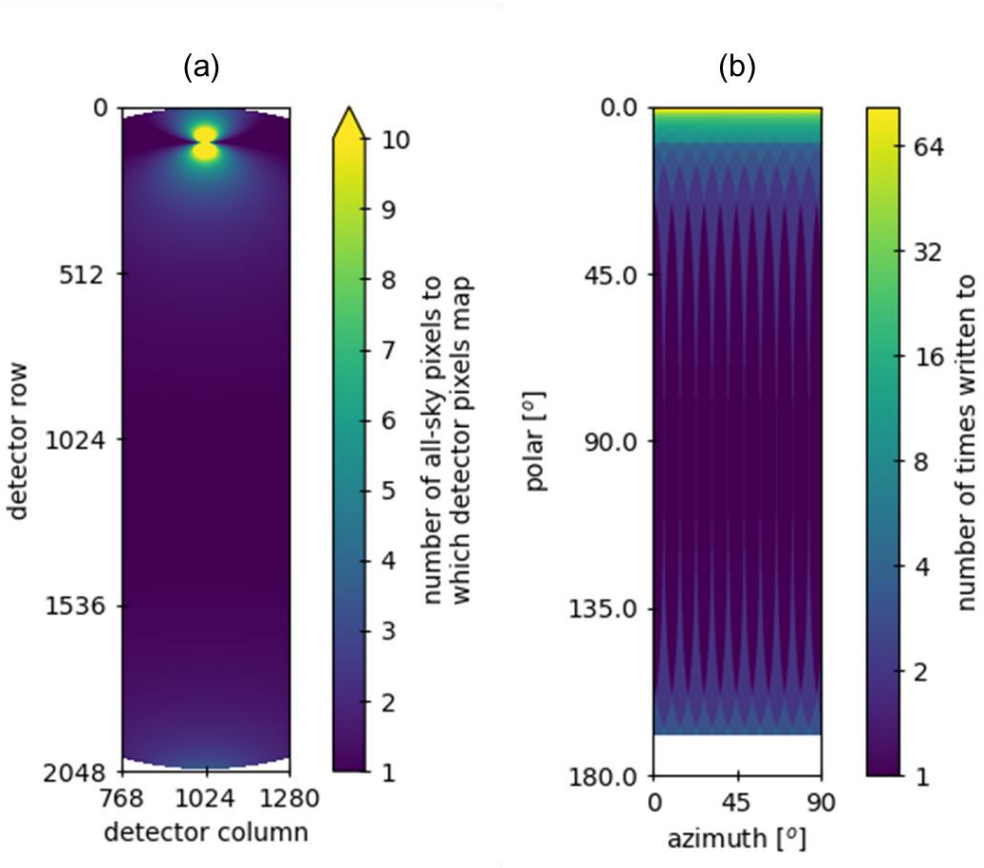

Figure 6.19: (a) the approximate number of cylindrical projection pixels to which each of EnVisS' pixels maps, for a central portion of the detector. (b) the approximate number of times cylindrical projection pixels are written to when projecting a full rotation's set of framelets to a single all-sky projection, assuming minimum imaging frequency.

The cylindrical reprojection of framelets involves both one-to-many and many-to-one mapping. The many-to-one mapping displayed in Figure 6.19 (b) means that the number of observations represented by each pixel of a cylindrical projection is not uniform. The exact numbers of how many times pixels are written to depends on the imaging frequency. A higher imaging frequency than the minimum will yield higher numbers than are plotted in Figure 6.19 (b), and this is discussed in section 6.4 in the context of its impact on signal and noise.

The mapping of multiple raw framelet pixels to a single cylindrical projection pixel is the mechanism by which the reprojection reduces the data volume, and this reduction is a major motivation for performing it prior to transmission. However, the mapping of a single framelet pixel to multiple cylindrical projection pixels has the opposite effect, increasing the data volume required to represent that framelet. To avoid the one-to-many mapping

of the cylindrical projection, one option is to employ an equal-area all-sky projection, rather than equirectangular. With equal-area mapping, the number of pixels required to represent a region of sky is proportional to the region's solid angle.

Many equal-area projections exist, used mostly for world maps, and each projection has its own characteristics and merits. Whether an equal-area projection will be employed in the EnVisS data handling pipeline, and if so which one, is a matter for the future development of the instrument, but an example is presented here to demonstrate its distinction from the equirectangular representation. The Mollweide equal-area projection is often used for representing astronomical data covering the whole sky (e.g. Argudo-Fernandez et al. (2015), Richard Gott III et al. (2005)), as well as global planetary surfaces (e.g. Běhounkova et al. (2017), Hirt et al. (2012)). Figure 6.20 gives an illustrative example of a Mollweide representation of an EnVisS all-sky scan, showing the same data as Figure 6.18.

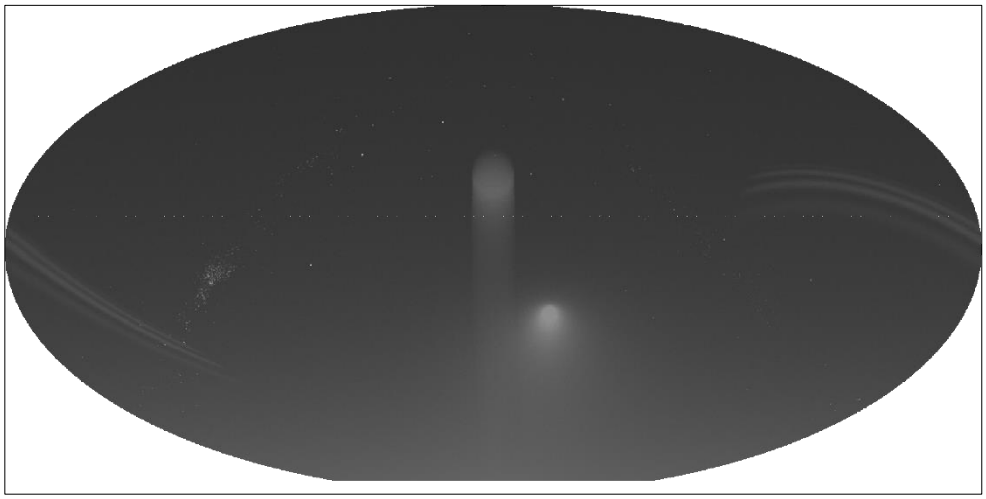

Figure 6.20: Mollweide projection of simulated EnVisS data from Figure 6.17 and Figure 6.18.

A full resolution cylindrical all-sky EnVisS image contains 8.4 Mpix, whilst a full resolution Mollweide representation of the same data contains 6.5 Mpix. A set of 82 raw framelets at full resolution also contains 8.4 Mpix, so its reprojection to cylindrical coordinates provides no data volume reduction. If minimum imaging frequency is used, no onboard reprojection should be performed, and the raw framelets should be transmitted (this mitigates the impact of changes in the spacecraft's spin axis or period during integration). However, if framelet sampling frequency is increased, the number of framelets grows, and the data-saving benefit of employing the equirectangular projection rather than keeping raw framelets becomes more significant.

Because of the non-uniform mapping of EnVisS framelets, implementing the projection of raw framelets to all-sky coordinates in EnVisS' data handling will be a non-trivial task. Some level of assumptions will have to made be regarding the stability of the spacecraft's spin axis and period, which if not satisfied, could lead to misalignment and inaccurate projection of the images. This is not a subject for this thesis, but will no doubt be an important part of the instrument's future development.

6.3 Signal and Noise in EnVisS Data

Section 6.2.2 describes the restrictions imposed on EnVisS' exposure time by spacecraft B2's spin motion, and the very short exposure times that are required to avoid motion blur within its images. An important implication of this is that the amount of signal that the instrument will be able to collect will be very small, and will impact the scientific viability of the instrument's objectives. This section characterises the signal to noise ratio (SNR) of EnVisS' observations, and the following section (6.4) analyses a number of methods by which low SNRs can be mitigated.

6.3.1 Comparison with Previous Cometary Cameras

A good place to begin with assessing the signal levels EnVisS will measure is its comparison with other cameras that have previously imaged comets during spacecraft encounters. The brightnesses of comets vary in many regards, and given that Comet Interceptor's target will not be selected until after its launch, there is a wide range of possible scene brightnesses that EnVisS could observe. However, the previously flown cometary cameras discussed in section 4.4, all of which successfully performed imaging of cometary environments, provide benchmarks with which EnVisS can be quantitatively compared.

6.3.1.1 Relative Instantaneous Signal

The amount of light measured by a camera is often represented by its f-number. A larger entrance pupil allows more light to the detector, and signal is proportional to the square of aperture diameter (i.e. proportional to aperture area). The larger the image on the detector, the more spread out the energy is, meaning signal is also inversely proportional to the square of the camera's focal length (i.e. proportional to the area on the detector over which the signal is distributed). The ratio of focal length to pupil diameter therefore provides a measure of the light captured by the system, and is called the f-number (with larger f-numbers meaning a lower signal is collected) (Smith, 2000).

The f-number's convenience comes from its ability to directly compare the image brightness of optics of different focal lengths and aperture sizes, hence its widespread use and popularity in photography, both scientific and amateur. However, when considering the signal in a single pixel of an array detector as we are here, the f-number does not account for the portioning of the total signal amongst pixels. If a camera's detector is swapped for one with the same physical dimensions, but with half as many pixels, its f-number remains the same. However, the pixels of the camera are now collecting twice as much signal (because they each occupy twice as much area on the focal plane) and pixel SNRs will be larger (assuming the two detectors have the same intrinsic noise characteristics).

Because we're interested in the SNRs of individual pixel measurements, we'll here instead characterise signal collection using the camera's aperture area and the solid angle sampled by the pixels. If pixels in two different cameras (camera 1 and camera 2) view the same extended source of uniform surface brightness L, the signal (or counts) N measured by each conforms to

$$N_i \propto a_i \, \sigma_i \, \tau_i \, L \tag{6.17}$$

where $i = \{1, 2\}$ denotes the camera, and a, σ and τ are the aperture area, solid angle sampled by the pixel, and exposure time respectively. Note that signal, N is a function of many more variables than just these (see chapter 7), but if it is assumed that the other variables are the same for both cameras it can be said that

$$\frac{N_1}{N_2} = \frac{a_1 \, \sigma_1 \, \tau_1}{a_2 \, \sigma_2 \, \tau_2} \tag{6.18}$$

If the two cameras accumulate their signal for the same duration $\tau = \tau_1 = \tau_2$, then the ratio of equation 6.18 becomes

$$\frac{N_1}{N_2} = \frac{a_1 \, \sigma_1}{a_2 \, \sigma_2} \tag{6.19}$$

The right hand side of equation 6.19 will here be called the *relative instantaneous signal* (RIS) of the cameras. The RIS gives the relative signal accumulated per unit time per pixel by two different cameras, assuming they're imaging an extended source (larger than the pixel's extent) of identical brightness.

Table 6.6 lists the pixel instantaneous FOV (IFOV) and aperture diameter of seven cameras, from five different spacecraft, all of which imaged comets during some form of encounter (both flybys and escorts). For comparison, the same parameters are stated for EnVisS. Table 6.7 lists the RIS of EnVisS with respect to each of these seven

cameras, as well as the ratios of EnVisS' pixel solid angle and aperture area to those of the listed instruments.

Table 6.6: List of parameters for cometary imagers.

Instrument	IFOV μ rad	Aperture diameter mm	See
EnVisS	1533	1.1	Pernechele et al. (2019)
OSIRIS NAC	18.6	90	Keller, et al. (2007) Naletto, et al. (2002)
OSIRIS WAC	101	25	
Deep Impact HRI	2	300	Hampton, et al. (2005) Farnham, et al. (2007)
Deep Impact MRI	10	120	
Giotto HMC	22.4	130	Keller, et al. (1987) Keller, et al. (1986)
Stardust Imaging Camera	59.4	58	Newburn, et al. (2003)
Deep Space 1 MICAS	13	100	Beauchamp, et al. (2000) Soderblom, et al. (2002)

Table 6.7: RIS of EnVisS with respect to the instruments of Table 6.6.

Instrument	σ_{enviss}/σ_i	a_{enviss}/a_i	EnVisS Relative Instantaneous Signal
OSIRIS NAC	6.79×10^{4}	1.49×10^{-4}	1.0
OSIRIS WAC	230	1.94×10^{-3}	0.45
Deep Impact HRI	5.88×10^{5}	1.34×10^{-5}	7.9
Deep Impact MRI	2.35×10^{4}	8.40×10^{-5}	2.0
Giotto HMC	4680	7.16×10^{-5}	0.34
Stardust Imaging Camera	666	3.60×10^{-4}	0.24
Deep Space 1 MICAS	1.39×10^{4}	1.21×10^{-4}	1.7

Note from Table 6.7 that EnVisS' RISs have quite a large range, and for several of the reference instruments they are on the order of 1. EnVisS' aperture area is significantly

less than all seven instruments. This is no surprise, given that EnVisS is a compact, low mass instrument on a small spacecraft, in contrast to the large cameras of Table 6.6. Conversely, the solid angle sampled by each of EnVisS' pixels is significantly larger than that of the reference instruments' pixels. Again this is to be expected. The reference instruments had narrow fields of view, designed to resolve nucleus surface features, whereas EnVisS, with a similar number of pixels to the reference instruments, will cover a significantly larger field of view in order to capture the whole sky. The net result is that EnVisS' instantaneous light capture is fairly close to that of the reference instruments, differing by no more than an order of magnitude in either direction.

6.3.1.2 Exposure Time

From the discussion of section 6.3.1.1, one might assume that given EnVisS captures light at a similar rate to other cometary imagers, it is likely to measure similar signals in its images. However, this does not account for the significant differences in the cameras' exposure times. Six out of the seven reference instruments of Table 6.6 were frame cameras mounted on 3-axis stabilised spacecraft (the exception being Giotto's HMC). 3-axis stabilisation allowed their accurate and steady pointing such that images with long exposure times could be captured. For example, Bertini et al. (2017) describe coma observations made with numerous filters of the OSIRIS NAC and WAC, for which exposure times ranged from 80 ms to 146 s (their Table 1). Even with these exposure times, significantly larger than those possible with EnVisS, on chip stacking was employed and some filters' SNRs were found still to be too low. The Stardust Imaging Camera used exposure times of 10 ms and 100 ms during its close encounter imaging.

With HMC being mounted on the spin stabilised Giotto spacecraft, its exposure time was limited to durations comparable to EnVisS. Exposure times ranged from 14.2 μs to ~30 ms, with TDI used to increase these values by a factor of up to 8 (Keller, et al., 1987). TDI could be implemented for HMC with little to no motion blur because scene motion was close to linear across the instrument's small FOV. The fisheye FOV of EnVisS is very different, and scene motion is not linear, complicating the application of TDI (see section 6.4.3 for analysis of this).

The exposure time associated with capturing full resolution EnVisS images is therefore significantly smaller than was used by any of the reference cameras, by a factor of 10-100,000 depending on the instrument. This alone indicates that signals will be especially low in EnVisS images, but it is valuable to corroborate this with quantitative estimates of the actual signals EnVisS will measure.

6.3.2 Estimating the Expected Brightness of the Comet

6.3.2.1 Previous Observations

Whilst the brightnesses of the comet's environment and structures that EnVisS will measure cannot be precisely constrained, a likely radiance range that the instrument will observe can be inferred from data collected by the reference instruments of Table 6.7. Radiometrically calibrated images from these instruments provide a means by which to estimate what EnVisS will observe, allowing quantitative analysis of the SNRs it will achieve in its measurements. Table 6.8 displays the parameters of several observations from a selection of the reference instruments.

Table 6.8: Brightnesses of dust comae observed by cometary cameras, with accompanying observation parameters: λ - effective wavelength at which spectral radiance was calculated; $\mathbf{R_{obs}}$ - distance from comet nucleus to observer; $\mathbf{R_{\odot}}$ - distance from comet nucleus to Sun; **impact parameter** - observation impact parameter.

Comet	Instrument	λ nm	R_{obs} km	$R_{\odot}$ AU	Impact parameter km	Observed Spectral Radiance $Wm^{-2}sr^{-1}nm^{-1}$	Reference
9P/Tempel 1	Deep Impact HRI	653	5.6e6	1.5	360	$8x10^{-8}$	McLaughlin, et al. (2014)
					720	$3x10^{-8}$	
	Deep Impact MRI	610	1.4e4		20	$2x10^{-6}$	
103P/Hartley 2	Deep Impact HRI	653	4.3e4	1.0	11	$9x10^{-6}$	Lindler, et al. (2012)
		375	7e3		4.4	$2x10^{-5}$	
81P/Wild 2	Stardust NAVCAM	666	2.8e3	1.9	38	10^{-6}	Carcich (2014)
			1.1e3		3.3	$2x10^{-5}$	
67P/C.-G.	OSIRIS NAC	650	3.9e4	3.7	270	$2x10^{-6}$	Tubiana & ESA (2017)
	OSIRIS WAC	613	160	1.4	5	$2x10^{-5}$	
1P/Halley	Vega 2 TVS	~690	1.4e7	0.9	3e4	10^{-5}	Sagdeev et al. (1986a)

The observations listed in Table 6.8 demonstrate the wide range of brightnesses that comets can exhibit. Lines of sight with lower impact parameters generally observe higher radiances than those viewing the more distant coma, due to the higher density of material being observed. The cometary cameras listed in Table 6.6 and whose data are shown in Table 6.8 were designed to study the nucleus and inner coma of their target comets,

using narrow FOVs. Therefore, there are few high quality, low noise images of the distant coma from which to estimate the brightnesses to expect farther from the nucleus. But observing distant and largescale structures is a significant part of EnVisS' objectives, and these structures will constitute the majority of what the instrument sees. It is therefore important to estimate the expected signals at larger distances from the nucleus.

The data of Giotto's OPE instrument (discussed in section 4.4.6) complement the data of Table 6.8 for estimating the brightnesses EnVisS will image. OPE imaged in the opposite direction to the comet nucleus as the spacecraft approached from a large distance, and measured a coma radial brightness profile, described by Levasseur-Regourd et al. (1999). Brightnesses were measured at three different wavelengths: 441 nm, 576 nm and 718 nm, and at all three wavelengths brightnesses were well within an order of magnitude of each other. Figure 6.21 plots the approximate coma brightness measured by OPE at 718 nm.

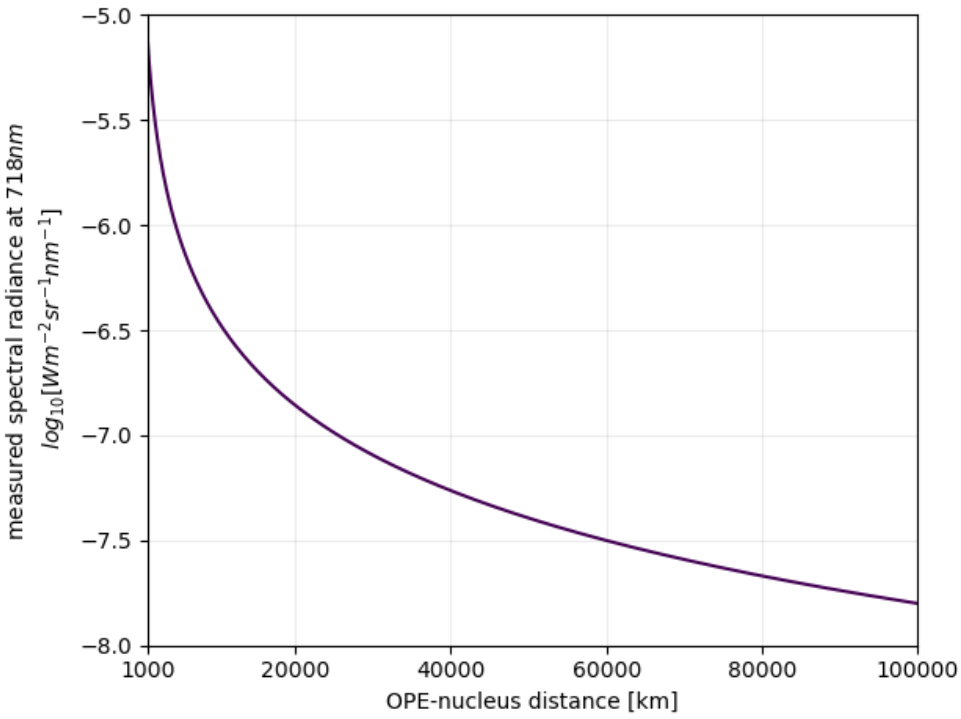

Figure 6.21: Approximate coma brightness measured by OPE at visible wavelengths, as a function of the instrument's distance from 1P/Halley's nucleus (Levasseur-Regourd, et al., 1999).

Similar coma brightnesses are indicated by Tubiana et al. (2015) and their measurements of comet 67P's dust environment. Their Figure 3 presents results broadly

indicating coma observations with impact parameters between 1000 and 10,000 km to have observed brightnesses from 10^{-6} to 10^{-7} $Wm^{-2}sr^{-1}nm^{-1}$.

The data from Table 6.8 and from OPE tell the same story, that EnVisS will likely, over the course of its entire flyby, view brightnesses ranging from 10^{-8} to 10^{-5} $Wm^{-2}sr^{-1}nm^{-1}$. However, without further qualification, this range could be interpreted overly optimistically. As displayed in Figure 6.21, the high end of this brightness range will only fill EnVisS' FOV when it is deep into the comet's coma, near closest approach. For the rest of the flyby it will occupy only a small portion of the instrument's FOR, or will even be obscured by the spacecraft's dust shield. For most of the flyby, the majority of the camera's FOV will view distant coma, and thus will see the low end of the brightness range, i.e. $\lesssim 10^{-7}$ $Wm^{-2}sr^{-1}nm^{-1}$.

6.3.2.2 Comet Interceptor's Engineering Dust Coma Model

Separately to the work for this thesis, an Engineering Dust Coma Model (EDCM) was developed by a working group within the Comet Interceptor mission for use by its hardware teams in order to provide a consistent, physically based model of the dust environment that the three spacecraft are expected to encounter (Tubiana & Marschall, personal communications, 2021). It provides dust number densities and column densities along the flyby trajectories of each spacecraft. For EnVisS, the EDCM provides another mode by which coma brightnesses can be estimated.

For spacecraft B2, the EDCM gives dust number density values along a straight trajectory, perpendicular to the Sun-comet line, with a closest approach of 200 km occurring on the day side of the nucleus. From these dust number densities along the B2 flyby trajectory, the brightness of the coma as viewed by EnVisS for a LOS along the trailing spin axis was calculated. This is the equivalent viewing geometry to that of OPE, and represents the LOS of minimum signal and maximum signal during inbound and outbound portions of the flyby respectively. The method by which the brightnesses were calculated is described in chapter 7. To reflect the uncertainty in the nature of the comet that the mission will encounter, the EDCM provides a distribution of values for each data point, allowing the 10^{th}, 50^{th} and 90^{th} brightness percentiles to be estimated. The calculated brightness profile is plotted in Figure 6.22.

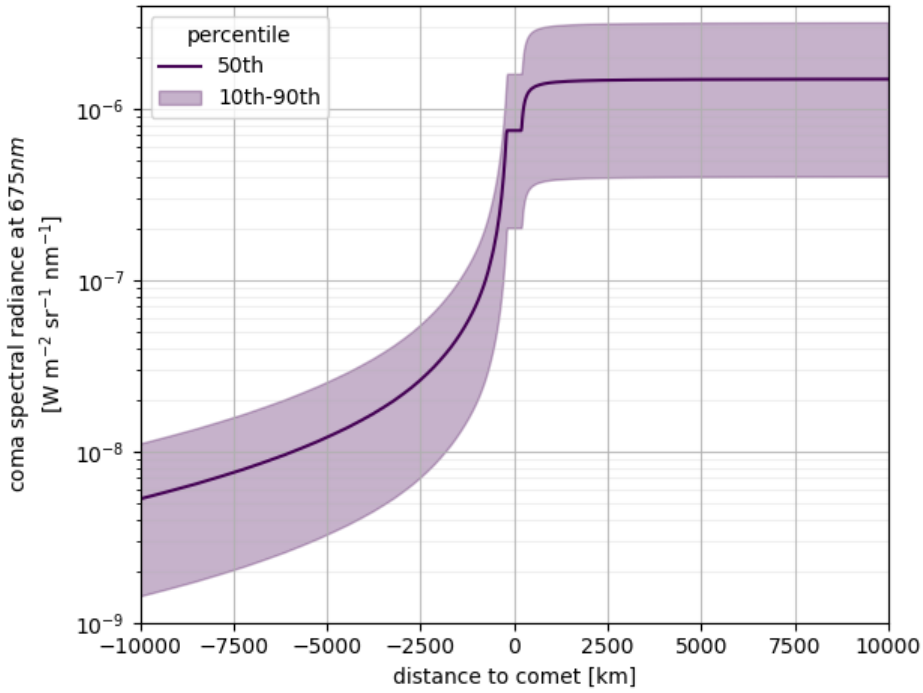

Figure 6.22: Observed spectral radiance of the dust coma in the B2 trailing spin axis direction, as a function of B2's distance to closest approached, calculated from the Comet Interceptor EDCM. The discontinuity at closest approach is a feature of the EDCM data point spacing.

The EDCM data do not cover the full spatial extent over which OPE collected measurements, but allow the estimation of the observed brightness for distances ±10,000 km from closest approach. Comparing Figure 6.22 and Figure 6.21, the overall EDCM brightness profile is lower than that measured at Halley (Halley was a relatively bright comet, so this is expected), but it is still largely within the 10^{-8} to 10^{-5} $Wm^{-2}sr^{-1}nm^{-1}$ range that is predicted by the cometary observations covered in section 6.3.2.1.

The EDCM provides dust densities along the three spacecraft flyby trajectories, but does not give a full 3D dust distribution throughout the coma. For this thesis, in order to facilitate the estimation of brightnesses viewed by EnVisS along any arbitrary LOS (something which is necessary to produce physically accurate simulations of EnVisS data), a simple, spherically symmetric dust coma model (SSDCM) was employed, whose absolute dust densities were linked to the EDCM. Specifically, the SSDCM uses the spherical dust density distribution

$$n = \frac{k}{r^2} \tag{6.20}$$

where n is the dust number density, r is the distance from the nucleus and k is a constant. To match the SSDCM to the EDCM, k was set such that the dust column density along spacecraft A's flyby trajectory was the same in both the SSDCM and the EDCM.

An example all-sky map of the spectral radiance of the SSDCM is shown in Figure 6.23. This map was generated at a viewpoint 10^5 km from closest approach, on a trajectory perpendicular to the Sun-comet line and with a closest approach of 500 km. The full simulation details are described in chapter 7.

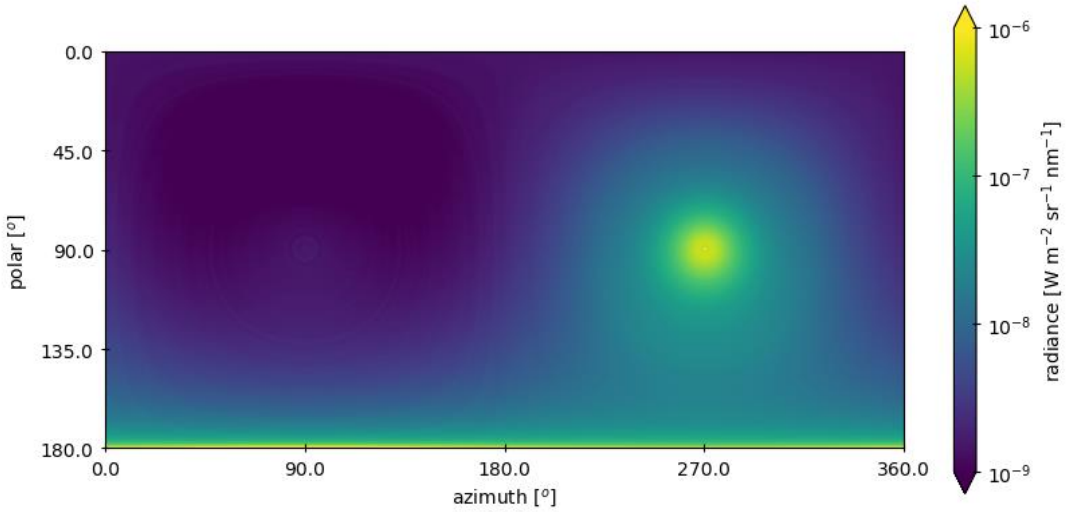

Figure 6.23: Simulated spectral radiance of a spherical dust coma over the whole sky (B2-centric coordinates). Viewpoint is 10^5 km from closest approach, with a nucleus phase angle of 90°. The Sun is at 270° azimuth. Brightness is displayed with a logarithmic scale.

Figure 6.23 represents the sky in B2-centric coordinates, so the brightest inner coma and nucleus are stretched across the bottom of the image. LOSs towards this region are relatively bright, due to the large amount of material they observe. The Sun is located at 270° azimuth, 90° polar. LOSs close to the Sun are also relatively bright, due to the dust's higher scattering efficiency at large phase angles (see chapter 7). The brightnesses predicted by the SSDCM are consistent with the EDCM and the observational data of 6.3.2.1.

6.3.3 Baseline SNR of EnVisS Images

The signal to noise ratio of EnVisS' images as a function of the observed scene brightness can be estimated by modelling the major optical, spectral and electronic characteristics of the instrument. The 'true' electron count N (i.e. without the influence of

noise) in a single pixel in response to observing a spectral radiance L with an exposure of duration τ is given by

$$N = a\,\Omega\,\tau \int L\,T_o\,T_f\,q\,\frac{\lambda}{hc}\,d\lambda \tag{6.21}$$

where a is the area of the camera's entrance pupil, Ω is the solid angle sampled by the pixel, T_o and T_f are the spectral transmissions of the optics and filter respectively, and q is the quantum efficiency of the detector. Accounting for three noise sources: shot noise, dark current and read noise (see section 2.3.1.1), the SNR of the pixel's measurement is given by

$$SNR = \frac{N}{\sqrt{N + N_d\tau + N_r}} \tag{6.22}$$

where N_d is the dark current and N_r the magnitude of the read noise.

Imaging modes whose objective is to increase the SNR of EnVisS' final images above that described by equations 6.21 and 6.22 will be discussed in the following sections, but the baseline case conforming to these equations will first be presented. This baseline case corresponds to operating EnVisS at its full resolution, as described in section 6.2.2. The parameters used for modelling EnVisS' signal are presented in Table 6.9 and the accompanying plot of Figure 6.24. It is assumed that the light observed by EnVisS has a solar spectrum, consistent with scattering and reflection of sunlight by dust particles.

Table 6.9: Parameters used for the EnVisS instrument model.

Parameter	Value	Note
Image size	2048 pixels	Across-track axis.
Pixel IFOV	1.533 mrad	Corresponds to the above image size.
Entrance pupil diameter	1.2 mm	Assumed constant over FOV.
Spacecraft spin period	4 s	Possible range: 4-15s.
Exposure time	1 ms	Linked to resolution and spin period (as in 6.2.2.1).
Optics transmission	0.8	Assumed spectrally constant.
Bandpass	~550-800 nm	See Figure 6.24.
Detector read noise	13 e^-	CMV4000 detector.
Detector dark current	125 $e^-\,s^{-1}$	CMV4000 detector.
Detector quantum efficiency	0.5 at 675 nm	See Figure 6.24.

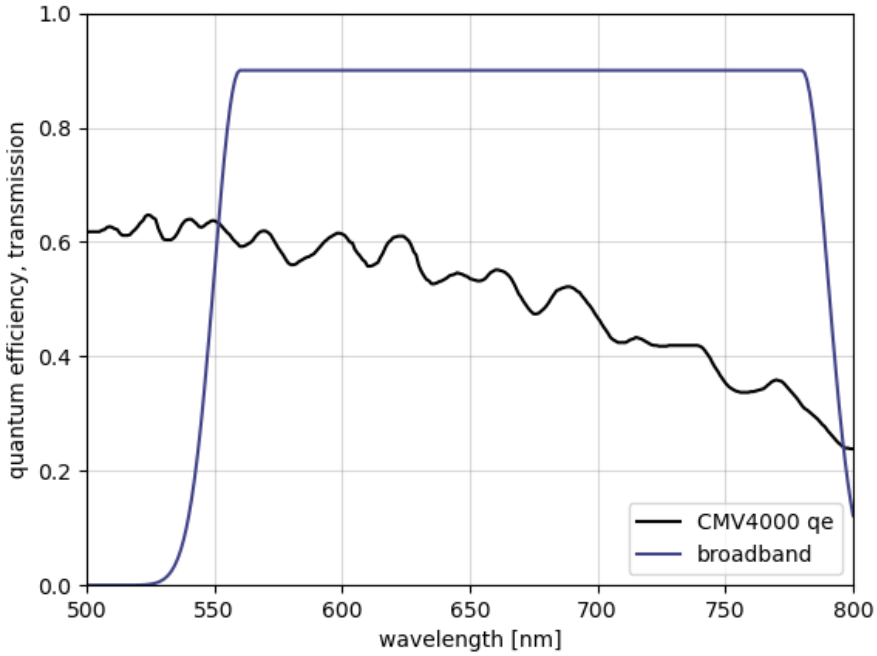

Figure 6.24: Broadband filter transmission curve and detector quantum efficiency curve used for EnVisS instrument model.

The SNR that individual EnVisS pixels will measure as a function of the radiance they observe is plotted in Figure 6.25, calculated as described above.

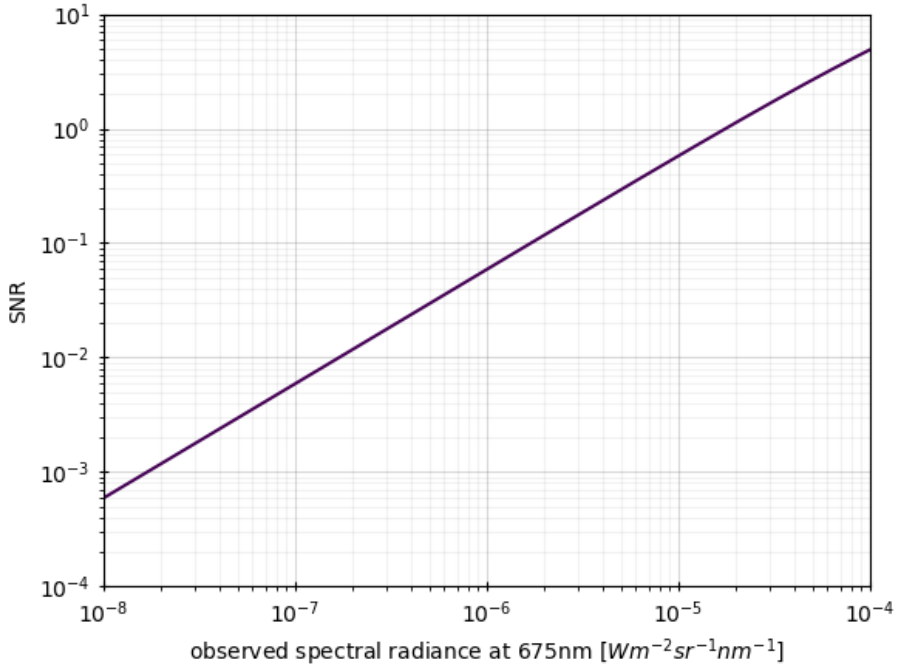

Figure 6.25: EnVisS baseline SNR as a function of observed radiance.

Figure 6.25 demonstrates that signal is a profoundly limiting factor for EnVisS. For the baseline case, the vast majority of the expected range of observed brightnesses result in an SNR significantly less than one, i.e. the signal received by the pixels is less than the noise floor, and no meaningful detection is made. This is largely caused by the very short exposure times – significantly shorter than are conventionally used for cometary coma observations – to which EnVisS is restricted. Methods to maximise the SNR and increase it well above the baseline case will be necessary if EnVisS is to obtain scientifically valuable data, and such methods are discussed in section 6.4.

6.3.4 Stellar Signal

EnVisS will utilise measurements of the Sun crossing its FOV to determine its actual spin rate (which may differ from the intended value, and could be altered by dust impacts) and provide a timing reference for image scanning throughout the B2 flyby. Measuring variations in B2's spin rate and spin axis orientation throughout the flyby and during image capture will assist with the processing of its framelets into all-sky scans, but at the time of writing, the degree to which this will be possible, and by what mechanism, is an area of uncertainty requiring further study as part of the instrument's ongoing development. Whether or not stars within the camera's captured framelets can be used

to assist with determining B2's orientation is, in any case, an important factor. Therefore, a brief estimation of this is presented here.

Because of the coarse angular resolution of EnVisS, its point spread function (PSF) is well contained within the area of a single pixel over the majority of its FOV and spectral range. Taking a conservative value, approximately 70% of the energy of a point source observed by EnVisS is focussed within a single pixel, assuming the point source is centred on that pixel (Pernechele 2020, personal communication). Identifying a star from a single pixel is challenging when noise of comparable magnitude is present. Zakharov et al. (2013) discuss, in the context of star trackers, the minimum SNR required to measure stars' positions. It is a function of instrument specifications, and the desired angular accuracy, but they state that an absolute minimum SNR of 3 is required.

From a minimum acceptable SNR, the minimum detectable point source flux density F can be approximated with an adapted version of equation 6.21:

$$F = \frac{N}{a\,\tau\,q\,T_{o,\lambda'}\,T_{f,\lambda'}\,\eta_{psf}}\,\frac{h\,c}{\Delta\lambda\,\lambda'} \tag{6.23}$$

where $T_{o,\lambda}$ and $T_{f,\lambda}$ are the transmissions of the optics and filter respectively at the effective wavelength of measurement (λ'), $\Delta\lambda$ is the spectral range over which the flux is being measured, η_{psf} is the fraction of energy focussed within the pixel and N is the electron count required to achieve the desired SNR, given by equation 6.24:

$$N = \frac{SNR^2 + \sqrt{SNR^4 + 4\left(N_d\tau + {N_r}^2\right)SNR^2}}{2} \tag{6.24}$$

For EnVisS, assuming an effective wavelength of 675 nm and a passband of 250 nm width, a minimum flux density of 1.6x10^{-10} W m^{-2} nm^{-1} is required to achieve an SNR of at least 3. Given that we are considering the detectability of stars, it is more convenient to deal with their apparent magnitude than flux density. The apparent magnitude m of a point source with flux density F_m is given by Pogson's law:

$$m = 2.5\log_{10}\frac{F_0}{F_m} \tag{6.25}$$

where F_0 is the zero-point flux (the flux density associated, by convention, with a point source of magnitude 0) at the appropriate wavelength (Heck, 1978). Zero-point fluxes are available for a wide range of standard photometric systems, one of the most applicable of which to EnVisS is the Vilnius S band, which has a central wavelength of 652 nm (Mann & von Braun, 2015). At this wavelength, the zero-point flux is 1.86x10^{-11}

W m^{-2} nm^{-1}. From the minimum required flux density of $1.6x10^{-10}$ W m^{-2} nm^{-1}, equation 6.25 gives that EnVisS has a minimum detectable point source apparent magnitude of -2.34. This is significantly brighter (~1 magnitude) than even the brightest stars, meaning that all stars are well below EnVisS' detectable limit, and will not be visible within its framelets. Indeed, aside from the Sun, only Venus and Jupiter are likely to be sufficiently bright to approach detectability. As a result, stars will not be available within EnVisS' images to provide orientation information, or to assist with the stitching or mapping of framelets into all-sky scans.

6.4 Increasing the SNR of EnVisS Data

6.4.1 Pixel Binning

A commonly employed method of increasing a camera's measured SNR is the reduction of its images' angular resolution by means of pixel binning (i.e. grouping multiple pixels' measurements to a single value, see section 2.3.1.6). Its advantage comes from the flexibility it provides to capture images with a range of resolutions and SNRs with a single instrument. As well as providing increased SNR, binning has the benefit of reducing data volume, though this is not the motivation for its consideration here. This section discusses the application of pixel binning to EnVisS data, and the expected SNRs that will be achievable through its implementation.

6.4.1.1 On-Chip Binning

Many spaceborne cameras, including cometary imagers, have utilised on-chip pixel binning to provide a flexible mechanism by which to control the signal of their images. The major advantage of on-chip binning is that it occurs prior to reading the signals out from the camera's detector, and thus minimises the readout noise in the final image, leading to maximum SNR gains. However, not all detectors support on-chip binning. It is a common functionality of spaceborne CCDs but is much less often implemented in APS detectors. At the time of writing, the capabilities of EnVisS' detector are not fully characterised, though it is anticipated that on-chip binning will not be possible with the instrument's expected design. However, on-chip binning is described here in order to contextualise the performance of alternative binning methods.

To quantify the SNR gains provided by on-chip binning, let us assume EnVisS views a scene of uniform brightness, such that all pixels would accumulate signal (electrons) at the same rate J if noise were not present. From equation 6.22, the SNR in the pixels of an un-binned image is given by

$$SNR = \frac{\tau_0 J}{\sqrt{\tau_0 (J + N_d) + {N_r}^2}} \tag{6.26}$$

where τ_0 is the exposure time used for capturing a full resolution (i.e. un-binned) image. For a binned image with binning factor n_{bin}, the signals of all pixels in an $n_{bin} \times n_{bin}$ square are combined on the detector, creating a binned pixel of size $n_{bin} \times n_{bin}$ pixels. For EnVisS, whose exposure times are linked to its imaging angular resolution (see section 6.2.2.1), the $1/n_{bin}$ reduction in resolution caused by the on-chip binning permits a larger exposure time of $n_{bin}\tau_0$. The SNR of an on-chip binned pixel, which is the sum of ${n_{bin}}^2$ un-binned pixels, is therefore given by

$$SNR_{oc} = \frac{{n_{bin}}^3 \, \tau_0 J}{\sqrt{{n_{bin}}^3 \, \tau_0 (J + N_d) + {N_r}^2}} \tag{6.27}$$

The SNR gains achieved by this binning method result not only from the merging of pixels' signals, but also from the longer exposure time. When the signal is dominated by Poisson noise, on-chip binning increases SNR by a factor of ${n_{bin}}^{3/2}$ over the un-binned case, whilst low SNRs, where read noise dominates, are increased by a factor of approximately ${n_{bin}}^3$.

6.4.1.2 Post-Readout Binning

With the option of on-chip binning not available to EnVisS, the simplest alternative is to perform a software binning of the pixels in its framelets once they have been read out from the detector. This could be done onboard spacecraft B2, prior to transmission (bringing the advantage of reduced data volume) or after images have been received on Earth. It can also be performed on select portions of images, potentially increasing SNR in darker regions whilst retaining higher spatial resolution in brighter regions. The drawback of post-readout binning is that the possible gains in SNR are (relative to on-chip binning) small. Because the binning occurs after readout, every un-binned pixel already contains its own readout noise, whose contributions add in quadrature when pixels' signals are combined. The SNR of pixels in a post-readout binned image is given by

$$SNR_{pr} = \frac{n_{bin} \, \tau_0 J}{\sqrt{\tau_0 (J + N_d) + {N_r}^2}} \tag{6.28}$$

Post-readout binning yields a fixed SNR increase of n_{bin} for all SNR domains.

6.4.1.3 EnVisS Binning

Whilst the SNR gains of on-chip binning cannot be achieved with EnVisS' expected hardware, binning can still be implemented with the instrument in a way that significantly outperforms post-readout binning. This method will here be referred to as 'EnVisS binning', and it is assumed throughout this chapter that this will be the actual method by which the EnVisS instrument will implement binning in order to mitigate low SNRs.

When the capture of a binned image (with binning factor n_{bin}) is planned, exposure time can be increased to $n_{bin}\tau_0$, as described for on-chip binning. This results in the capture and read out of full resolution framelets in which pixels are smeared, but signal is higher due to the increased exposure time. These smeared framelets are then software binned to reduce their angular resolution, remove the pixel smear, and increase the signal further. As with on-chip binning, the image's SNR benefits from both the combining of pixels' signals, and the increase in exposure time (but unlike on-chip binning, it's not possible to avoid every individual pixel obtaining its own read noise prior to binning). The resulting SNR is given by

$$SNR_{En} = \frac{{n_{bin}}^2 \, \tau_0 \, J}{\sqrt{n_{bin} \, \tau_0 (J + N_d) + {N_r}^2}} \tag{6.29}$$

Note that the only difference between on-chip and hybrid binning is the quantity of read noise, and as such, when Poisson noise dominates they both yield the same factor of ${n_{bin}}^{3/2}$ increase in SNR. In the low SNR domain, when read noise dominates, SNR is increased by a factor of approximately ${n_{bin}}^2$.

6.4.1.4 Binning Method Comparison

A comparison of the three binning methods discussed above, alongside the case of using no binning, is shown in Table 6.10.

Table 6.10: Comparison of the binning methods discussed in section 6.4.1.

Binning Method	Exposure Time	Final Pixel Signal	Final Pixel Noise	SNR Increase Factor	
				Poisson dominates	Read noise dominates
None	τ_0	$\tau_0 J$	$\sqrt{\tau_0(J+N_d)+{N_r}^2}$	-	
On-chip	$n_{bin}\,\tau_0$	${n_{bin}}^3\,\tau_0 J$	$\sqrt{{n_{bin}}^3\,\tau_0(J+N_d)+{N_r}^2}$	${n_{bin}}^{3/2}$	${n_{bin}}^3$
Post-readout	τ_0	${n_{bin}}^2\,\tau_0 J$	$n_{bin}\sqrt{\tau_0(J+N_d)+{N_r}^2}$	n_{bin}	
EnVisS	$n_{bin}\,\tau_0$	${n_{bin}}^3\,\tau_0 J$	$n_{bin}\sqrt{n_{bin}\,\tau_0(J+N_d)+{N_r}^2}$	${n_{bin}}^{3/2}$	${n_{bin}}^2$

As an example of the methods' relative performance, Figure 6.26 shows the SNR obtained with each method as a function of the un-binned SNR, for a binning factor $n_{bin} = 4$. Full resolution exposure time and noise values are as given in Table 6.9.

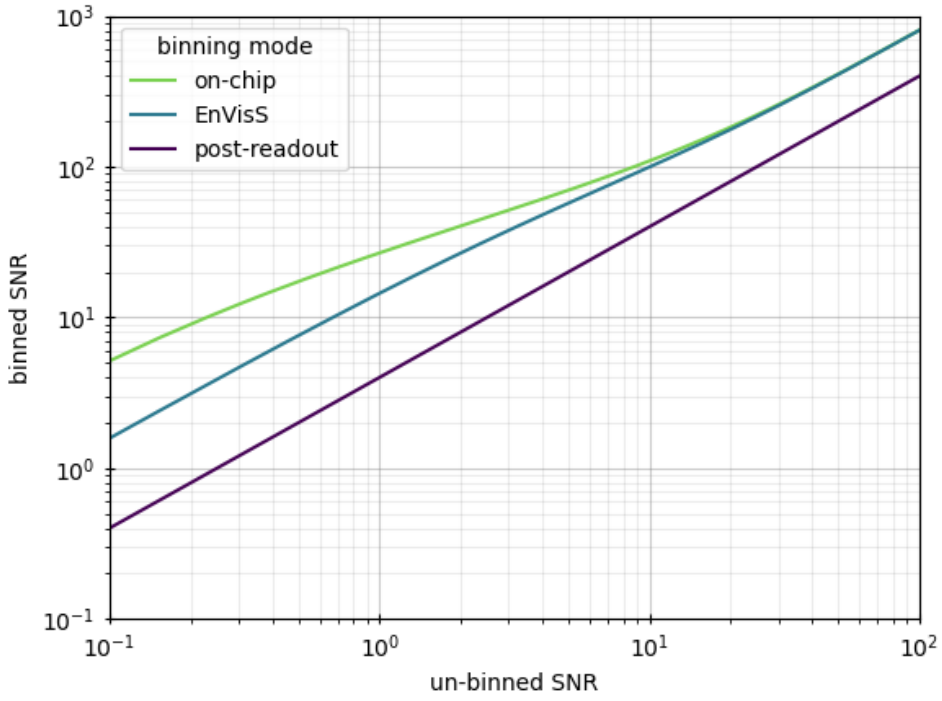

Figure 6.26: SNR of binned EnVisS pixels ($n_{bin} = 4$) as a function of the un-binned SNR for the three binning methods discussed in the main text.

Note from Figure 6.26 the convergence of on-chip and EnVisS binning as un-binned SNR increases (i.e. Poisson noise starts to dominate). EnVisS binning is significantly more effective than post-readout binning at both low and high signal levels.

6.4.1.5 Expected SNRs Implementing EnVisS Binning

The SNRs that can reasonably be achieved through the use of the EnVisS binning technique have been calculated for the range of expected radiances that EnVisS will observe. Table 6.11 details these SNRs for the lower radiance values of 10^{-8}, 10^{-7} and 10^{-6} W m^{-2} sr^{-1} nm^{-1} that will dominate the camera's FOV for the majority of the flyby. Binning up to a factor of 100 is considered, because although this significantly reduces the angular resolution from the baseline case, it still provides the opportunity to make novel imaging observations from within a cometary coma at close to a full 180° range of phase angles. Figure 6.27 plots the expected SNRs over the full range of anticipated radiances, for the same binning scenarios as are covered in Table 6.11.

Table 6.11: SNRs of EnVisS images for a range of binning factors and observed radiances L.

Binning factor	2 x 2			5 x 5			10 x 10			50 x 50			50* x 100		
L [W m^{-2} sr^{-1} nm^{-1}]	10^{-8}	10^{-7}	10^{-6}	10^{-8}	10^{-7}	10^{-6}	10^{-8}	10^{-7}	10^{-6}	10^{-8}	10^{-7}	10^{-6}	10^{-8}	10^{-7}	10^{-6}
Angle per binned pixel [°]	0.2			0.5			1			5			10		
Exposure time [ms]	2			5			10			50			100		
SNR	10^{-3}	10^{-2}	0.2	10^{-2}	0.1	1	10^{-1}	0.6	5.8	1.5	14	130	4	39	340

*Binning along the x-axis is limited by the filter's width on the detector, taken to be 50 pixels.

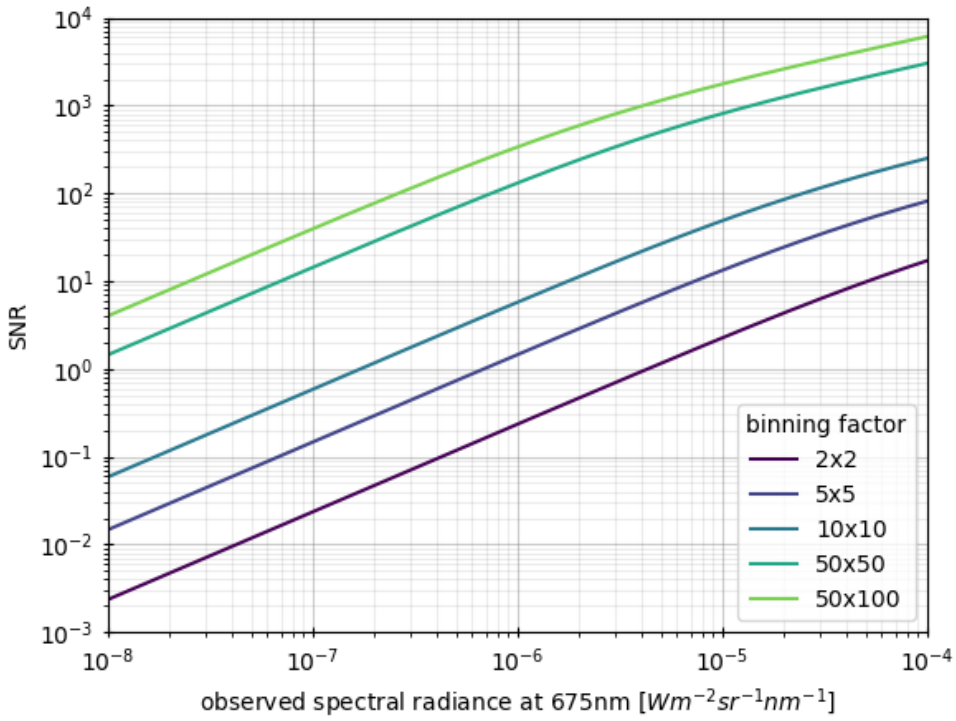

Figure 6.27: SNRs of EnVisS images for a range of binning factors and observed radiances.

Both Table 6.11 and Figure 6.27 demonstrate that EnVisS SNRs can be increased substantially (up to two orders of magnitude) from the baseline case (Figure 6.25) if large levels of binning are employed. Even in the maximum binning case considered, the bottom end of the expected brightness range is still limited to low single figure SNRs, whilst the smallest binning factors may only provide sufficient SNR for the very brightest regions of the scene.

6.4.1.6 Simulated Data

Figure 6.28 presents three examples of binned EnVisS data. These images were simulated using the instrument model described in section 6.3.3 and the SSDCM described in section 6.3.2.2. They are therefore a physically representative approximation of the data EnVisS would collect. They are displayed with a log brightness scale, where the observed radiances are derived from the images' electron counts.

The effect of binning on the SNR is very visible from comparison of the three images. The presence of noise in the 50x100 binned images is much less significant and widespread than in the 10x10 binned image. In all three images, bright dust at large phase angles and the bright inner coma can be seen. The angular resolution of the

50x100 binned image is very coarse (≈10°), but the higher SNR would help with characterising the dust's phase function over as large a range as possible. Lower levels of binning (e.g. 10x10) are limited to capturing only the brightest dust features, and are only suitable for high resolution study of the inner most coma and nucleus when the spacecraft is near to closest approach.

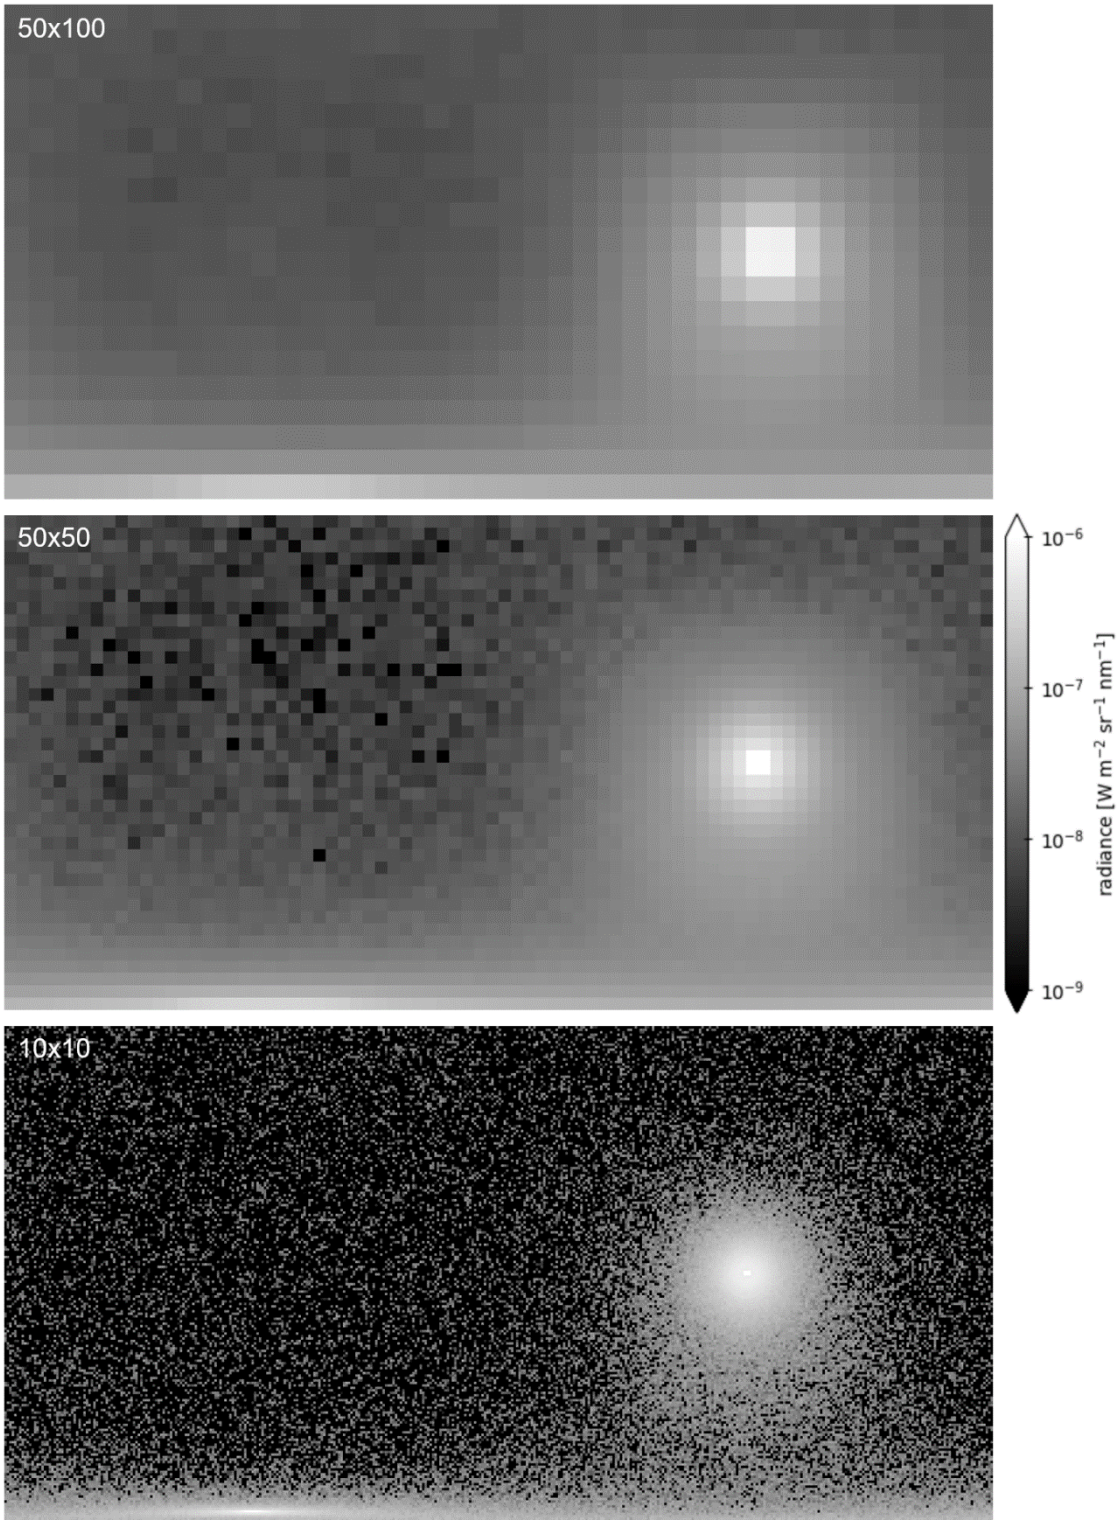

Figure 6.28: Radiometrically representative simulations of binned EnVisS data, captured 10^4 km from closest approach, with the same flyby geometry as Figure 6.23. Labels indicate the binning levels of each image (see Table 6.11 for more details). Images are displayed with a log brightness scale, derived from their electron counts.

6.4.2 Image Sequence Co-addition

Co-addition (also often called image stacking) of multiple low signal images of the same scene can increase SNR and reveal features not visible in a single image, and it is a technique used commonly in astronomical observations (e.g. Gwyn (2008); Meech et al. (2004); Moreno et al. (2016)). The repeated scanning of the sky by EnVisS will produce sequences of potentially large numbers of images which may be conducive to co-addition. Ideally, as many of EnVisS' raw data as possible would be returned to Earth for processing, including individual framelets of each sky scan. This would allow the most rigorous co-addition, as the quality of alignment between co-added frames could be manually checked and adjusted. The large data volume which EnVisS produces, and the significantly smaller data budget available for the whole of the B2 spacecraft, mean that this approach is not possible. Co-addition will likely largely have to be performed automatically onboard the spacecraft. This has the advantage of reducing the camera's data volume whilst ideally simultaneously producing higher signal images. However, it carries the inherent risk of irreversibly degrading image data if co-addition is inaccurately executed (e.g. misalignment of co-added images). This section makes an assessment of the effectiveness of co-addition for increasing SNR, by approximating the number of images that it will be possible to co-add throughout B2's flyby.

6.4.2.1 SNR Improvement

By combining images, co-addition adds together both the signal and noise contained within the pixels' measurements. If n_{coadd} measurements each with signal N are co-added, the total signal will be $n_{coadd}\,N$. The random noise N_{noise} in individual images will add in quadrature, leading to a total noise in the co-added image of $\sqrt{n_{coadd}}\,N_{noise}$. Given that the SNR of an individual image is given by

$$SNR_1 = \frac{N}{N_{noise}} \tag{6.30}$$

and that of a co-addition of n_{coadd} images is given by

$$SNR_{coadd} = \frac{n_{coadd}\,N}{\sqrt{n_{coadd}}\,N_{noise}} \tag{6.31}$$

the increase to SNR resulting from coaddition is

$$\frac{SNR_{coadd}}{SNR_1} = \sqrt{n_{coadd}} \tag{6.32}$$

The minimum number of images required to achieve a co-added SNR of SNR_{coadd} with individual images each with an SNR of SNR_1 is therefore

$$n_{req} = \left(\frac{SNR_{coadd}}{SNR_1}\right)^2 \tag{6.33}$$

Even before performing analysis on the ranges of values of n_{coadd} that will be possible with EnVisS, it is clear from equation 6.33 that SNR gains greater than one order of magnitude are unlikely to be possible by co-addition alone, as this would require hundreds of images to be combined. A major barrier to this is the constant changing position of EnVisS with respect to the comet environment. A key requirement of co-addition is that when two images are combined, all pixels whose values are added together measured signal from the same portion of the scene. Where this condition is not met, spatial resolution is degraded. The worse the correspondence between the co-added images, the more the resolution will be degraded.

6.4.2.2 Static Scene Assumption

It will be assumed here that co-addition will be carried out with individual all-sky scans, where a single all-sky scan is an already constructed image of the full sky, comprising framelets from a full rotation (i.e. a whole sky projection of the types presented in Figure 6.18 and Figure 6.20).

Consider two all-sky scans captured consecutively by EnVisS during the flyby. Assuming that the spin axis and period of B2 was the same, and unchanged, during both image integrations, the only change in their viewing perspective of the scene results from the spacecraft's flyby motion. In this case, co-addition can be achieved, in its most basic form, by simply adding the two images together. However, because the spacecraft moved between capturing the two all-sky scans, objects' positions will have shifted (along the vertical axis of e.g. Figure 6.18), and will not be in exactly the same place in the two images. This results in a smearing of image features. If the motion between the two images is suitably small (e.g. less than one pixel), smearing will be negligible.

EnVisS will, for much of its flyby, be imaging at large distances, and scene motion between consecutively captured all-sky scans will be small. To achieve significant improvement in SNR, co-addition of two images will not be sufficient, and a larger stack will be required. This leads to the question of how many images can be consecutively captured before the scene motion, relative to its position in the first image of a stack, exceeds an acceptable length for the images to be combined. This dictates a co-addition limit, and is a function of the scene's motion on the camera's detector, which will vary as

the flyby progresses. To assess this, let us rely again on the v/R ratio first outlined in section 6.2.2.1 to approximate the rate of scene motion across the detector.

Figure 6.29 (top) plots the rate of scene motion across EnVisS' detector for an example v/R value of 10^{-4} s^{-1}. Figure 6.29 (bottom) shows a map of the corresponding maximum number of images which can be consecutively captured and co-added before pixel smear exceeds 1 pixel length, as a function of pixel position on the detector. A 4 s B2 spin period is assumed.

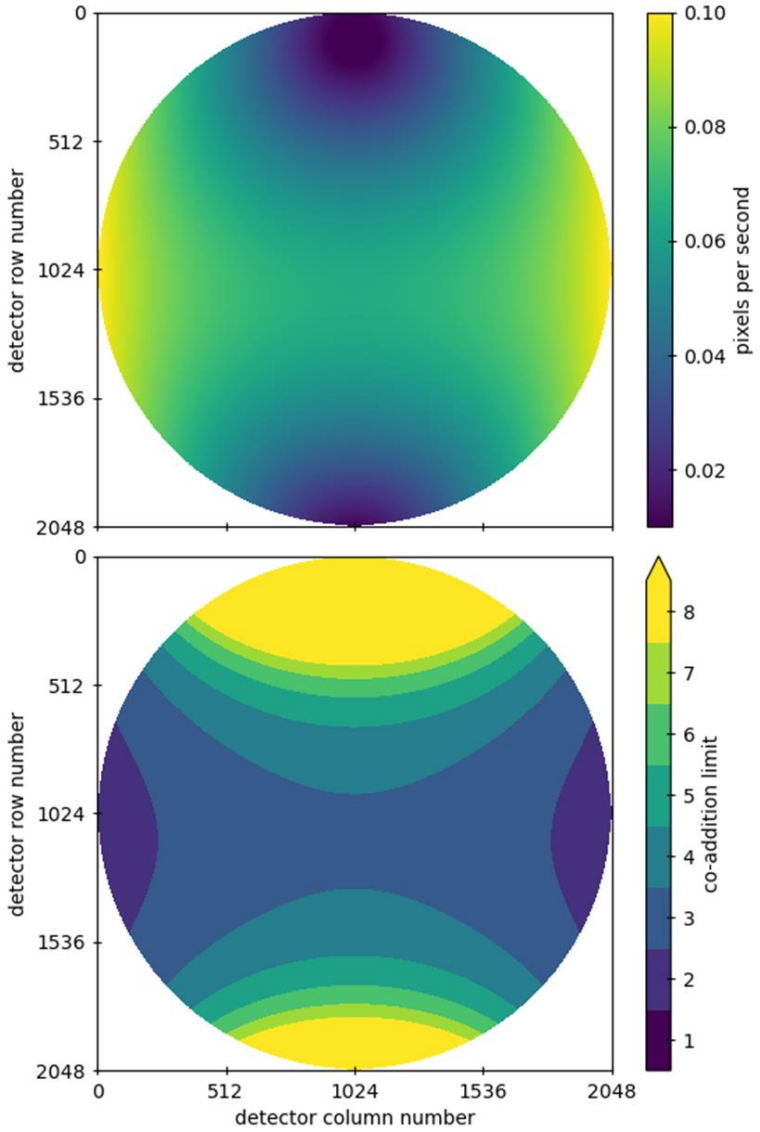

Figure 6.29: (top) Rate of scene motion across EnVisS detector for $v/R = 10^{-4}$ s^{-1}. (bottom) Corresponding maximum number of consecutive images which can be co-added without smear >1 pixel (imaging at full resolution, 4s B2 spin period).

It is clear from Figure 6.29 that when performing co-addition of frames, motion blur does not arise uniformly. To prevent motion blur from occurring anywhere within a co-added image, the number of stacked images should be dictated by the most limiting area of the imaging region (for example, imaging with columns 1000-1048 would limit co-addition to 3 images in the case of Figure 6.29). Alternatively, blurring could be accepted over a limited portion of a co-added image so as to prioritise achieving higher SNR at its polar regions (e.g. co-adding 8 frames in the case of Figure 6.29).

Because the co-addition limit is dictated by the relative size of the image pixels and image motion, it is larger for binned images. Figure 6.30 shows the co-addition limit for the range of v/R values that EnVisS is expected to encounter throughout its flyby, for a range of possible binning factors. These curves are calculated for a 4 s spin period, assuming that imaging is occurring with the camera's central filter, and by taking the most limiting value anywhere within the filter's boundary.

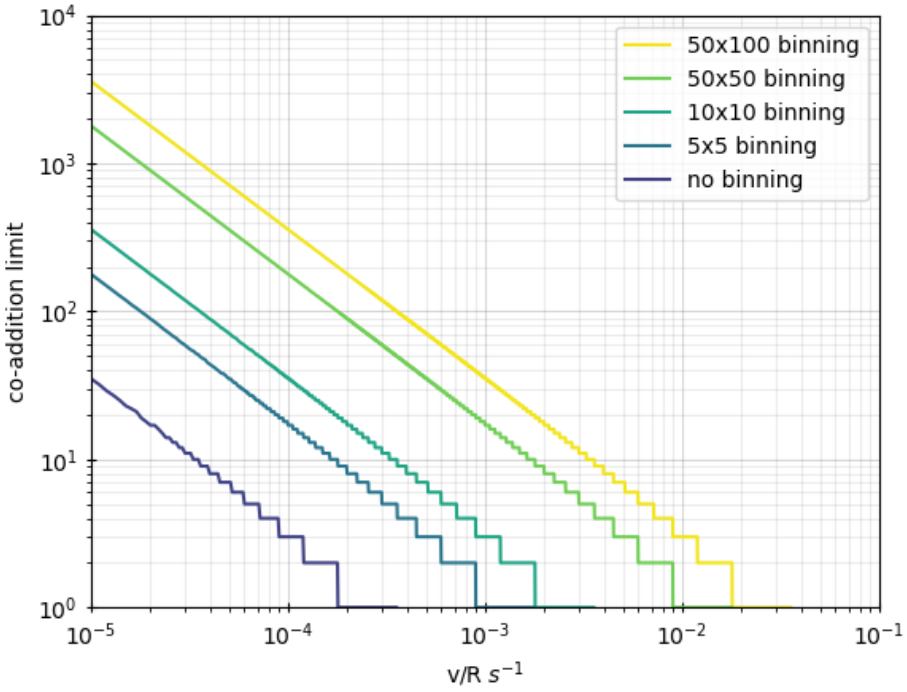

Figure 6.30: Co-addition limit as a function of flyby v/R for a number of image binning factors.

Figure 6.30 demonstrates that for all binning levels, the co-addition limit varies markedly over the expected range of v/R ratios, affecting when and by how much SNR can be increased. Figure 6.31 plots the maximum factor by which co-addition can increase SNR as a function of B2's distance to closest approach when enforcing the co-addition limit of

Figure 6.30. This assumes a constant flyby speed of 50 km s^{-1}, and approximates the v/R ratio by taking R to be the spacecraft-nucleus distance.

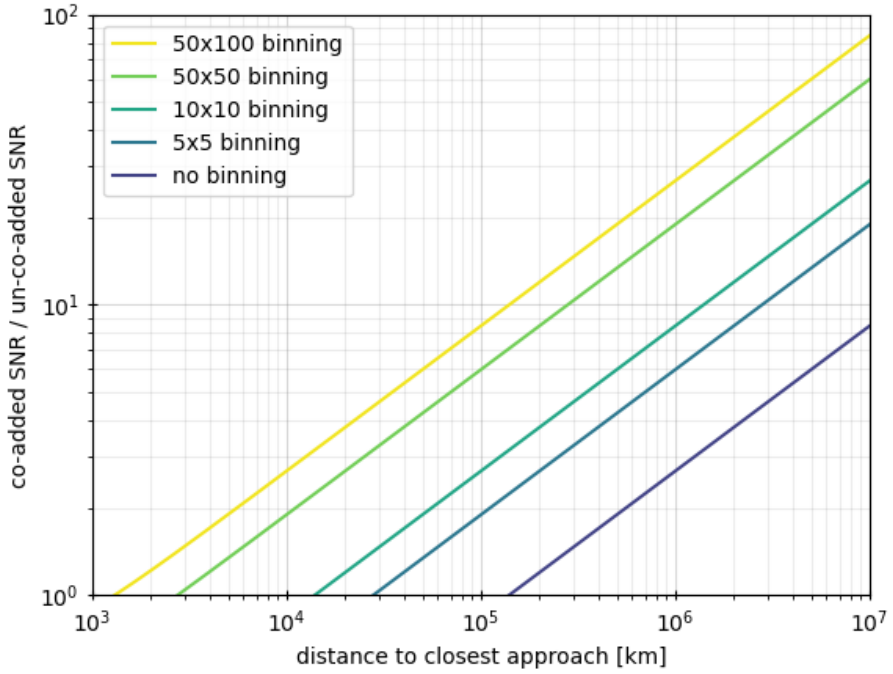

Figure 6.31: Maximum factor by which co-addition increases SNR as a function of B2 distance to closest approach (assuming a 50 km s^{-1} flyby speed), when limiting the number of co-added frames according to Figure 6.30.

From Figure 6.31, a number of general conclusions about the usefulness of co-addition can be drawn. In the early stages of the flyby, SNR could be increased by factors ranging from approximately 10-100 for a range of binning cases. Based on the expected SNRs, such improvements could contribute significantly to the quality of images with binning factors of 10 or greater. Given the very low SNRs anticipated for un-binned images (as low as 10^{-3}, see Figure 6.25), co-addition will not be sufficient to increase the un-binned image SNRs to acceptable levels, and given the small number of un-binned images that can be co-added at any stage of the flyby, there is little justification for performing co-addition of un-binned images. Within 10,000 km of closest approach, the rate of scene motion is likely to be too great for co-addition to yield any meaningful increase in SNR for any level of binning.

It is important to note that v/R has been approximated here by setting R to the spacecraft-nucleus distance, but in reality it will be a function of the viewing direction and structure of the comet's environment (after all, EnVisS' pixels will be viewing the large

scale coma, not the nucleus), and could be both smaller and larger than predicted here. A conservative estimate of v/R - limiting co-addition to a smaller number of frames - helps to mitigate misalignment and blurring but restricts its usefulness. Co-adding larger numbers of frames increases the risk that motion blur occurs, and also requires stability of the spacecraft orientation for a greater duration. Highly binned images will be less susceptible to this due to their coarser angular resolution, and therefore a more favourable product for co-addition.

6.4.3 Increasing The Effective Exposure Time

Section 6.2.2.2 describes the frequency at which framelets must be captured in order for EnVisS to build up a complete scan of the sky. In essence, framelets with along-track extents of 5° need to be captured every 5° of spacecraft spin. For a 4 s spin period, this equates to a sampling frequency of 18 Hz (i.e. a framelet needs to be captured every 55.6 ms). Separately, the exposure time, as discussed in section 6.2.2.1, is dictated by the imaging resolution and spin period. For a spin period of 4 s, an image captured at the full resolution of EnVisS' detector requires a maximum exposure time of ≈1 ms.

Under these imaging mechanics, the detector exposes for 1 ms and subsequently reads out the desired filter's columns in order to capture a framelet, and then waits ≈55 ms for the spacecraft to rotate to the correct position for the next framelet to be captured. Thus, the duty cycle (the proportion of the total scanning time that is spent collecting signal) is ~2%. Image SNR would benefit from utilising more of the available time to collect signal, but the imaging geometry and mechanics limit this.

Two techniques by which the effective exposure time (and the duty cycle) can be increased are discussed here. It should also be noted that the technique of binning (section 6.4.1) itself increases the duty cycle, because the exposure time grows as the angular resolution drops, whilst the sampling frequency remains the same (plotted below in Figure 6.32). Once $n_{bin} \geq 50$ (the width of the filter), duty cycle reaches 100% because exposure time and sampling period become equal. The two methods presented in this section are therefore only beneficial when applied to images captured at full resolution, or with low values of n_{bin}.

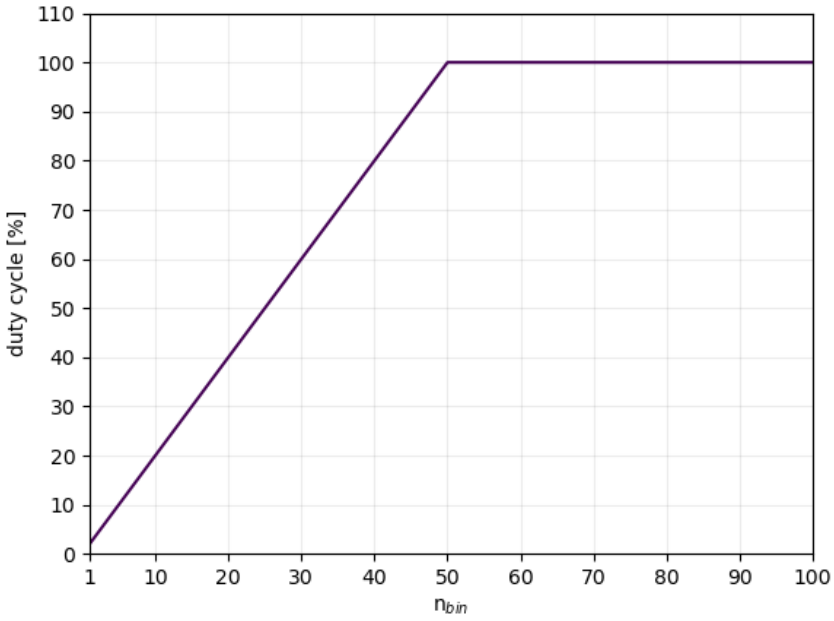

Figure 6.32: EnVisS imaging duty cycle for different levels of pixel binning.

6.4.3.1 Time-Delay Integration

Because push-frame and push-broom cameras are necessarily moving relative to the scene they are imaging, it is common that their short exposure times limit the SNRs they can achieve. As a result, the technique of time-delay integration (TDI, introduced in section 2.3.1.7) exists to mitigate this, and is commonplace on spaceborne scanning cameras. Its use with a wide-angle imager however, especially with the likes of EnVisS, would be unconventional, due to the distortion of the fish-eye optics. However, analysis is presented here to demonstrate that TDI can still be implemented with such an imaging system.

Recall from section 2.3.1.7 that TDI works by shifting the charges accumulated in a detector along its rows, one column at a time (in synchronisation with scene motion), in order to accumulate signal for an extended period without loss of spatial resolution. After a certain number of shifts, the accumulated charge in each of the pixels is read out. If the motion of the shifting charges does not remain synchronised with the motion of the scene, image smear occurs. Consider a single pixel with position (in pixel coordinates) $\vec{p}_0$ on the detector at the beginning ($t = 0$) of an exposure. This pixel views the point or points in the scene which project to position $\vec{p}_s(t = 0) = \vec{p}_0$ on the detector. As the spacecraft spins, $\vec{p}_s$ will move according to

$$\vec{p}_s(t) = \vec{p}_0 + \int_0^t \frac{d\vec{p}_s}{dt} dt \tag{6.34}$$

Meanwhile, over the course of an exposure with n_{TDI} steps, charges will shift $n_{TDI} - 1$ columns from their original position, and the charge which began accumulating within pixel $\vec{p}_0$ at time $t = 0$ will have a position $\vec{p}_c$ on the detector given by

$$\vec{p}_c(t) = \vec{p}_0 + t \begin{bmatrix} f_{TDI} \\ 0 \end{bmatrix} \tag{6.35}$$

where f_{TDI} is the frequency at which TDI charge shifts occur. The maximum number of useable TDI steps before smear is introduced is dictated by the relationship

$$|\vec{p}_s - \vec{p}_c| < r_{max} \tag{6.36}$$

where r_{max} is the maximum acceptable distance on the detector by which the projected scene point and accumulating charge may separate. Figure 6.33 illustrates this concept. The acceptable value of r_{max} for a specific imaging system may depend on factors such as the camera's optics (e.g. its PSF), the required spatial resolution and the required SNR of the final images.

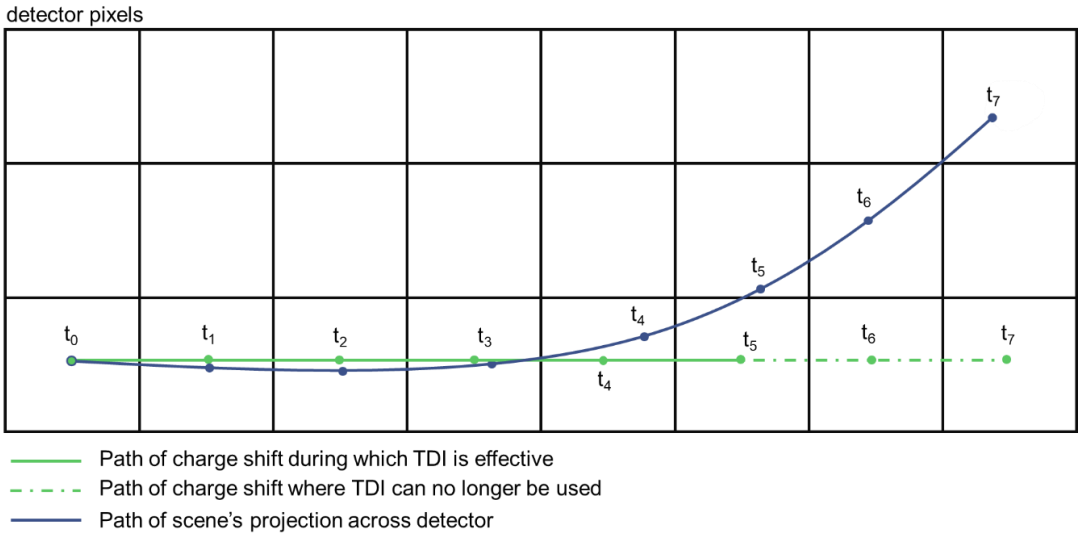

Figure 6.33: Illustration of TDI being performed with curved scene motion. Here, five steps of TDI can be performed. t_0, t_1...t_7 indicate times separated by a constant interval.

The fish-eye optics and large FOV of EnVisS mean that objects in the camera's scene move across the focal plane with curved, rather than straight paths. Figure 6.34 plots the paths taken by the scene across the detector as B2 rotates (whilst assuming that the spacecraft has zero flyby velocity relative to the comet). The vertical component of the motion is colour coded. The green and pink regions, where the scene moves downwards

and upwards respectively, are clearly not compatible with TDI. In the regions coloured white, there is no vertical motion of the scene, and it travels parallel to the detector's rows.

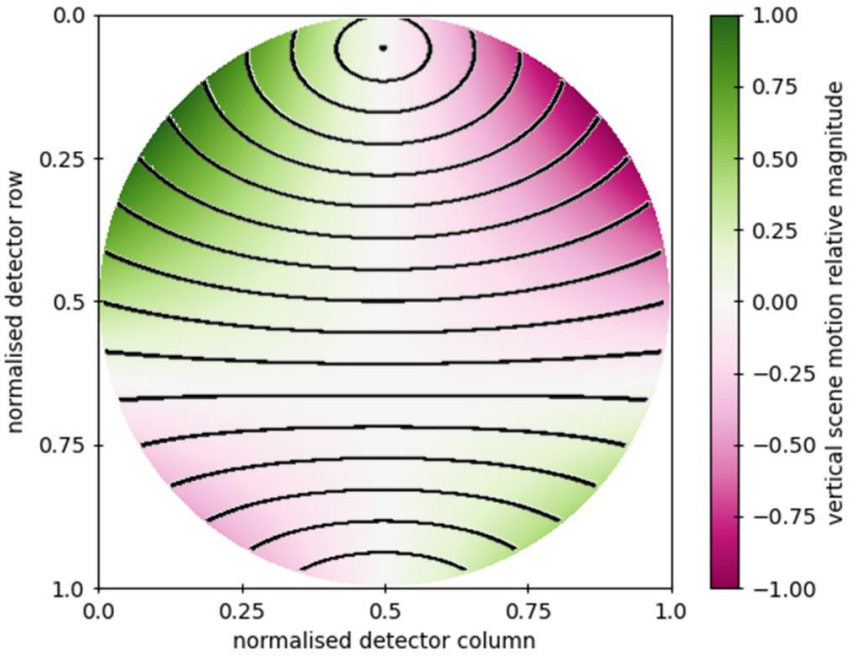

Figure 6.34: Illustration of scene motion on detector due to spacecraft spin. The black lines are paths taken by points in the scene. Green represents positive (downward) vertical velocity and pink negative. White regions have zero vertical velocity.

Applying equations 6.34-6.36 to a specific EnVisS pixel gives the maximum number of TDI steps that can be implemented starting from that pixel. Figure 6.35 shows the result of performing this calculation for all of EnVisS' pixels, taking $r_{max} = 1$ pixel.

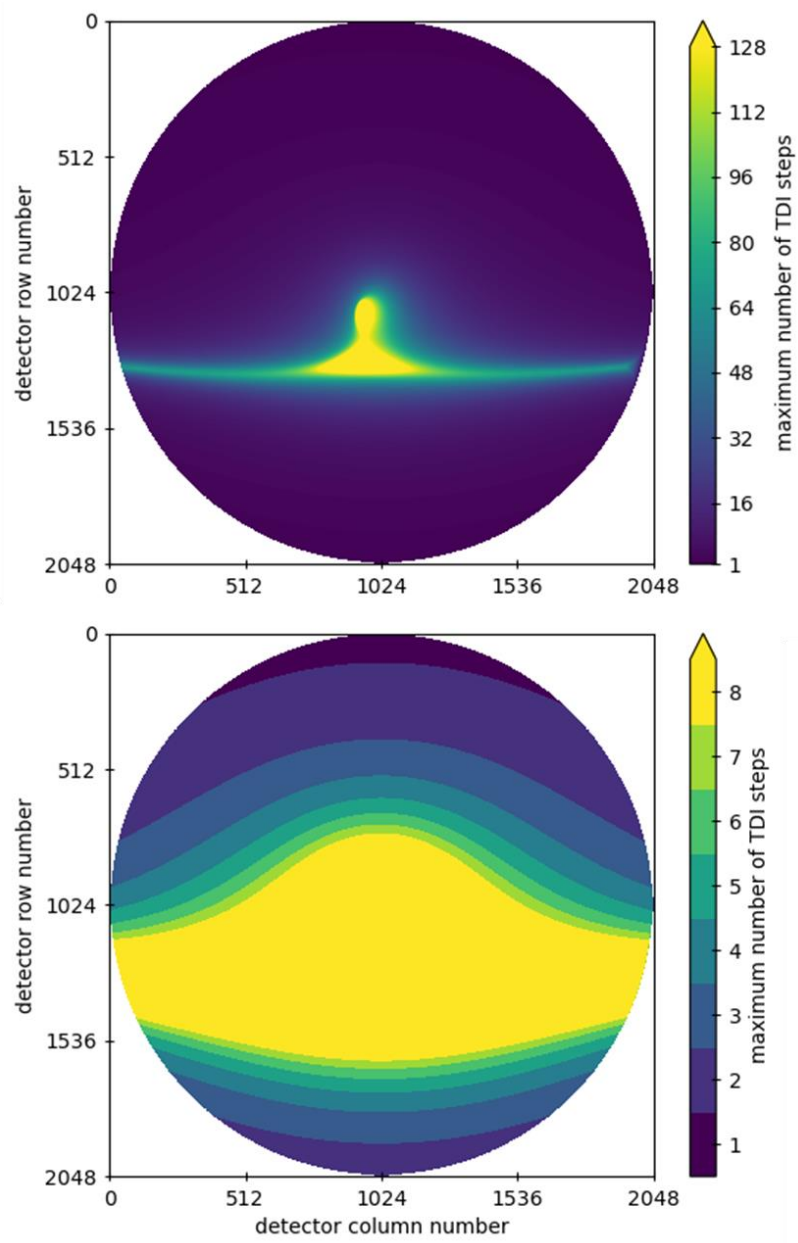

Figure 6.35: The maximum number of TDI steps which can be performed starting from each of EnVisS' pixels. For clarity, the same data are represented with two different scales: 1-128 TDI steps (top) and 1-8 TDI steps.

The top panel of Figure 6.35 indicates that only a very small number of pixels, lying close to the region of the detector where scene motion is parallel to pixel rows, can implement a large number of TDI steps without image smear. However, when focusing on smaller numbers (≤8) of TDI steps (bottom panel), it's clear that sizeable portions of the instrument's FOV could be compatible with the technique. Again, these areas are focussed around the region of horizontal scene motion, and they grow in extent as TDI steps decrease. An n_{TDI} of 2 is compatible with almost the entire EnVisS FOV, but any

higher values result in a reduction in the filters' across-track FORs. This relationship is plotted in Figure 6.36 for a central region of the detector spanning columns 512-1536.

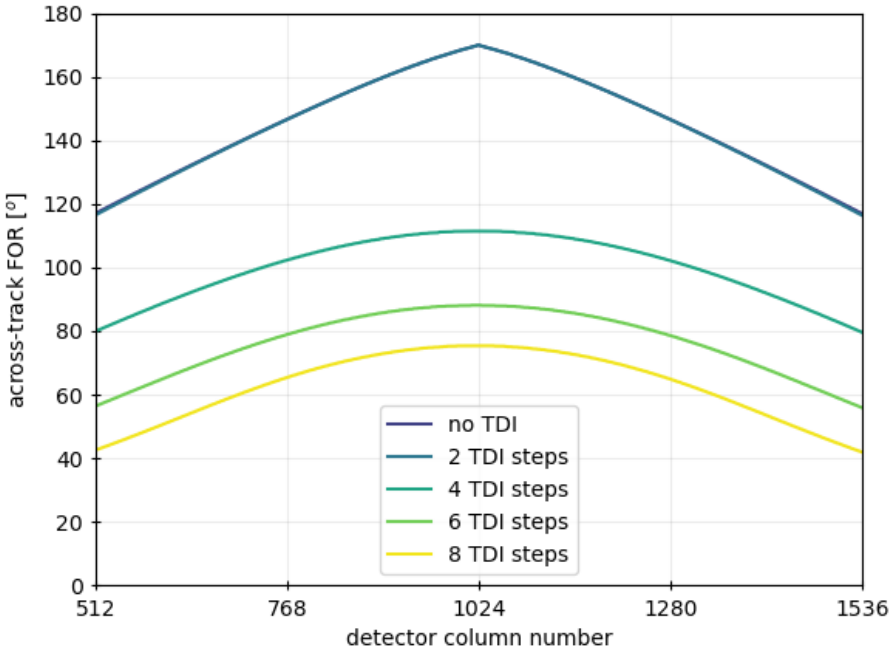

Figure 6.36: Across-track FOR of EnVisS when using TDI, as a function of detector column and n_{TDI}. The line corresponding to 'no TDI' is largely obscured by the '2 TDI steps' line, because use of $n_{TDI} = 2$ has minimal impact on the instrument's FOR.

Note from Figure 6.36 that the FOR when using 2 TDI steps is the same as when TDI is not utilised, meaning a doubling of signal via TDI comes with no FOR trade-off. The assumption of zero flyby velocity used for calculating the data of Figure 6.35 will not hold for all v/R ratios, and the FORs of Figure 6.36 may not in fact be achievable throughout B2's flyby. To test this, equations 6.34-6.36 are again applied to each of EnVisS' pixels, this time accounting for both B2 spin and flyby motion in $\vec{p}_s(t)$ for a range of v/R ratios. A sample of the results are plotted in Figure 6.37.

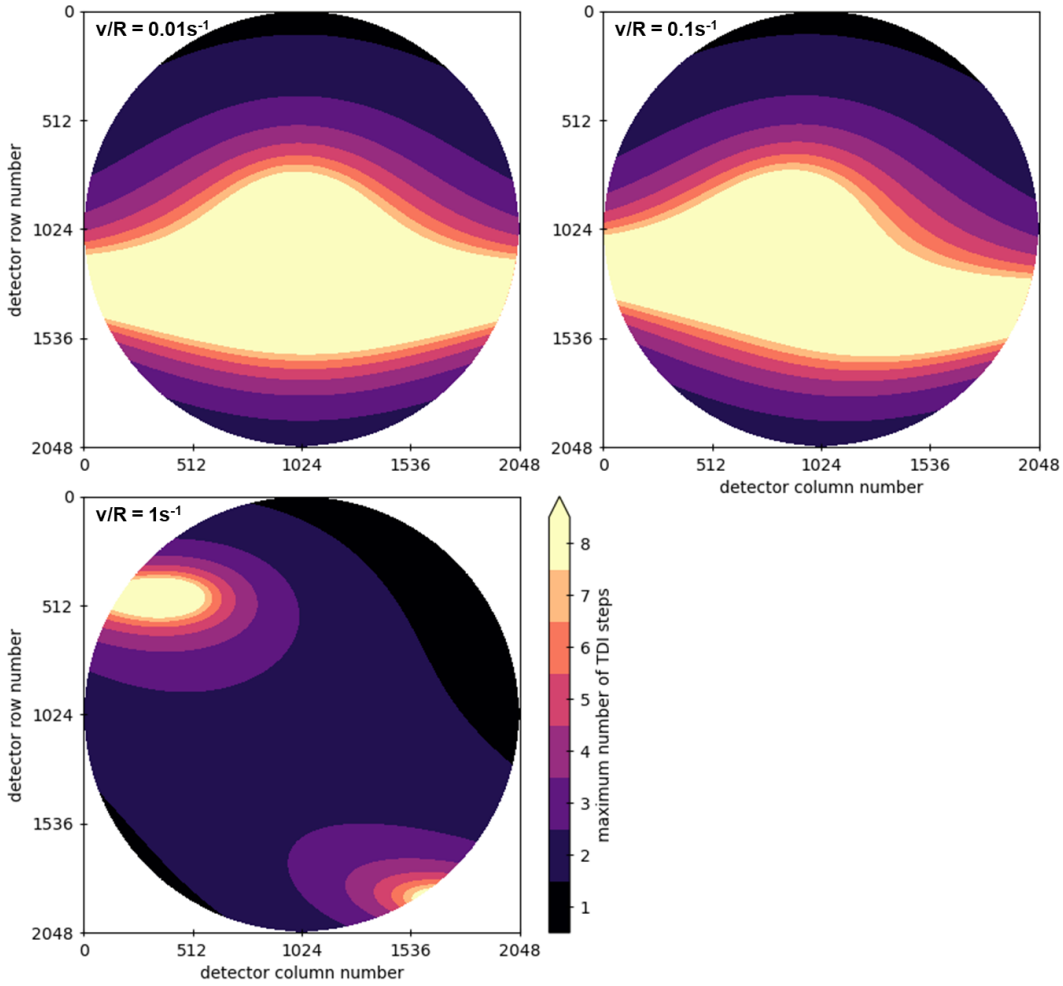

Figure 6.37: The maximum number of TDI steps which can be performed starting from each of EnVisS' pixels for three different v/R ratios, and a 4 s spin period.

At v/R ratios less than 0.01 s^{-1}, flyby velocity has a negligible impact on the distribution of maximum possible n_{TDI} values across the instrument's FOV, whilst $v/R = 1$ s^{-1} can be seen to render TDI ineffective. The FOR of the detector's central column as a function of v/R ratio and n_{TDI} is plotted in Figure 6.38.

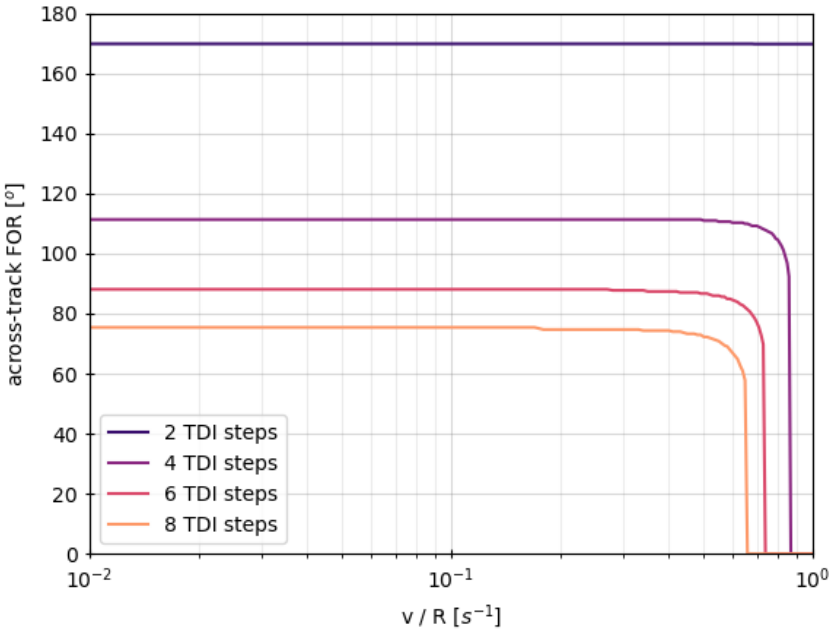

Figure 6.38: Across-track FOR of the central column of EnVisS' detector as a function of v/R when using TDI.

Figure 6.38 shows that flyby velocity has no effect on the achievable FORs when using TDI for v/R ratios up to ~0.4 s^{-1}. Given that 0.2 s^{-1} is expected to be the largest v/R observed by EnVisS, B2's flyby velocity is unlikely to impact the usefulness of TDI.

Whilst TDI with $n_{TDI} > 2$ is not compatible with capturing scans covering the entire sky, it can potentially still provide observations supporting EnVisS' objective of characterising the full phase function of the comet's coma. Let the vectors $\vec{v}$ and $\vec{R}_{\odot}$ define the B2 spacecraft flyby trajectory and the nucleus-Sun line respectively, and let P be the plane containing the vectors $\vec{R}_{\odot}$ and $\vec{v} \times \vec{R}_{\odot}$, and the location of B2 (illustrated in Figure 6.39). If $\vec{R}_{\odot}$ intersects the camera's FOR (i.e. the FOR sees the Sun), then P also lies wholly within the FOR, meaning EnVisS can image along any direction that lies in the plane P. Therefore, a full 180° range of phase angles can be observed by EnVisS over the course of a full 360° rotations of B2, providing the Sun sits within the instrument's FOR.

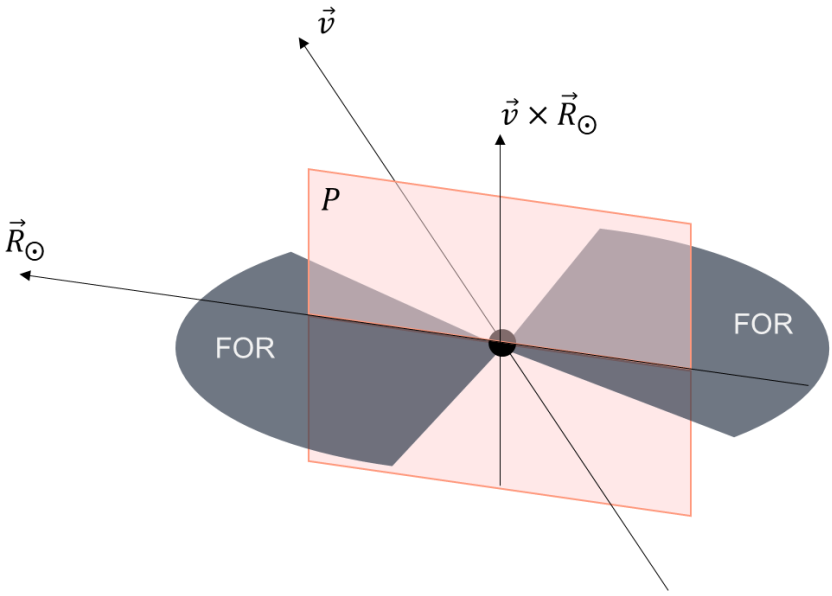

Figure 6.39: Illustration of the geometry described in the main text whereby a full 180° range of phase angles can be observed with incomplete coverage of the sky. $\vec{v}$ is the B2 flyby velocity vector (and spin axis), $\vec{R}_{\odot}$ is the nucleus-Sun line.

ESA's design requirements for the B2 spacecraft nominally constrain the solar aspect angle of its flyby (the angle between $\vec{v}$ and $\vec{R}_{\odot}$) to a range of 45°-135°. This means, in B2-centric coordinates, the Sun is expected to have a polar angle of the same range. Figure 6.40 visualises the FORs associated with different values of n_{TDI} for the untilted (top) and tilted (bottom) mounting configurations of EnVisS, alongside the possible range of solar positions relative to the spacecraft flyby trajectory. It demonstrates that the FORs associated with $n_{TDI} \leq 5$ will be compatible with observing a full phase function for any Sun position, whilst a large range of possible flyby geometries will be compatible with n_{TDI} as high as 8. In fact, for solar aspect angles close to 90°, n_{TDI} could feasibly be larger (though the width of the filter places a limit on the maximum possible number of TDI steps). The acquisition of full phase function observations using TDI could be planned and uploaded to EnVisS in response to the determination of B2's specific flyby geometry after the target comet and intercept trajectories are defined.

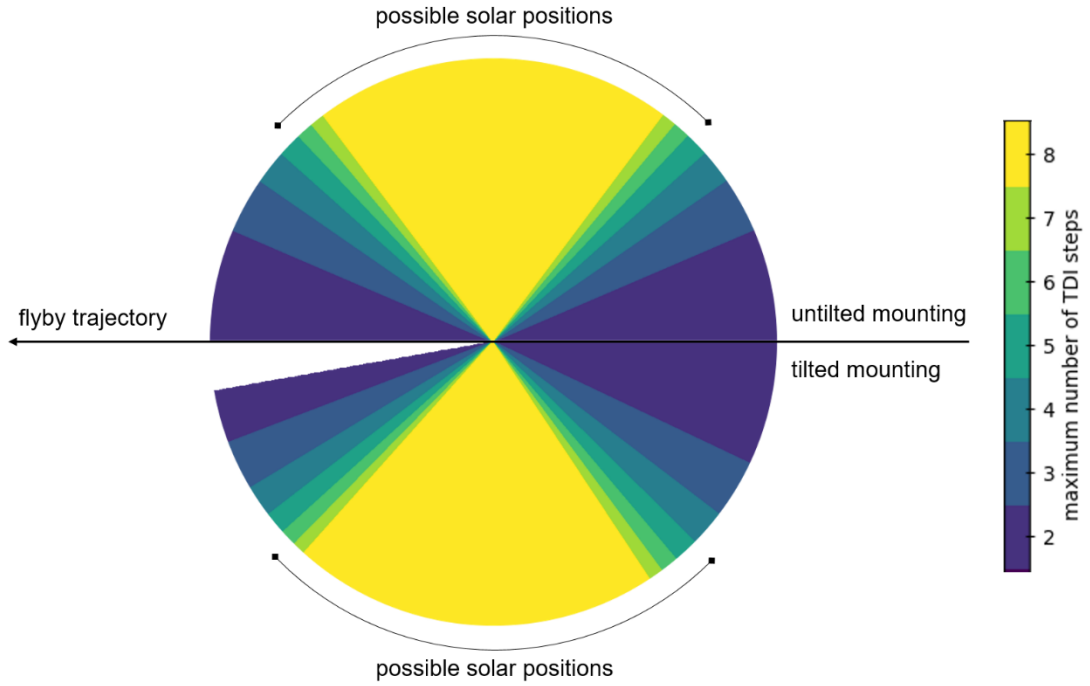

Figure 6.40: Illustration of the ranges of polar angles viewed by FORs associated with different values of n_{TDI}, for both untilted (top) and tilted (bottom) mountings of EnVisS.

Figure 6.41 plots the expected SNRs when utilising TDI, as a function of the observed scene brightness. It is evident that the SNR increases of multiple orders of magnitude that are required over the majority of the expected brightness range cannot be achieved with TDI. The only viable application of TDI would likely be at closest approach, when scene brightness should be at a maximum and it may be desirable to take advantage of TDI's retention of the camera's full imaging resolution, in order to resolve the nucleus and near-nucleus structures.

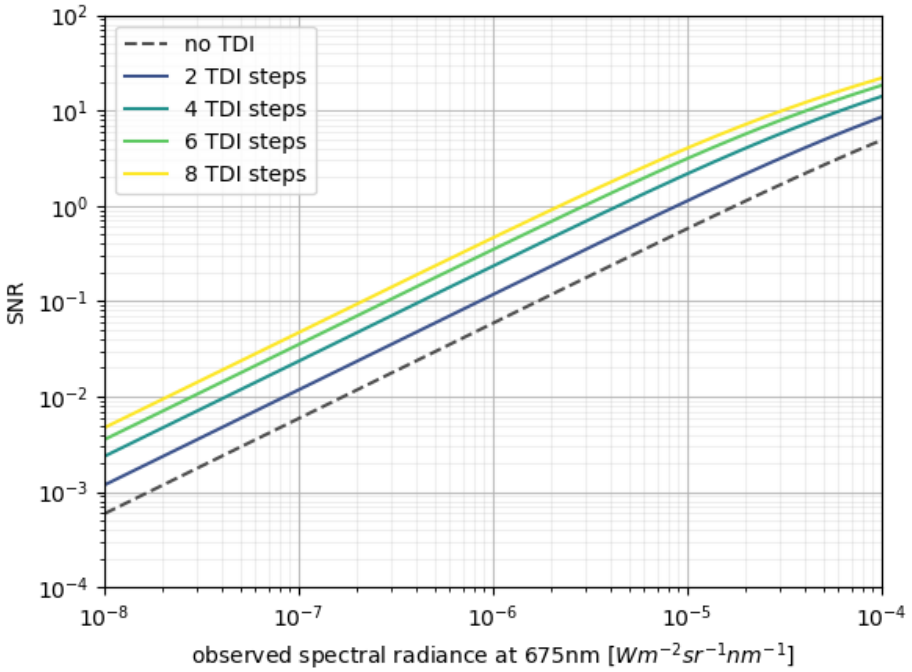

Figure 6.41: SNR of EnVisS images as a function of observed radiance for a range of TDI scenarios.

The TDI analysis carried out here focusses on the limitations imposed by the camera's imaging geometry, and the expected geometry of the spacecraft's flyby. But, as with the pixel binning technique, the camera's hardware is another factor affecting the suitability of TDI. When the analysis presented here was first carried out, the detector that EnVisS would use had not been defined, and was expected to be a CCD. With the design of EnVisS now relying on a CMOS, it will not be possible for the instrument to implement TDI. The following section discusses an alternative, albeit less effective, technique that is compatible with EnVisS' hardware.

6.4.3.2 Repeat Sampling

Repeat sampling provides an alternative means of increasing the duration for which EnVisS collects image signals. As B2 rotates, and the FOVs of its filters advance across the sky, any given point in the sky is viewed by a given filter for a limited period of time, called the filter's dwell time. The dwell time of a filter is a function of polar angle, and some examples are plotted in Figure 6.42.

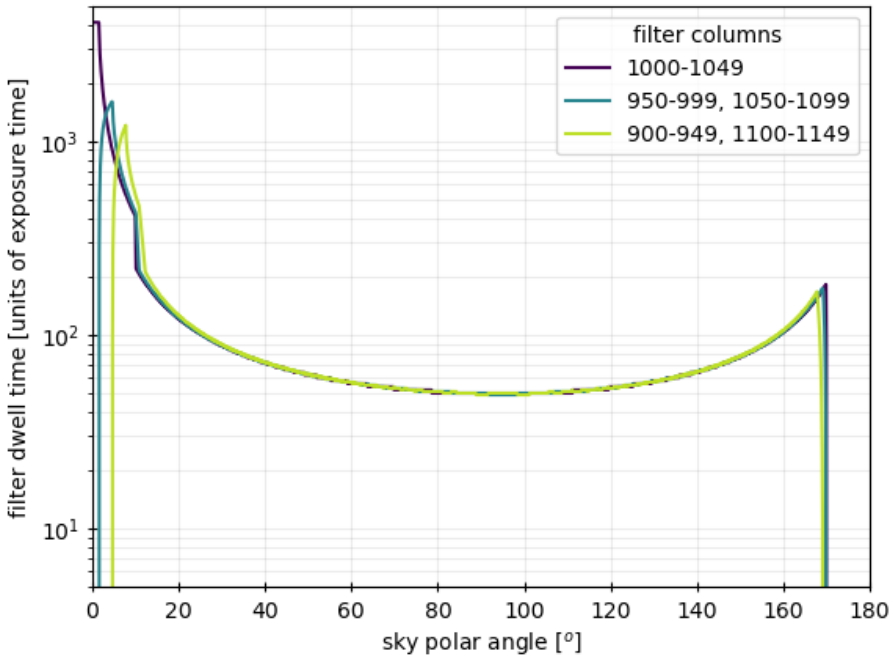

Figure 6.42: Dwell times of five EnVisS filters. Dwell time is plotted in units of full resolution exposure time.

Filter dwell times are smallest at the B2-centric equator, because this is where the filters have the smallest azimuthal extent (e.g. Figure 6.11). In this region of the sky, filter dwell times are approximately 50 times larger than the exposure time (as we already knew from section 6.2.2). This means that there is the potential to image any given point in the sky up to 50 times (i.e. sample it repeatedly). This is achieved by increasing the exposure frequency from its minimum value, resulting in overlap of a filter's consecutive framelets, as illustrated in Figure 6.43. When the imaging frequency is a non-integer multiple of the minimum imaging frequency, some regions of sky are sampled more than others (b). For imaging frequencies which are an integer multiple of the minimum frequency (c), repeat sampling occurs uniformly.

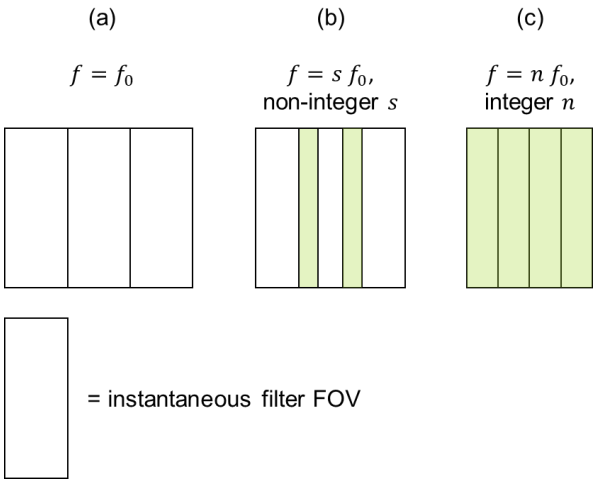

Figure 6.43: An illustration of repeat sampling for three different cases. (a) minimum imaging frequency results in no repeat sampling. (b) imaging at a frequency which is a non-integer multiple of minimum frequency achieves partial repeat sampling of the sky. (c) imaging at a frequency which is an integer multiple of minimum frequency achieves uniform repeat sampling. Green shading indicates regions of increased sampling.

When overlapping framelets are mapped to an all-sky projection their signals can be combined in order to increase the image SNR. Because these combinations of signals are occurring after they have been read out from the detector, read noise is also present. Therefore, for a region of sky which is sampled n_s times, the approximate factor by which its SNR is increased over the baseline case is $\sqrt{n_s}$. It is worth noting from Figure 6.42 that some regions of sky (particularly polar angle <20°) are sampled repeatedly even at minimum imaging frequency.

Figure 6.44 plots the number of times a given point in the sky with polar angle between 80° and 100° (i.e. the region of sky for which there is no repeat sampling at minimum frequency) is sampled by any one of EnVisS' filters as a function of imaging frequency, up to the camera's maximum data rate of 30 Mpix s^{-1}. The separate maximum and minimum lines reflect the fact that some azimuths are imaged more times than others when imaging frequency is a non-integer multiple of the minimum (Figure 6.43 (b)). The corresponding approximate factor by which SNR is increased from the minimum frequency case is also indicated.

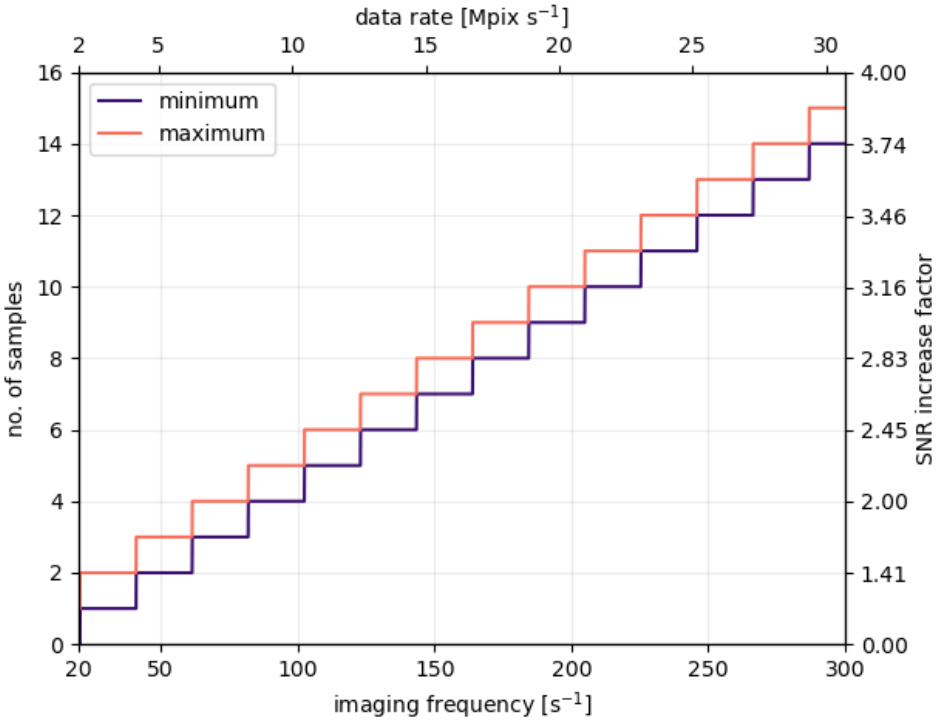

Figure 6.44: Number of times an equatorial (in B2-centric coordinates) point in the sky is imaged by a filter during an all-sky integration as a function of imaging frequency and data rate. The corresponding factor by which SNR is increased verses the minimum frequency case is also indicated.

The detector's maximum data rate and the expected filter widths of 50 pixels constrain the possible SNR gains of repeat sampling to a factor of less than four. This is not sufficient to make a significant difference to the range of brightnesses that EnVisS can usefully observe, and certainly is too small an increase to assist with observations of brightnesses below 10^{-6} $Wm^{-2}sr^{-1}nm^{-1}$. The only advantage the technique offers over binning is that, in principle, it can retain full image resolution, and so its usefulness would be limited to increasing the SNR in bright images, where it is also desirable to maintain a high spatial resolution. This would likely be limited to imaging at closest approach, where the bright nucleus and near-nucleus features (e.g. dust jets) could potentially be resolved.

6.4.4 Summary

Table 6.12 provides a summary of the approximate factors by which the SNR enhancement techniques discussed above, and combinations thereof, increase image SNR from the baseline case. Grey cells indicate combinations which are not possible, feasible or effective to implement. For example, TDI and repeat sampling, which are only

effective if maintaining a high resolution is required, would not be compatible with the co-addition of 50 images, due to the impact of motion blur.

Table 6.12: A summary of the approximate factors by which different combinations of imaging techniques increase image SNR. Incompatible combinations are shaded grey.

		Binning				Co-addition			8x TDI	Repeat Sampling
		5x5	10x10	50x50	50x100	5	10	50		
Binning	5x5	25				50	80	170*		
	10x10		100			200	300	500*		
	50x50			2500		5000	8000	15000		
	50x100				5000	10000	15000	30000		
Co-addition	5	50	200	5000	10000	2			15	7
	10	80	300	8000	15000		3*		25	12
	50	170*	500*	15000	30000			7*		
8x TDI						15	25		8	
Repeat Sampling						7	12			4

*Will not be possible for the majority of the flyby due to scene motion and image smear.

Table 6.12 demonstrates a number of key points:

- Pixel binning is the only technique that can, on its own, produce significant and sufficient SNR increases (i.e. >100-1000). Co-addition is valuable, and can boost SNR significantly when used alongside binning, but on its own is ineffective.
- Observing the lowest expected brightnesses will require large levels of binning and co-addition to achieve sufficient SNRs. At the beginning of the flyby, it may be possible to increase SNR by a factor ~10^4, facilitating large scale coma observations.
- TDI and repeat sampling produce minimal gains in SNR compared to binning. Their value comes from their retention of full resolution. Therefore, their use should be considered only for imaging bright, small scale structures (e.g. the nucleus and its immediate surroundings at closest approach).

Mitigation of low SNRs will be an essential aspect of EnVisS' imaging design, and this has had a profound impact on the evolution of the instrument since its first concept. Data will be of a different form (significantly reduced resolution) to that originally envisaged. The analysis that has been presented here addresses only the camera's broadband filter, ignoring any narrowband spectral filters. This is because narrowband filters, which accept ~1/100th as much light as the broadband, suffer even more so from low SNRs,

and will likely be unusable. The EnVisS design at time of writing no longer includes narrow band filters, due both to a need to reduce the number of filters on the detector, and the analysis of this section.

6.5 Imaging Polarimetry

The design of EnVisS is unique amongst spaceborne cometary cameras not just because of its imaging mechanism and the all-sky coverage of its images, but also because it includes a technique that is commonly employed by Earth-based cometary observations, but has not been used from within a comet's coma: imaging polarimetry. Polarimetry – the study of light's polarisation – provides information that studying its intensity and wavelength cannot. Cometary dust is responsible for polarising light, and measurement of the polarisation therefore probes the nature of the dust. Observations of polarisation have an advantage over intensity observations in that they do not need to be normalised to account for distances between the Sun, the comet and the observer (Hadamcik & Levasseur-Regourd, 2003), and together, intensity and polarisation observations give a more complete description of the light and the processes through which it has gone.

In imaging polarimetry, the spatial variation of polarisation is recorded, just as images normally record the spatial variation of intensity. The temporal, spectral and comet-to-comet (and other astronomical bodies) variation of polarisation all also provide valuable insight, and each can itself be combined with imaging polarimetry to better understand the physical properties of the scattering media.

6.5.1 Polarisation and the Stokes Parameters

Light waves comprise oscillating electric and magnetic fields, whose directions of oscillation are mutually perpendicular, and also perpendicular to the light's direction of propagation. Light from many sources, including the Sun, is an ensemble of many randomly oriented electromagnetic waves, with no preferential directions for the electric or magnetic field oscillations. Such light possesses no polarisation, and is known as unpolarised. In contrast, some emission, reflection and scattering processes produce light with preferentially oriented electric and magnetic fields (the reflection of light by the surface of water, and atmospheric scattering, for example), which by definition is known as polarisation.

Consider an arbitrary 2D Cartesian coordinate system with axes $\hat{x}$ and $\hat{y}$ both perpendicular to a propagating beam of polarised light. At any instant, the net electric

field vector of the light (by convention polarisation is defined by the electric field direction, rather than magnetic) can be imagined to draw a straight line on the x-y plane. For polarised light, in the most general case, the time-evolution of the tip of this electric field vector describes an ellipse on the x-y plane (Figure 6.45 (a)), and this is known as elliptically polarised light. Two special cases of elliptical polarisation exist: circular and linear polarisation. For circular polarisation, the path described by the electric field vector is circular (Figure 6.45 (b)), and for linear polarisation, the electric field vector's path describes a straight line on the x-y plane (Figure 6.45 (c)).

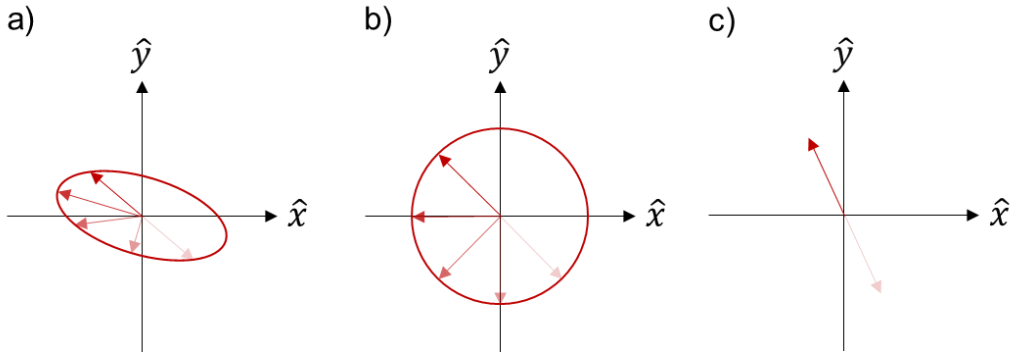

Figure 6.45: Elliptical (a), circular (b) and linear (c) polarisation, illustrated by the time-evolution of the light's electric field vector on the x-y plane, where $\hat{x}$ and $\hat{y}$ define an arbitrary Cartesian coordinate system, and are both perpendicular to the light's direction of propagation (out of the page).

For unpolarised light, at any instant in time the electric field vector will also form a line on the x-y plane, but its time-evolution will be random and rapid, such that averaging over even a short time interval results in no overall electric field direction. Light can be fully polarised, whereby all constituent waves' oscillations conform to the polarisation, or partially polarised, meaning that its intensity can be described as a polarised contribution combined with an unpolarised contribution.

The polarisation state of light is often described using the four Stokes parameters, favoured for their mathematical convenience and direct measurability. The Stokes parameters are defined in terms of time-averaged electric fields $\langle E \rangle$ as

$$s_0 = \langle {E_x}^2 \rangle + \langle {E_y}^2 \rangle = \langle {E_a}^2 \rangle + \langle {E_b}^2 \rangle = \langle {E_l}^2 \rangle + \langle {E_r}^2 \rangle,$$

$$s_1 = \langle {E_x}^2 \rangle - \langle {E_y}^2 \rangle,$$

$$s_2 = \langle {E_a}^2 \rangle - \langle {E_b}^2 \rangle,$$

$$s_3 = \langle {E_r}^2 \rangle - \langle {E_l}^2 \rangle \qquad (6.37)$$

where the electric field subscripts denote three different coordinate bases: $(\hat{x}, \hat{y})$ as already described; $(\hat{a}, \hat{b})$, which is a 45° rotation of the $(\hat{x}, \hat{y})$ basis; and circular basis $(\hat{l}, \hat{r})$, associated with the clockwise and anti-clockwise rotation of the electric field vector. The $\hat{x}$ and $\hat{y}$ directions are referred to as horizontal and vertical respectively, whilst the directions of $\hat{a}$ and $\hat{b}$ are referred to as 45° and 135° respectively. These coordinate bases are illustrated in Figure 6.46.

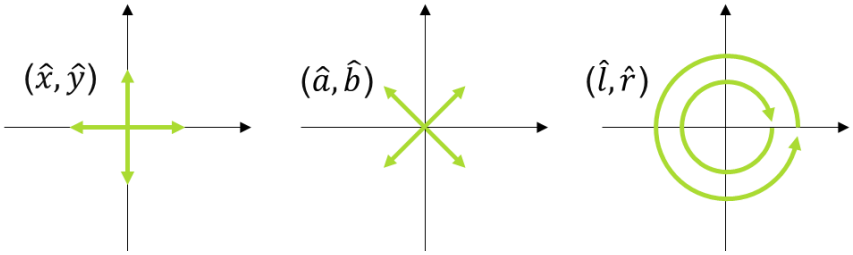

Figure 6.46: Diagram illustrating the three coordinate bases associated with the Stokes parameters, as described in the main text.

Each Stokes parameter describes a different element of the light's polarisation state:

- s_0: the light's total intensity.
- s_1: the prevalence of horizontally linearly polarised light (indicated by positive values) versus vertically linearly polarised light (indicated by negative values).
- s_2: the prevalence of 45° linearly polarised light (indicated by positive values) versus 135° linearly polarised light (indicated by negative values).
- s_3: the prevalence of right (clockwise) circularly polarised light (indicated by positive values) versus left (anticlockwise) circularly polarised light (indicated by negative values).

Often, when circular polarisation is known to be negligible, the fourth Stokes parameter is omitted, leaving s_0, s_1 and s_2, known as the linear Stokes parameters. The light scattered by cometary dust generally contains only low levels of circular polarisation, and the use of only the linear Stokes parameters is common practice in cometary polarimetry (e.g. Hadamcik & Levasseur-Regourd (2003); Jockers et al. (2005)). Indeed, EnVisS is designed to measure only the linear Stokes parameters, and the presence and measurement of circular polarisation will be ignored in all analysis that follows.

6.5.2 Measuring the Polarisation State of Light

The simplest possible polarimetric measurement of light is to measure its intensity after it has passed through a polariser, a polariser being an optical filter which only transmits

light whose plane of polarisation has the same orientation as the filter's polarising axis. A perfect polariser, whose polarising axis is oriented at an angle θ from the $\hat{x}$ axis, will transmit an intensity I_θ given by

$$I_\theta = \frac{1}{2}(s_{0in} + s_{1in}\cos 2\theta + s_{2in}\sin 2\theta) \tag{6.38}$$

when placed in a beam of light with linear Stokes parameters s_{0in}, s_{1in} and s_{2in} (Chen, et al., 2014). Determining all three linear Stokes parameters of the input light requires a minimum of three such intensity measurements to be made, each with a different value of θ. Polarimetry of the light scattered by comets normally aims to determine two characteristics of its polarisation: its degree of linear polarisation (DOLP), P_{lin} (the ratio of the polarised light's intensity to the total intensity of the light) and its angle of polarisation (AOP), Θ (the orientation of the plane in which the light is linearly polarised). These are related to the linear Stokes parameters by equations 6.39 and 6.40 respectively:

$$P_{lin} = \frac{\sqrt{{s_1}^2 + {s_2}^2}}{s_0} \tag{6.39}$$

$$\Theta = \frac{1}{2}\tan^{-1}\frac{s_2}{s_1} \tag{6.40}$$

Imaging polarimetry of comets concerns the spatial mapping of DOLP and AOP, and has so far exclusively been performed with distant observatories (either Earth-based or from spacecraft such as STEREO (Thompson, 2015)). EnVisS is designed to implement the technique in a way that is novel to cometary polarimetry in two regards: performing it over the entire sky, and performing it during a flyby through the comet's coma.

The major categories of observational techniques used to perform imaging polarimetry in a range of scientific and engineering fields are described by Tyo et al. (2006). The vast majority of cometary polarimetry observations are obtained using what is known as the division-of-time technique (e.g. Hadamcik & Levasseur-Regourd (2003); Jockers et al. (2005)). In the division-of-time technique, a single camera sequentially acquires multiple images, and a polariser in front of or within the imager is rotated to a different orientation for each image. This is equivalent to the approach used for many multispectral spaceborne observations, whereby multiple monochromatic images are captured by a single camera, but a different spectral filter is placed in front of the detector for each image (e.g. Keller et al. (2007)). The technique is simple to implement, but it bears the complication that the images are not contemporaneous, and therefore not well suited for observation of a rapidly changing scene (which may result from movement of the

camera, or evolution of the scene itself). For observations made from distant telescopes, there is minimal variation in the observed image over short periods, and this method is favoured.

6.5.3 Cometary Polarimetry

Whilst recorded observations of comets date back millennia, polarimetry of comets, which is now a significant and important field in modern cometary science, dates back to just the 19th century. In 1811, Francois Arago built a polariscope, and with it made the first observations of extra-terrestrial polarised light, by observing the disk of the Moon (Dougherty & Dollfus, 1989). Arago refined his measurement technique, and in 1819 measured the polarisation in the tail of comet C/1819 N1 (Tralles) (Kiselev & Rosenbush, 2004). Subsequent observations of comet Halley confirmed comets as sources of polarised light, and Arago postulated that the polarimetry was induced by the scattering of sunlight (which is originally unpolarised) by cometary particles (Arago, 1858). In the majority of observations, the plane in which cometary light is polarised was found to be perpendicular to the scattering plane (defined by the positions of the Sun, the comet and the observer) (Dobrovolsky, et al., 1986).

In the late 19th century, as photographic polarimetry developed, polarisation was found to vary from comet to comet, and within individual comets both spatially and temporally. Öhman (1941) made spectroscopic polarimetric (spectropolarimetric) observations of comets and showed that emissions from cometary gas are also polarised, though to a lesser degree than the dust's scattered light. Observations made by Öhman and others throughout the 20th century showed that the polarisation of both scattered and emitted cometary light depends on the observation phase angle. For each, maximum polarisation is observed at phase angles close to 90°, where the DOLP of scattered light usually reaches up to 20-30%. Whilst the polarisation of molecular emission is the subject of some cometary polarimetry studies, the vast majority focus on the light scattered by cometary dust. This is the objective of EnVisS, too, so the focus here will be on dust scattering.

The polarimetric phase curves of comets (i.e. the variation of polarisation as a function of observer phase angle) comprise two sections, known as the positive and negative polarisation branches. The positive branch, which covers phase angles larger than α_{inv} (called the inversion angle) is characterised by positive polarisation, meaning the plane of polarisation is perpendicular to the scattering plane. The negative branch covers phase angles below α_{inv}, and exhibits negative polarisation, where the plane of polarisation is parallel to the scattering plane. The inversion angle, α_{inv} is found

consistently across comets to have a value ≈22° (Kiselev, et al., 2005). The cause of the negative branch, which reaches a minimum polarisation of approximately -1.5% at a phase angle of ≈10%, is largely attributed to scattering by aggregate particles (Petrova, et al., 2001). Comets' positive polarisation branches are poorly observed at large phase angles (≥120°), because of the challenges associated with making measurements directed towards the Sun.

In polarimetry, comets are often split into two categories: 'dust-rich' and 'gas-rich', due to differences in the polarisation they exhibit. Gas-rich comets generally have lower polarisation than dust-rich. The reason for this is thought to be a difference in the scattering properties of the dust in the two classes of comets (caused for example by a difference in particle size distribution) (Levasseur-Regourd, et al., 1996), but it is also suggested that the cause could be the depolarisation of the light due to scattering by the comets' gasses (Kiselev, et al., 2005).

Within individual comets, polarisation is often observed to exhibit significant spatial variation. Observations of seven comets made by Hadamcik and Levasseur-Regourd (2003) found jets and other ejection structures to have higher polarisation than the surrounding coma, indicating a significant difference in the dust properties within these structures. Imaging polarimetry of comet C/2011 W3 (Lovejoy) observed the DOLP of its dust tail to increase with distance from the nucleus, which the author suggested indicated stratification of dust particles by radiation pressure (Thompson, 2015).

6.5.4 EnVisS Imaging Polarimetry

The polarimetric imaging design of EnVisS is the same as that of its intensity imaging technique. Narrow polarising filters in front of the detector will be used to capture push-frame scans of the sky, and each filter will be mounted with a differently oriented polarising axis, in order to capture intensity images with different θ values (equation 6.38). During the development of EnVisS and its polarimetry concept, two filter arrangements have been considered: 3-filter and 4-filter. Their difference is the number of individual measurements (and filters) required to perform them (3 and 4 for 3-filter and 4-filter methods respectively), and the associated mathematics required to derive the linear Stokes vectors.

6.5.4.1 3-filter Polarimetry

Linear polarimetry involves the characterisation of the first three Stokes parameters, and thus requires a minimum of three measurements. Minimising the number of measurements required by a polarimetry instrument may be desirable when observing

time is limited, or as is the case for EnVisS, there is limited physical space on the detector for accommodating filters.

Sparks and Axon (1999) describe the 3-filter method. Three separate measurements of intensity (I_1, I_2, I_3) are made through polarisers of three different orientations (θ_1, θ_2, θ_3). Assuming the polarisers to be perfect, the Stokes parameters are found by the equation

$$\begin{pmatrix} s_0 \\ s_1 \\ s_2 \end{pmatrix} = c\, M \begin{pmatrix} I_1 \\ I_2 \\ I_3 \end{pmatrix} \tag{6.41}$$

where

$$c = \frac{2}{\sin(2\theta_2 - 2\theta_1) + \sin(2\theta_3 - 2\theta_2) + \sin(2\theta_1 - 2\theta_3)} \tag{6.42}$$

and

$$M = \begin{bmatrix} \sin(2\theta_3 - 2\theta_2) & \sin(2\theta_1 - 2\theta_3) & \sin(2\theta_2 - 2\theta_1) \\ \sin 2\theta_2 - \sin 2\theta_3 & \sin 2\theta_3 - \sin 2\theta_1 & \sin 2\theta_1 - \sin 2\theta_2 \\ \cos 2\theta_3 - \cos 2\theta_2 & \cos 2\theta_1 - \cos 2\theta_3 & \cos 2\theta_2 - \cos 2\theta_1 \end{bmatrix} \tag{6.43}$$

In principle, any three distinct polariser orientations can be used to achieve linear Stokes measurement through these equations. However, measurement accuracy is maximised with an optimal setup in which the polarisers are oriented with their polarising axes mutually offset by 60° (e.g. 0°, 60° and 120°) (Chen, et al., 2014). This implementation has been used on spaceborne instruments including the Hubble Space Telescope's (HST) Faint Object Camera (Jedrzejewski, et al., 1998) and the STEREO SECCHI COR1 coronagraphs (Thompson & Reginald, 2008).

For the optimal polariser arrangement of $\theta_1 = 0°$, $\theta_2 = 60°$ and $\theta_3 = 120°$, equations 6.41-6.43 simplify to:

$$s_0 = {}^{2}/_{3}\,(I_1 + I_2 + I_3)$$

$$s_1 = {}^{2}/_{3}\,(2I_1 - I_2 - I_3)$$

$$s_2 = {}^{2}/_{\sqrt{3}}\,(I_2 - I_3) \tag{6.44}$$

6.5.4.2 4-filter Polarimetry

By measuring the intensity of light at four different polarisation angles rather than three, the recovery of the linear Stokes parameters is simplified and some redundancy is added

to the system. However, the 4-filter method offers no flexibility in the orientation of the filters, and requires either more space or time to obtain the full set of intensity measurements. It follows from the definition of the Stokes parameters (equation 6.37) that with four measurements of intensity through polarisers with orientations 0°, 45°, 90° and 135°, the linear Stokes parameters are given by:

$$s_0 = I_0 + I_{90} = I_{45} + I_{135}$$

$$s_1 = I_0 - I_{90}$$

$$s_2 = I_{45} - I_{135} \tag{6.45}$$

Measurement of cometary polarisation from ground based telescopes commonly utilises this approach (e.g. Hadamcik & Levasseur-Regourd (2003); Kiselev, et al. (2004)), and HST's Wide Field Planetary Camera 2 also performed imaging polarimetry by this method (Burrows, et al., 1994).

6.5.5 Polarimetry Error Analysis

The signal challenges faced by EnVisS, as described in section 6.3, will be an important and limiting factor for all of its imaging modes. SNR will no doubt influence the accuracy with which polarimetric measurements can be acquired, especially given that polarimetric measurements require higher SNRs than intensity measurements.

It is valuable to have a quantitative estimate of the accuracy that EnVisS is expected to achieve in its polarimetric measurements. This helps to constrain the SNR requirements for EnVisS' polarimetry. Additionally, with two possible focal plane designs (3-filter and 4-filter) for performing imaging polarimetry under consideration for use in the EnVisS instrument, it is beneficial to assess whether the achievable measurement accuracies are significantly different for the two methods. In this vein, numerical analysis of both the 3-filter and 4-filter methods was carried out to estimate their measurement accuracies.

6.5.5.1 Error Estimation Method

The purpose of the polarimetry error estimation is to derive the relationship between the SNR of EnVisS' images and the accuracy of the measurements (DOLP and AOP) derived through its imaging polarimetry for both the 3- and 4-filter methods. To achieve this, a Monte Carlo simulation of EnVisS' polarimetric imaging was carried out. Feng et al. (2013) introduce a simple method of estimating the impact of intensity measurement uncertainty on polarimetric measurement accuracy, whilst Chen et al. (2014) perform simulated polarimetry measurements to estimate the relationship between camera detector noise and polarimetry accuracy. A similar methodology is employed here, using

the EnVisS camera and noise model developed for section 6.3, in order to estimate EnVisS' specific polarimetry accuracy characteristics, and crucially, to compare the performance of the 3- and 4-filter polarimetry methods.

Sets of images of a uniform source of polarised light were simulated as being captured by EnVisS, through its polarising filters (orientations 0°, 60°, 120° for 3-filter; 0°, 45°, 90°, 135° for 4-filter). From these images, the linear Stokes parameters of the light source were calculated (with equations 6.44 and 6.45), followed by its DOLP and AOP (using equations 6.39 and 6.40 respectively). Additionally, intensity images of the light source were simulated, captured in an identical way to the polarised images except for the omission of polarising filters. These intensity images were used to define the SNR associated with the polarimetric measurements. The simulation process is illustrated in Figure 6.47.

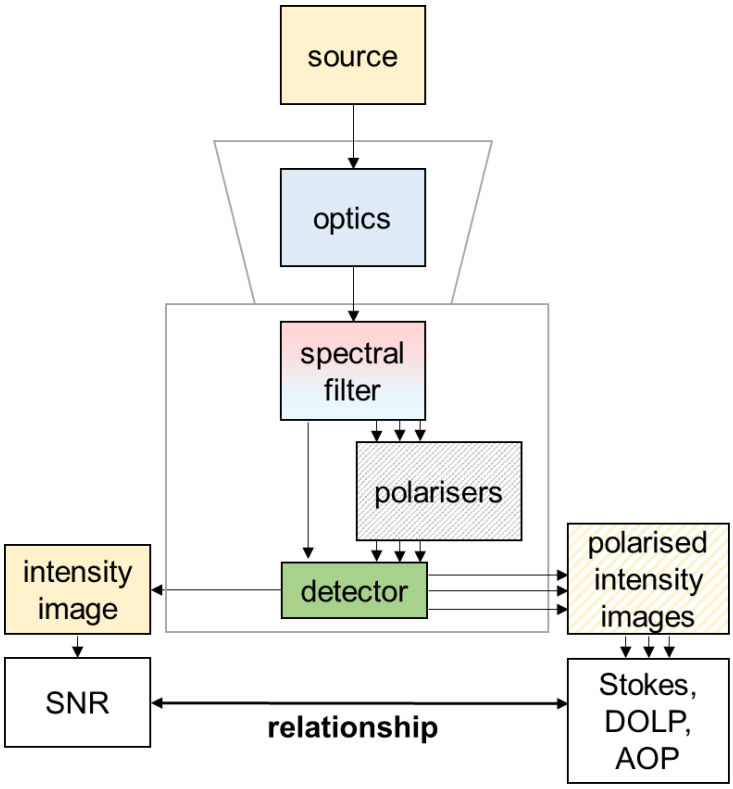

Figure 6.47: Simulation design for estimation of the relationship between SNR and polarimetric accuracy.

SNR and DOLP/AOP accuracy were estimated for a range of source brightnesses, DOLPs and AOPs. For each combination, simulated measurements were repeated 1000 times, each with the addition of random noise (Poisson and Gaussian) according to the EnVisS instrument model. For example, for a single 3-filter measurement, 1000 images

through each of the three filters (3000 images in total) were simulated, along with another 1000 un-polarised intensity images.

6.5.5.2 DOLP Measurement Error

Random noise in the polarised intensity images used for deriving polarisation state leads to both systematic and random errors in the measurement of DOLP. The variation of systematic measurement error as a function of input DOLP and image SNR is plotted in Figure 6.48, for both 3-filter and 4-filter methods.

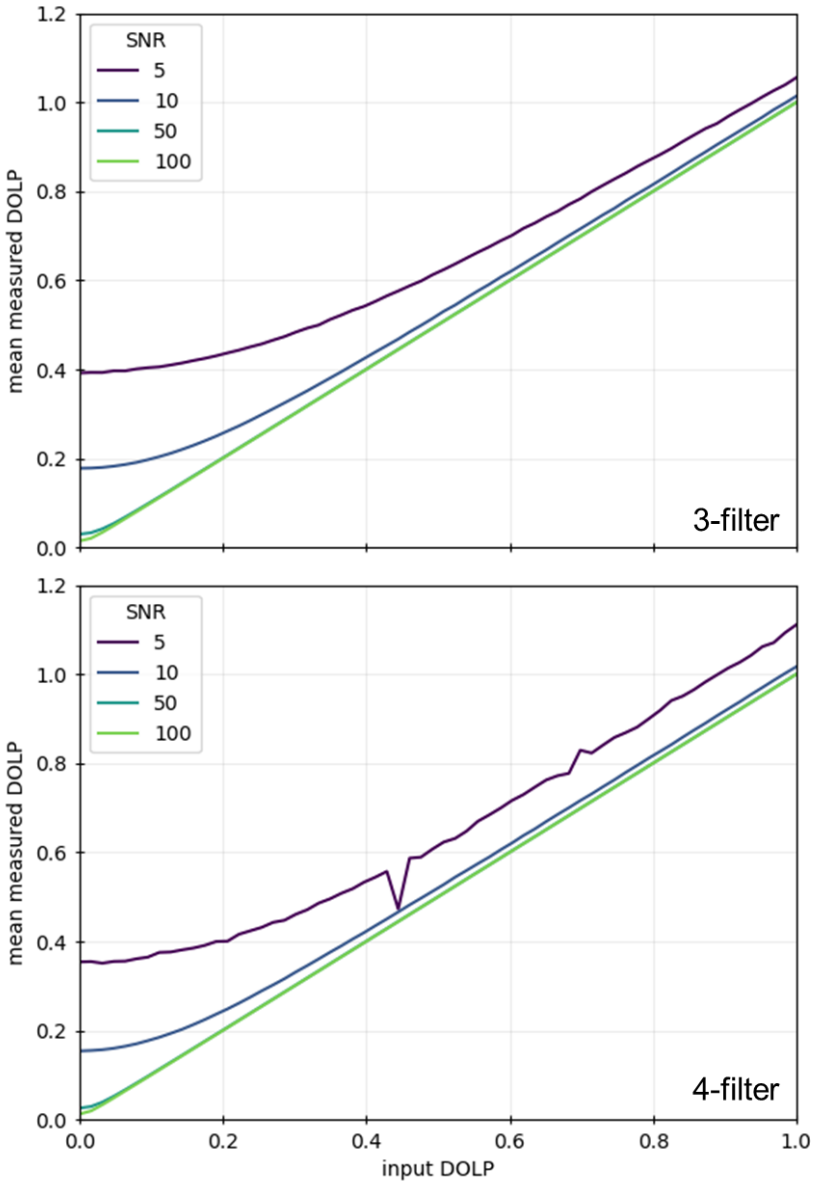

Figure 6.48: Mean measured DOLP versus input DOLP and SNR, for 3- and 4-filter methods.

There is no marked difference between the systematic error of the 3- and 4-filter methods (the latter has a fractionally smaller error). Image noise leads to an over estimation of the DOLP, which is due to the quadrature addition of s_1 and s_2 in equation (6.39) (Patat & Romaniello, 2006). The magnitude of this systematic error is larger both when measuring a small input DOLP, and when SNR is lower. As input DOLP increases, mean measured DOLP approaches the true value, though notably for SNR=5, systematic error remains significant (>6% relative error) for all values of DOLP.

Also significant is the random error inherent to DOLP measurements. The standard deviation of measured DOLP as a function of image SNR is plotted in Figure 6.49 for 3- and 4-filter measurement methods, and for three different input DOLPs.

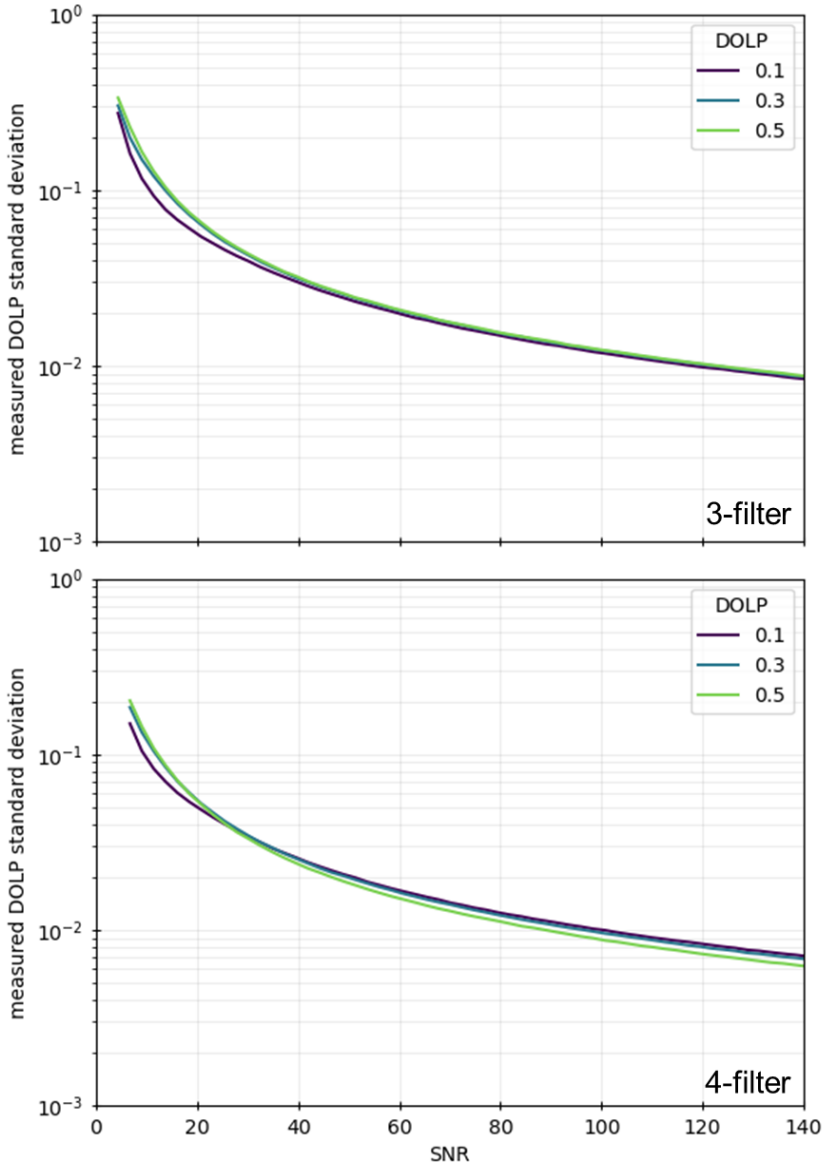

Figure 6.49: Standard deviation of measured DOLP as a function of image SNR and input DOLP for 3- and 4-filter methods (for an input AOP of 0°).

Figure 6.49 demonstrates the strong relationship between the standard deviation of DOLP measurement and image SNR, with DOLP measurement error decreasing by two orders of magnitude when SNR is increased from 5 to 140. As with systematic error, both 3- and 4-filter methods exhibit comparable error profiles, with 4-filter achieving a marginally better measurement accuracy over the plotted SNR range. The DOLP of the input light has little effect on its absolute measurement error, meaning that the relative measurement error is greater for smaller input DOLPs.

The variation of DOLP measurement error as a function of the light's true AOP is plotted in Figure 6.50 for the full 180° range of possible AOPs. Four different levels of SNR are plotted for both 3- and 4-filter methods. The input DOLP is 0.5.

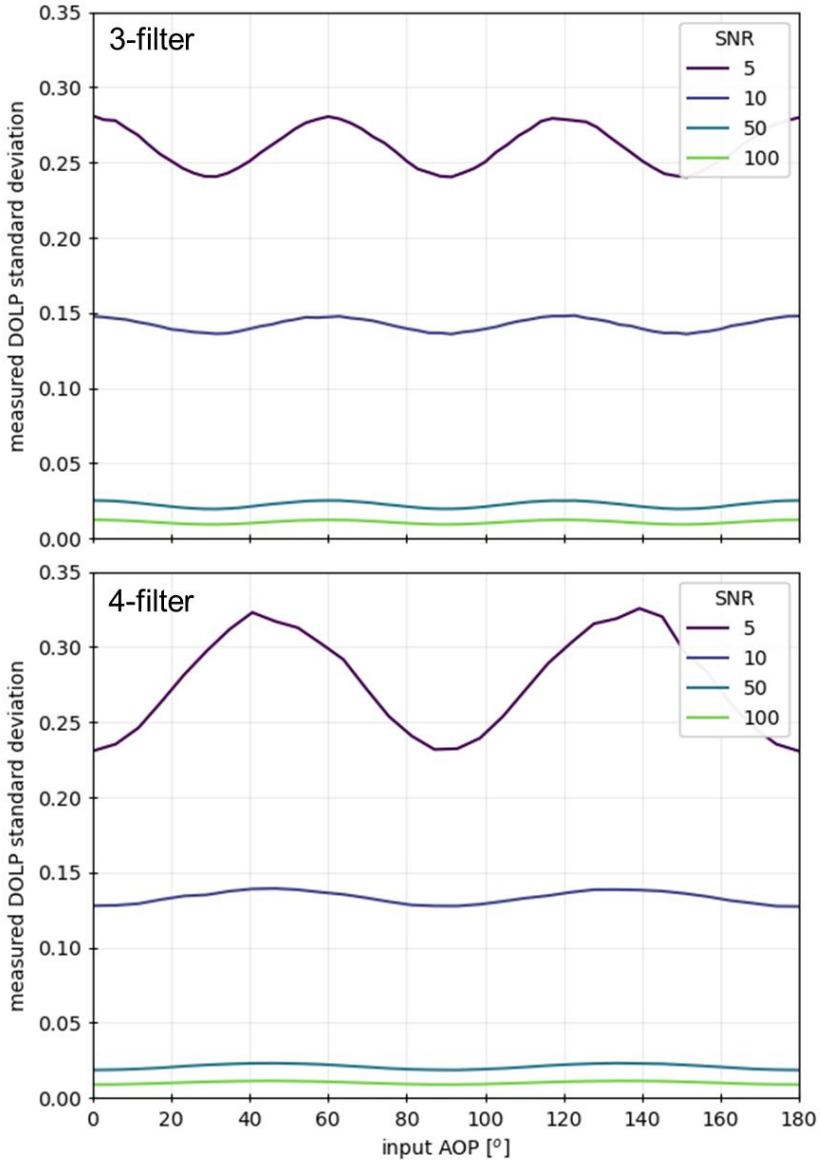

Figure 6.50: Standard deviation of measured DOLP as a function of light's true AOP and image SNR, for an input DOLP of 0.5.

It can be seen from Figure 6.50 that the DOLP measurement error varies periodically for both measurement methods. The 3-filter method's measurement error varies with a period of 60°, because of the 60° spacing of filters (a shift of 60° in AOP leaves the combination of measured intensities the same). Similarly, the accuracy of 4-filter measurements varies with a period of 90° due to the 90° spacing between filters. For a

given SNR, the amplitude of error variation is greater for the 4-filter method (reaching both lower and higher values than the error of the 3-filter method). For a given measurement method, the smaller the SNR, the larger the magnitude of the error variation. The mean value of the DOLP measurement error for a given SNR is comparable for both measurement methods.

6.5.5.3 AOP Measurement Error

The nature of the systematic error in the simulated AOP measurements is plotted in Figure 6.51.

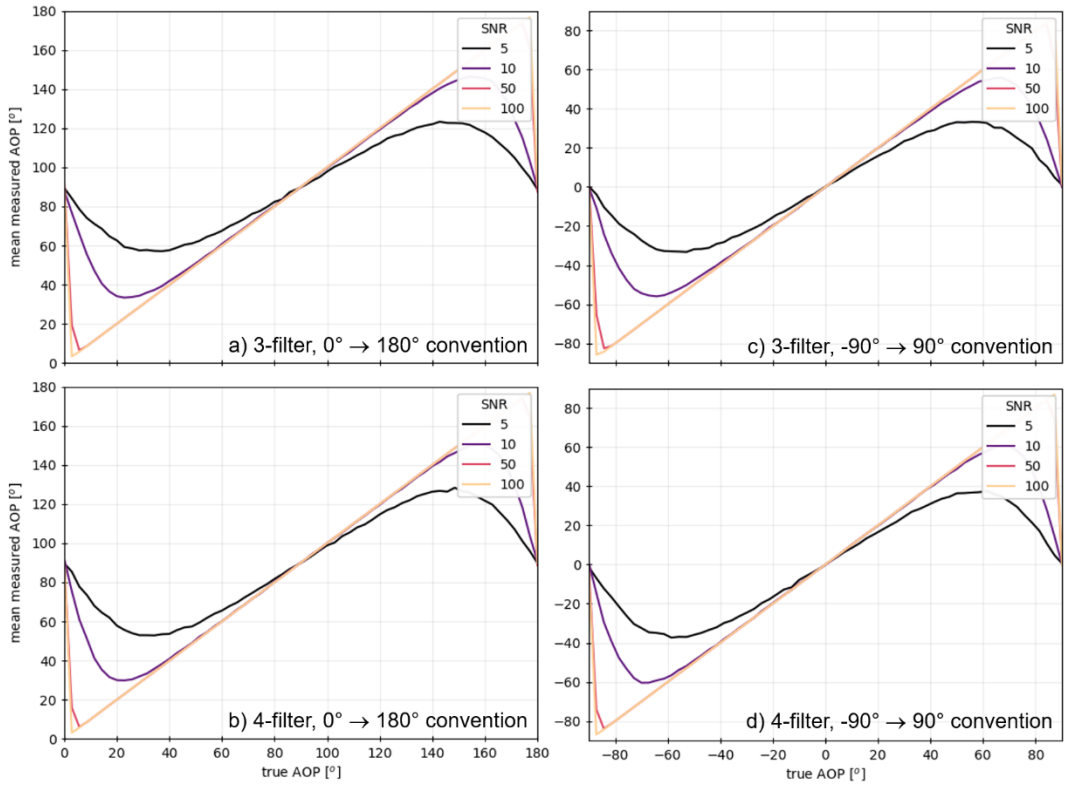

Figure 6.51: Mean measured AOP versus the light's true AOP for four different SNRs, both 3- and 4-filter methods, and two angular coordinate conventions whereby angles are represented by the cyclic ranges [0° → 180°] and [-90° → 90°] respectively. True DOLP = 0.3 for all plots.

For both the 3- and 4-filter methods, the mean measured AOPs of Figure 6.51 (a) and (b) equal the true AOP at a value of 90° (and a range of AOPs around 90° whose size depends on the SNR), but deviate significantly from the true AOP at and near values 0° and 180° (where the mean measured AOP approaches 90°). This is not an inherent feature of the measurement methods, but a result of the definition of the range of AOPs,

which has a discontinuity at 180°/0° (i.e. 180° = 0°, 181° = 1° etc.). This is demonstrated by the plots of Figure 6.51 (c) and (d), in which the convention has been changed to represent angles with the range -90° to 90°. The discontinuity is therefore moved to 90°/-90°, and as a result, the peaks in systematic error have moved from 0° and 180° to -90° and 90°.

Figure 6.52 plots the variation of the measured AOP's standard deviation as a function of SNR, for three different DOLPs and both 3- and 4-filter methods. The true AOP is arbitrarily 45°.

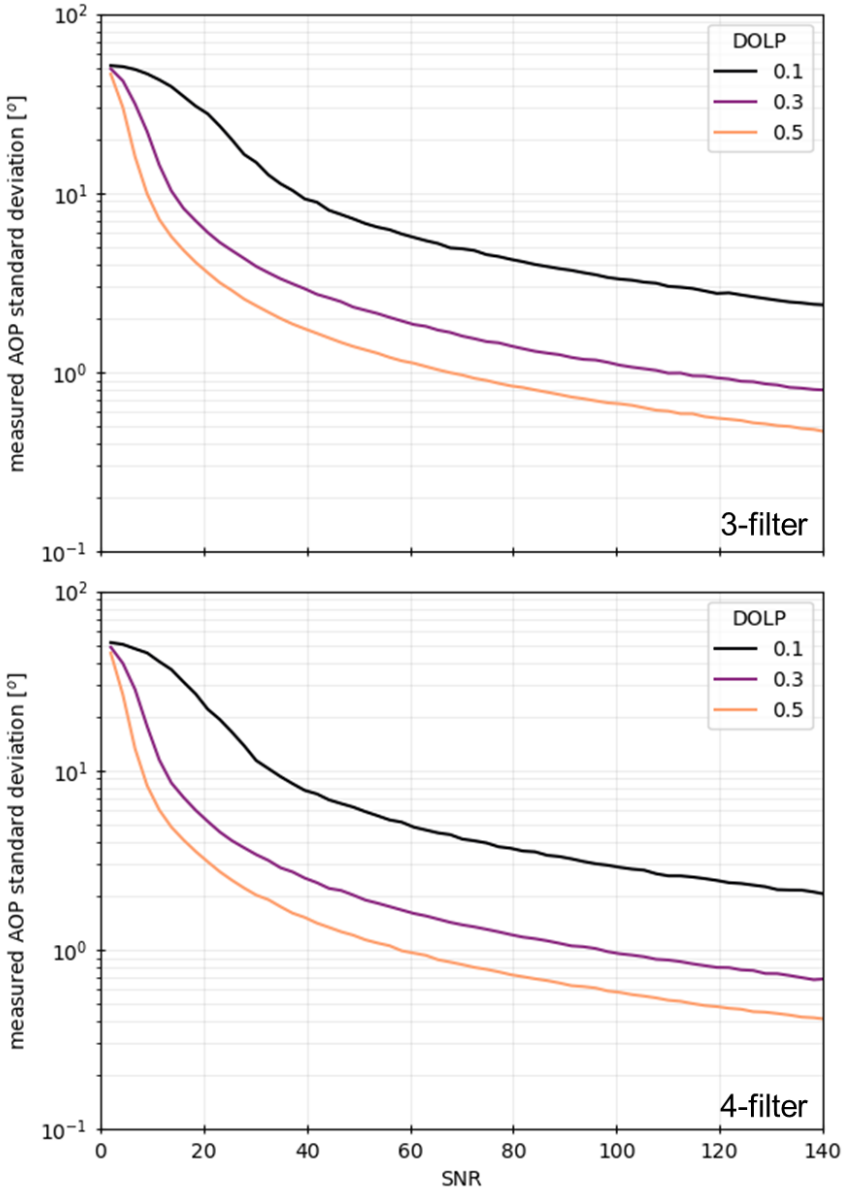

Figure 6.52: Standard deviation of measured AOP as a function of image SNR for three different DOLPs, and both 3- and 4-filter methods. True AOP = 45° for both plots.

Figure 6.52 shows very similar error behaviours for both 3- and 4-filter AOP measurements as a function of image SNR. Standard deviation grows rapidly at low SNRs, similarly to DOLP measurements (Figure 6.49). The measured AOP standard deviation increases as the input DOLP decreases (this is because the absolute intensity of the polarised light whose angle is being measured is lower for smaller DOLPs).

Finally, the behaviour of the measured AOP's standard deviation as a function of true AOP is plotted in Figure 6.53 for both 3- and 4-filter methods. In each case, four SNRs are plotted. DOLP is 0.3 for all measurements.

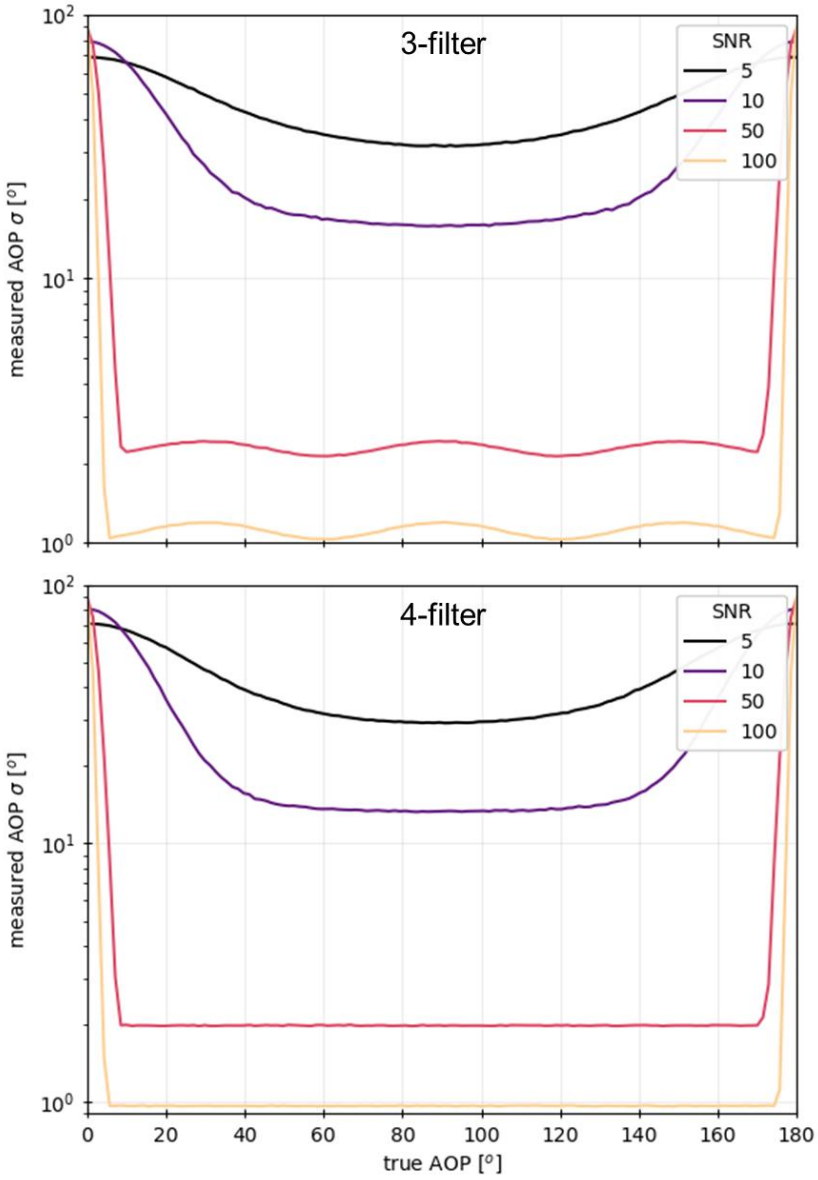

Figure 6.53: Standard deviation of measured AOP as a function of true AOP for four different image SNRs, and both 3- and 4-filter methods. True DOLP = 0.3 for both plots.

The standard deviations of Figure 6.53 peak at the 0°/180° AOPs (for the lower SNRs of 5 and 10, these peaks a very broad). This is the result of the discontinuity at 0°/180°, as discussed previously in relation to Figure 6.51. For SNRs of 50 and 100, the 3-filter method's standard deviation exhibits some significant periodic variation, which is not seen in the 4-filter methods. Once again, the 3-filter and 4-filter methods are found to behave very similarly, with no marked difference in their measurement accuracy.

6.5.5.4 Implications for Filter Selection

The motivation for comparing the measurement error characteristics of 3- and 4-filter methods is to determine whether either method is advantageous from a SNR and measurement accuracy perspective. This is driven by the very low expected SNRs (section 6.3), and the need to mitigate their impact on the scientific value of EnVisS data. If either method achieved significantly better measurement accuracy than the other, especially at low SNRs, it would be a strong argument in favour of its incorporation in the EnVisS design. However, both sections 6.5.5.2 (DOLP) and 6.5.5.3 (AOP) demonstrate that there is no appreciable difference in the achievable measurement accuracies of the 3- and 4-filter methods, or their behaviour as a function of SNR. EnVisS' imaging polarimetry is filter intensive, because unlike the camera's intensity observations, it requires multiple filters to obtain its end product. A large driver for the selection of the 3- or 4-filter method will be the smaller area on the detector required for the former, and the simplification this provides to the accommodation and mounting of EnVisS' filters. The fact that the 4-filter method provides no accuracy benefit helps to justify the use of the 3-filter arrangement.

6.5.5.5 Implications for Minimum SNR and Achievable Polarimetric Accuracies

Given that both systematic and random errors in the measurement of DOLP and AOP are related to SNR, this section discusses the implications this has on the likely accuracies EnVisS will be able to achieve, and the minimum requirements for the SNR of its polarimetry images. Having established that there is no significant difference between the errors of the 3-filter and 4-filter methods, and that 3-filter is preferable from a focal plane design point of view, the analysis of this section will be based on the 3-filter method.

With the light scattered from cometary dust generally having a DOLP between 0.0 and 0.3 (depending on the nature of the comet itself, and the viewing phase angle), it is necessary for EnVisS to be capable of measuring DOLPs within this range.

For EnVisS' observation of the comet's phase function, three main features will be important: the maximum magnitudes of its positive and negative polarisation branches, and the location of inversion angle, α_{inv}. Peak positive polarisation usually sits between 0.2 and 0.3 (though can be as low as 0.1 for gas rich comets, and has been observed to be as high as 0.5), whilst negative polarisation is likely to peak around -0.02. To achieve a DOLP measurement accuracy <50% of the expected negative branch peak will require SNRs of at least 120 (Figure 6.49), whilst SNRs will have to be higher still if the shape

of the negative branch and the location of α_{inv} are to be accurately determined (for example, see Kiselev et al. (2005) and their Figure 2). The larger and broader positive polarisation branch could be characterised with a larger DOLP measurement error than the negative polarisation branch (for example, a 10% error in the measurement of its peak polarisation could be achieved with SNRs ≈50).

Achieving these required SNRs will require use of the techniques described in section 6.4, particularly pixel binning. Angular resolution on polarimetric images will be in the region of 1° (for the brightest scenes) to 10° (for the more modal, dimmer views), whilst full resolution imaging polarimetry will not be possible. With these angular resolutions, and throughout the flyby, spatial resolutions at the nucleus and inner coma will vary from 18,000 km down to 18 km. Polarimetric telescope observations generally have spatial scales of 100-1000 km per pixel (e.g. Hadamcik et al. (2007)). Near closest approach, even if maximum binning levels were required, polarimetric images with spatial resolutions a factor of 2-10 better than those of telescope observations should be obtainable by EnVisS.

6.5.6 Approximating the OPE Measurement Method

In addition to its imaging polarimetry, EnVisS will also be capable of performing an approximate version of the photopolarimetry carried out by Giotto's Optical Probe Experiment (OPE) at comet Halley (see section 4.4.6). The OPE instrument measured brightness over a 3° diameter FOV, pointed parallel to the spacecraft's trajectory in the trailing direction. Polarised intensity measurements were captured at angles of 0°, 45°, 90° and 135° with a single fixed polariser, using the continuous rotation of the spacecraft to reorient the polariser. Eight spectral bands were measured simultaneously, four continuum and four gaseous emission.

EnVisS will view along spacecraft B2's trailing spin axis, and if mounted with the un-tilted configuration, will also view along the forward spin axis. In each case, a cone of similar diameter to OPE's FOV can be integrated using a small portion of EnVisS' detector, as illustrated in Figure 6.54.

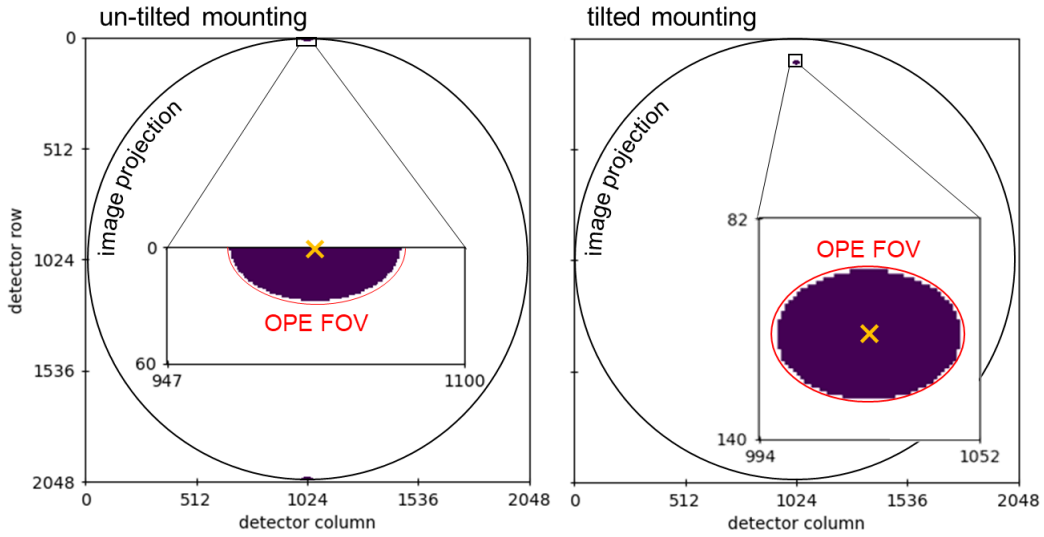

Figure 6.54: Illustration of OPE's FOV on EnVisS' detector for both possible camera mounting configurations. Purple regions indicate EnVisS pixels which share OPE's FOV. The yellow cross marks the position of the spin axis.

If a polarising filter is positioned over the portion(s) of the detector viewing the spin axis, then it captures that portion of the scene with a full range of orientations, 0°-180°, every half rotation of the B2 spacecraft. In line with the OPE measurement method, a single intensity measurement is obtained with an exposure time equal to an eighth of the spacecraft spin period. The pixels corresponding to the desired FOV are read out and binned to produce the measurement. Capturing four consecutive measurements in this way provides observations at 0°, 45°, 90° and 135°.

This method lacks the ability to resolve spatial structure in the polarisation, or to characterise the polarisation phase curve. However, on the premise of very low expected signals and SNRs, implementing the OPE measurement method with EnVisS is motivated by a need to maximise measurement signal. With the OPE method, exposure times can be significantly longer than with imaging polarimetry, even when compared to highly binned images. The 3° OPE FOV also allows a significant SNR gain from pixel binning. As a result, a significantly greater SNR can be achieved than is possible with any form of EnVisS' imaging polarimetry. Figure 6.55 plots the approximate expected measurement SNR when using the OPE polarimetry method, as a function of coma brightness. For comparison, the SNR is plotted for the EnVisS broadband filter as well as two of OPE's continuum filters.

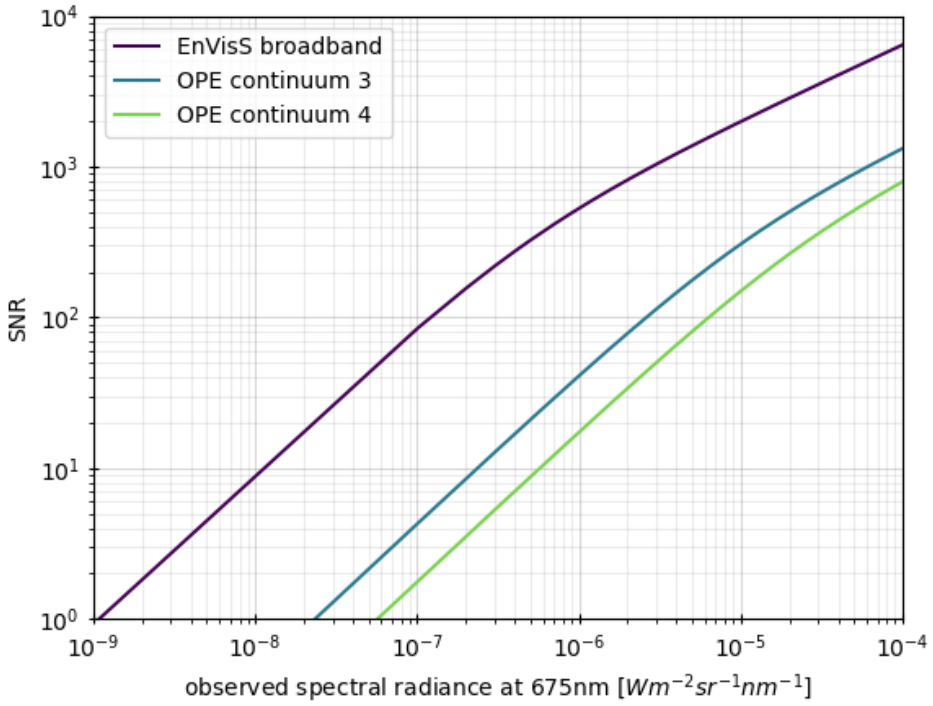

Figure 6.55: Approximate SNR of EnVisS OPE-like measurements (see main text) for three filters: the nominal EnVisS broadband filter and two of the continuum filters used on the OPE instrument.

Contrasting Figure 6.55 with Figure 6.25 reveals how significantly higher the OPE-like observations' SNRs are than those of EnVisS images. SNRs greater than 10 are achievable at brightnesses as low as 10^{-8} W m^{-2} sr^{-1} nm^{-1}, whilst SNRs greater than 100, which are necessary for measuring DOLP with an absolute accuracy greater than 0.01, will be obtained for brightnesses greater than 10^{-7} W m^{-2} sr^{-1} nm^{-1}. The OPE measurement method therefore extends the range of brightnesses over which polarimetric observations can be made, and increases the likelihood that a comprehensive set of scientifically valuable polarimetric observations can be obtained. Although the EnVisS implementation of the OPE measurement method would lack some of the original instrument's capabilities (specifically, the multispectral observations), it would provide the unique opportunity to use the measurement method to observe a dynamically new comet.

6.5.7 Summary

The intended polarimetric products from the EnVisS instrument will be unique in both their form and their scientific value. Particularly valuable will be the characterisation of the comet's phase function over close to the full 180° range, and first time observations

of large phase angles (>120°). Imaging structural variation of the polarisation from a viewpoint within the cometary environment, achieving unprecedented spatial and temporal resolution, will also be of high value. However, the low signals expected in EnVisS' images are a significant challenge, and do limit the scope of what can be achieved. Polarimetric measurement error is very sensitive to low signals, and SNRs in EnVisS' polarimetric images will need to meet higher standards than its intensity images. Error modelling indicates that SNRs will need to exceed ≈50 at a minimum, and accurate phase curve measurement will require SNRs on the order of 100. This will necessitate the use of significant pixel binning. Greater SNRs can be achieved through the OPE photopolarimetry method, but for a more limited set of measurements.

6.6 Discussion

This chapter has described the EnVisS imager, under development for ESA's Comet Interceptor mission, and has presented detailed analysis of the instrument that was carried out for this thesis, but has also supported the development of EnVisS. The objective of EnVisS is to capture images of a comet's large scale environment and structures during a spacecraft flyby, in order to link extended coma features and their evolution to activity at the nucleus. Push-frame imaging is used, facilitated by the spacecraft's spin stabilisation, to continuously scan a 180° FOV over the sky and build up all-sky images (images which independently contain the entire visible sky as seen by EnVisS).

The operation and capabilities of the camera are largely dictated by the imaging mechanism, which in particular imposes a maximum exposure time of 1-3 ms (depending on spacecraft spin rate), as is described in section 6.2. Analysing previous cometary observations and combining these with physical models of cometary comae and EnVisS, it is shown that imaging conditions under which EnVisS will operate are challenging, particularly with regards to image signal strength (section 6.3). A number of imaging and data processing techniques, adapted to the unique nature of EnVisS' operation, were analysed for their proficiency in mitigating low image signals (section 6.4). SNR can be controlled through the use of these techniques, and significant increases to signal can be achieved, but generally with a profound effect on the data form, mostly at the cost of image angular resolution. Use of one or more of these techniques will be required for all of EnVisS' images, with observations made at large spacecraft-nucleus distances likely to need the most, and images captured near closest approach more likely to be able to image with higher spatial resolutions, and less-intrusive measures to boost SNR. Sufficiently high SNRs are expected to be achievable for at least limited parts of the flyby

for both intensity and polarimetric measurements, though the former is limited to broadband spectral ranges, whilst error modelling shows the latter will be restricted to coarse angular resolutions and non-imaging photopolarimetry (section 6.5).

The analysis of EnVisS carried out for this thesis contributed significantly to the identification and characterisation of these challengingly-low expected SNRs, and it demonstrates the importance and value of modelling an instrument and its target in order to determine if any such issues exist, and whether they impact the instrument's ability to meet its science goals. Where possible, this allows the study and implementation of design modifications, observation strategies and data processing techniques to obtain suitable SNRs. For a new instrument, like EnVisS, whose imaging approach has minimal heritage, there is no guarantee that the original intended objectives of the instrument can all be met when such issues are encountered. Indeed, multispectral imaging of gaseous coma species and the ion tail were goals of the initial EnVisS concept, but are accepted now not to be possible in light of the camera's low signals. The techniques used in this chapter could be applied to any camera concept to assess the capabilities of its design, and to explore the impact of design modifications on performance. Equally, the fundamental concept of EnVisS – obtaining all-sky images of a complex and extended environment – could be applied to other missions, and realised with different imaging techniques.

As an example, a change in spacecraft design would facilitate an adapted version of EnVisS to obtain similar observations, both spectral and polarimetric, but with larger signals. Specifically, mounting EnVisS on a 3-axis stabilised spacecraft would remove the rotationally-induced motion blur that limits the instrument's exposure time (Figure 6.56). Use of multiple filters would be achieved by the inclusion of a mechanical filter wheel. A single FOV would obtain images covering half the sky, whilst an additional second FOV would increase coverage to the entire visible sky (with a blind spot caused by the spacecraft's structure).

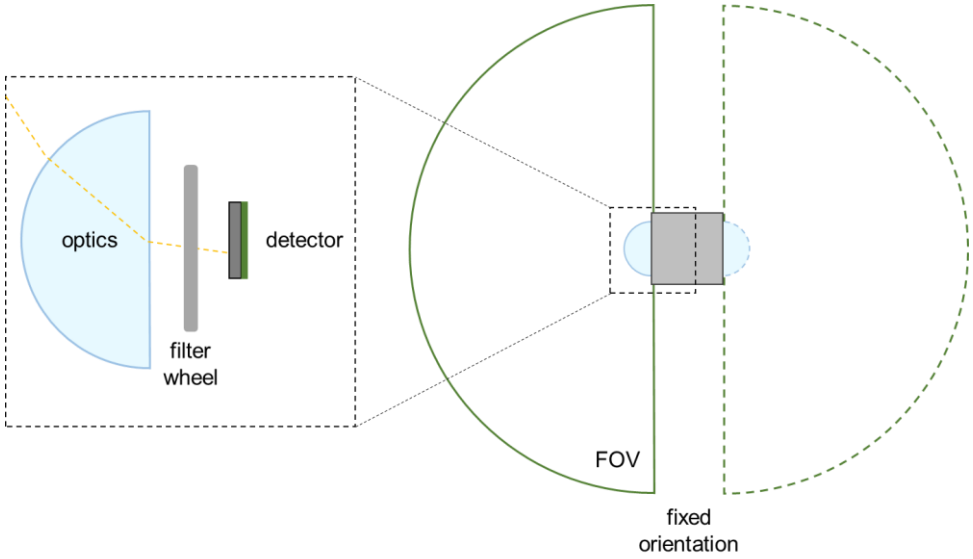

Figure 6.56: Illustration of an adapted implementation of EnVisS' all-sky imaging. Fisheye optics with a 180° hemispherical FOV (solid green semicircle) image through a filter wheel onto a detector to achieve hemispherical coverage of the sky. An optional second set of hemispherical-FOV optics (dashed green semicircle) imaging onto another detector (or the same detector via fibre optics, see Tomasko et al. (1999)) increases coverage to the entire visible sky.

With the instrument mounted on a stabilised platform, exposure times would be limited by flyby motion, and the values possible would range from 10 ms near closest approach to as high as 100 s at the earlier stages of the flyby (from section 6.2.2.1). With these larger exposure times, which are comparable to those used by other cometary cameras such as Rosetta's OSIRIS, higher-resolution higher-SNR images could be obtained (>100 when observing a brightness of 10^{-8} W m^{-2} sr^{-1} nm^{-1}). This would facilitate observations with narrower spectral ranges, allowing the observation of cometary ions and neutral gas. Other design factors, such as the complications of a filter wheel and two sets of optics, are also important considerations, but this adapted design could be a valuable inclusion on a future mission to a comet or other planetary body, and a worthy subject of further study.

7 Supporting Camera Design with Image Simulation

The purpose of this penultimate chapter is to further describe the details of the image simulation techniques used to facilitate much of the research presented in chapters 5 and 6. Both of these chapters dealt with novel conceptual imagers, for which there is no available hardware to produce observations of any kind. However, even in the very early study of an imaging concept, it is valuable to have access to example images, representative of the data that the instrument is intended to obtain. Indeed, a number and variety of such simulated images are used throughout chapters 5 and 6. This chapter will therefore provide a description of how these images were simulated. Section 7.1 introduces some background on image simulation, followed in section 7.2 by an overview of the approach to image simulation taken in this thesis. Sections 7.3 and 7.4 describe the methods employed to simulate images of planetary surfaces and cometary dust respectively. Specific code will not be presented (though the raw code can be accessed as described in the Appendix), but rather a broader description of the rendering method will be described, to give an idea of the requirements for, and capabilities of, a simple planetary image simulator. The approach can be readily implemented in a number of languages, and adapted to the specific needs of a range of planetary imaging studies.

7.1 Background

7.1.1 Motivation for Simulating Planetary Observations

Image simulation (whereby a 2D image is synthesised, or rendered, from a scene constructed of a set of objects or data) is used for a variety of applications in planetary science, including but not limited to the visualisation and study of planetary surfaces (e.g. software such as Pro3D (Barnes, et al., 2018)), education and outreach (e.g. planetaria (Hobson, et al., 2010)), and science operations planning (Pérez-Ayúcar, et al., 2018). Each of these applications imposes different requirements on the way in which simulated images are rendered. For example, outreach and visualisation images need to be clear and engaging, whilst the physical accuracy of their appearance may be less important.

Whatever the end application, the motivation for simulating images is that views which are otherwise unavailable can be produced. A planetarium's view as the audience flies past Jupiter for example, or the ability of a geologist to view a Martian outcrop from any direction. In the case of this thesis, the motivation of image simulation was to assist the study of camera concepts and their capabilities by producing data representative of what they would observe. This facilitated, for example, the testing of stereophotogrammetry techniques that could be applied to descent images (Brydon, et al., 2021), and the characterisation of signal to noise ratio in EnVisS' images. Similar image simulation has been utilised in previous studies of planetary imagers (e.g. Kirk et al. (2016)). It is worth noting also that image simulation need not be strictly limited to the computer generation of synthetic images. Malin et al. (2001) used simulated Martian descent images (simulated in the sense that they were captured with a real camera, mounted on a helicopter at Earth, but intended to replicate a descending spacecraft's view of Mars) to study descent imaging (similarly the drone data of chapter 5, section 5.5.2). This chapter will however focus on the simulation of images by computer rendering.

7.1.2 Rendering Images

Rendering an image is the process of generating a camera's view of a scene, given a description of the camera's viewing geometry and a description of the scene's appearance. The process comprises two major steps (see Figure 7.1): determining the visibility of the scene (i.e. which parts of the scene do the camera's pixels observe) and determining the appearance of the scene (i.e. what signals do the pixels receive). Scene visibility is a function of the 3D structure of the scene and the viewing geometry of the camera, whilst scene appearance depends on the nature of the objects in the scene (e.g. their material, colour) and the lighting of the scene. Approaches to rendering images can vary from very simple to highly complicated, and computing power has historically always placed an upper limit on what can feasibly be rendered (though this of course continually increases with the advancement of computing techniques and hardware). Here, two major methods by which image rendering is commonly achieved will be discussed: ray casting (and tracing), and rasterisation.

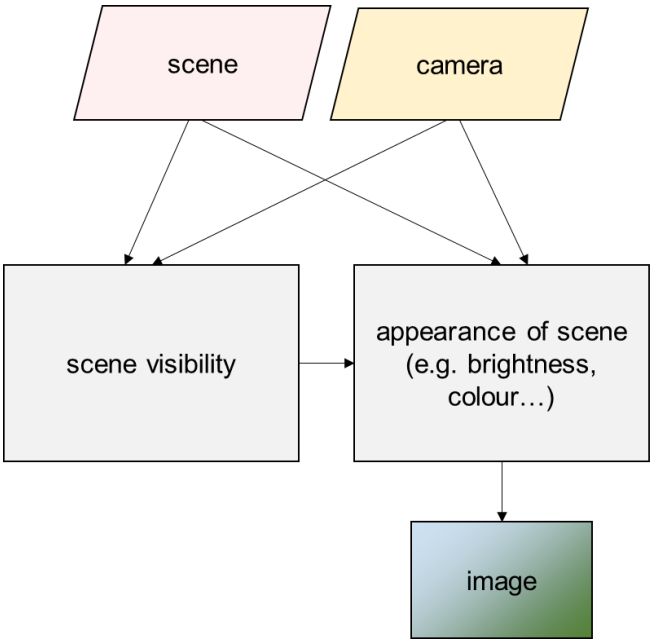

Figure 7.1: An illustration of the major features of image simulation.

7.1.2.1 Ray Casting and Tracing

Ray casting determines the visibility of scene objects to camera pixels by following a straight line (a ray) out of the camera and into the scene along the pixel's line of sight. Rays are directed according to the camera's outward mapping function (see section 2.1). Objects within the scene are tested for intersection, and those with which the ray intersects coincide with the pixels' view direction (see Figure 7.2). If a ray intersects multiple objects, the depths of intersection are used to determine which objects are occluded. The minimum time required to render an image using ray casting is proportional to the number of image pixels and the number of objects in the scene.

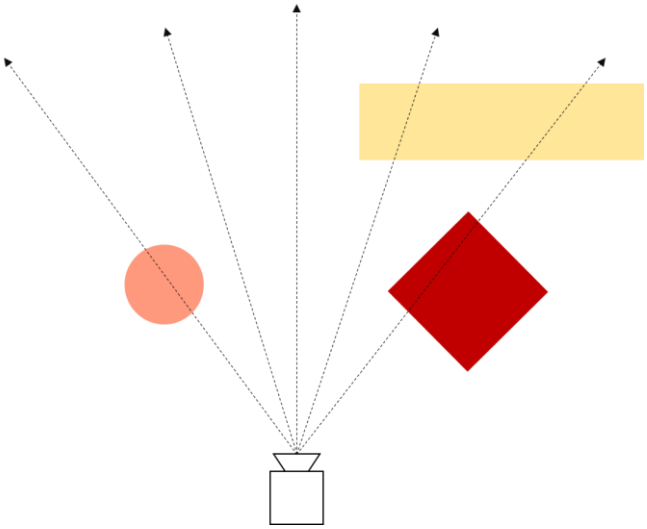

Figure 7.2: Illustration of the ray casting concept.

The fundamental principle behind ray casting is mathematically testing each pixel's ray for intersection with every object in the scene. The technique was first described by Roth (1982), in the context of an interactive computer aided design solid modelling system. Roth presented an approach whereby objects are constructed from one or more 'primitive' solids such as spheres, cubes and cylinders. Crucially, the algebraic surfaces of these primitives allow ray-surface intersections to be symbolically rather than numerically solved, helping to minimise computation time.

Roth's approach of using combinations of small numbers of primitive shapes to describe objects has the advantage that it keeps the total number of required intersection tests small, but it limits the complexity and variety of shapes that can be modelled. More complex surfaces are commonly modelled as tessellated meshes of triangles or polygons (Snyder, 1987). Testing a ray for intersection with such a surface requires determining which, if any, of the mesh's triangles the ray intersects.

Testing a ray for intersection with a triangle can be used to calculate a range of properties, such as a Boolean value indicating whether an intersection occurs, the intersection's depth (the distance from the camera to the point of intersection), the 3D position of the intersection, and the angle between the intersecting ray and the triangle's surface normal, all of which can be used to determine what colour the ray's pixel should display. Strictly speaking, ray casting refers to the techniques used to determine which regions of the scene a pixel views, whilst the accurate calculation of the signal (i.e. brightness, colour) is often achieved with a technique known as ray tracing (e.g. Pharr et al. (2016); Whitted (1980)).

Ray tracing is a recursive process: a ray is followed out from a pixel until it intersects an object. From the point of intersection, new rays are traced in directions corresponding to reflection, refraction or scattering (depending on the nature of the object), and rays are traced toward light sources to test for illumination and shadowing. Each of these rays is treated in the same way as the original ray, and is itself subject to reflection and scattering. Physical models of light propagation, scattering and reflection are used to calculate the nature of the light seen by each ray, and to inform the directions in which recursive rays should be traced. Processes such as diffuse reflection require multiple rays of different directions to be sampled, and a single pixel may require 100-1000 rays to be traced.

The rigour of ray tracing is well suited to accurately replicating the many physical processes (reflection, refraction, scattering, shadowing, emission) that determine a scene's appearance (Boulos, et al., 2007; Pharr, et al., 2016). However, checking rays for every pixel against every scene object quickly becomes expensive as image resolution and scene size grow, and leads to large render times. Scene objects are often organised into hierarchical tree structures, allowing significantly faster ray-scene intersection tests (Smits, 1998). However, as optimisations and acceleration structures are invoked, the implementation in code of a rendering program becomes more complex. Furthermore, even with optimisations, physically-based ray tracing remains computationally expensive, and a single detailed image can take minutes to render.

7.1.2.2 Rasterisation

Rasterisation is an image rendering method historically favoured for applications requiring short render times, particularly where a user needs to interact in real time with the rendered images. The principle behind the technique is the projection of each object in a scene onto the image plane, using the camera's inward mapping function (see section 2.1) (Garachorloo, et al., 1989; Pineda, 1988). Because the projection is not performed on a pixel-by-pixel basis, and each object in the scene only has to be projected once, the technique requires fewer computations than ray tracing and allows high-speed rendering of images. The concept is illustrated in Figure 7.3.

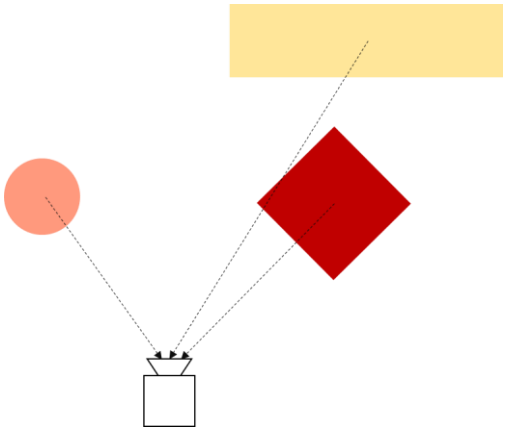

Figure 7.3: Illustration of the rasterisation concept.

As with ray tracing, rasterisation represents scene objects with primitive shapes, whose projections to 2D image coordinates can be algebraically calculated. Generally, polygons are mapped to image coordinates simply by the projection of each of their vertices, and as with ray tracing, triangular meshes are commonly used to describe the scene (e.g. Kugler (1996)). After rasterisation has been used to calculate the 2D position and shape of scene objects on the image plane, the corresponding colours of image pixels are determined by a technique called shading. Colours may be mapped directly from an image associated with the scene object, or physically calculated to account for effects such as illumination and reflection.

7.2 Rendering Approach

The development of a simple Python-based rendering method was motivated by the need to simulate the observations that the instruments discussed in chapters 5 and 6 would collect. This was particularly important given that both the instruments are very different to any previous space-borne cameras, and no representative images are available.

Because of the large number of computations involved in image rendering, and Python's relatively slow execution of extensive numerical calculations (due mainly to it being an interpreted, rather than compiled language), it would traditionally not be the chosen language for implementing an image renderer (Jun & Ling, 2010). However, the language is popular and widespread amongst scientists throughout all fields of research, including the planetary science community, thanks to its quickly and easily readable and writable syntax, a wide range of well-developed scientific libraries (e.g. NumPy (Harris, et al., 2020), SciPy (Virtanen, et al., 2020), Astropy (Robitaille, et al., 2013)), and its free

availability (Millman & Aivazis, 2011). Indeed, Python was used throughout this thesis' research to perform all the analysis presented in chapters 5 and 6. In order to ensure compatibility of the simulated data and all other analysis code, the decision was made to implement image simulation in Python. This decision also keeps the image simulation code readily modifiable, reusable and accessible.

From the outset, the image simulation method was developed such that it satisfied two major requirements:

- Easily implemented in Python, and relying only on additional libraries which are both commonly used in scientific computing and widely compatible with computers capable of running Python.
- Capable of rendering low detail images in <10 s, so that image simulation setups can be quickly assessed, and capable of rendering full detail images <600 s

Developing image simulation software was never the main focus of this thesis, but was required in order to facilitate much of its research. For that reason, this chapter will not describe all the details of the developed code, but will instead outline its general structure, approach and capabilities, so as to inform the use of similar techniques in other instrument studies. One of the main novelties of the rendering approach is its implementation in Python, and methods by which render times are kept short in spite of the language's slow execution are discussed. However, the specifics of using Python are generally not discussed, as the rendering techniques described in this chapter could be implemented in any language.

7.3 Simulating Planetary Surface Images

The first requirement for image simulation came from the penetrator descent camera, and a desire to produce examples of what its images would look like. This was followed by the need to produce high fidelity simulations of its descent images suitable for testing analysis techniques (e.g. Brydon et al. (2021)). In order to meet the needs of the penetrator descent camera study, there was a requirement for the following image rendering capabilities:

- Replicate the geometry, coverage and viewing perspectives of the penetrator descent camera.
- Simulate high resolution views of planetary surfaces.
- Be capable of simulating data with sufficient realism that they can be used to test image analysis techniques.

To achieve simulations representative of real planetary observations, the software utilises data obtained from real surfaces. Specifically, surface models are constructed from digital terrain models (DTMs) and orthorectified images (ORIs), which are products commonly produced by planetary imaging studies. An advantage of this approach is that DTMs and ORIs exist for a wide range and large number of solar system bodies, including all the terrestrial planets, many moons, some dwarf planets and several asteroids and comets. For well-observed bodies, such as Mars and the Moon, high-quality global datasets exist. A further advantage of DTMs and ORIs is that, because they are produced and used for scientific studies, their accuracy and fidelity is quantified, making them a suitable tool for physically accurate simulations.

DTMs and ORIs store planetary surface data (its elevation and appearance respectively) in a 2D frame, whose axes represent directions in a coordinate system linked to the planetary body. Generally, horizontal and vertical axes represent positions of longitude and latitude respectively, or distances from a reference location measured along lines of constant latitude and longitude respectively. A reference surface, or datum, provides a 'sea level' for the planetary body, from which DTMs' elevation values are measured. In its simplest form, the reference surface is that of an ellipsoid, but more complex equipotential surfaces are also used (e.g. Ardalan et al. (2010)). DTMs and ORIs may use planetocentric or planetographic coordinate systems (Figure 7.4), as well as a range of map projections (e.g. equirectangular, sinusoidal).

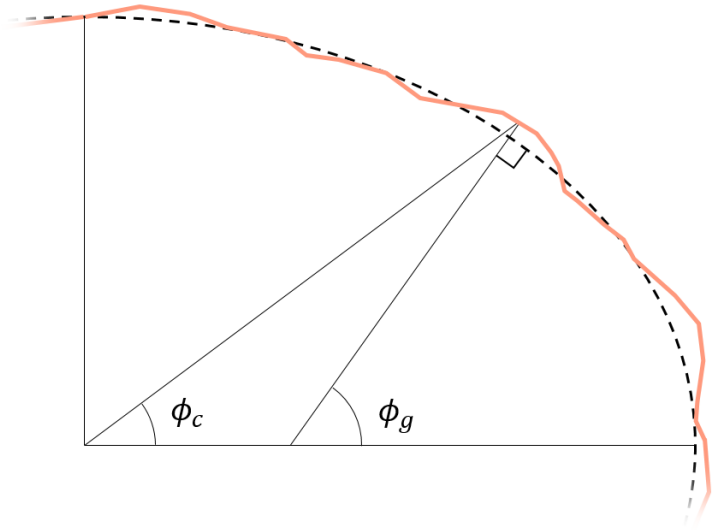

Figure 7.4: An illustration of the difference between planetocentric (ϕ_c) and planetographic (ϕ_g) latitudes and radii. The dashed line represents a reference surface ellipsoid, the orange line a planet's physical surface. The coordinates of a surface point can differ significantly depending on the scheme used.

DTMs comprise an ordered set of data points, each one storing an elevation value for a point on the surface (other information can be stored, too, but only elevation is used and considered here). Accompanying metadata defines the scale and bounds of the grid's horizontal and vertical axes (for example the minimum and maximum longitudes and latitudes of the DTM). Multiple schemes exist for the arrangement and storage of a DTM's data points, with the two most popular being regular-grids and triangulated irregular networks (TINs). Regular-grid DTMs are favoured for their simpler form, lower data volume and wider compatibility with analysis software, and are used much more commonly than TINs in planetary science. For this reason, only regular-grid DTMs will be considered here, and the simulation approach described in this section is only compatible with this format. It is worth noting though that TINs and other DTM formats can readily be interpolated to regular-grids (Kang, et al., 2015).

Each data point, or pixel, of a regular-grid DTM stores an elevation value (often mapped to a greyscale value when DTMs are visualised as images, as in Figure 7.5), which in conjunction with the pixel's location in the 2D grid defines a 3D position according to a mapping from DTM coordinates (latitude ϕ, longitude λ and elevation e) to 3D Cartesian coordinates x, y ad z:

$$\begin{bmatrix}\phi\\ \lambda\\ e\end{bmatrix} \rightarrow \begin{bmatrix}x\\ y\\ z\end{bmatrix} \tag{7.1}$$

As an example, a spherical reference surface with radius R has a planetocentric mapping function

$$\begin{bmatrix}x\\ y\\ z\end{bmatrix} = \begin{bmatrix}(R+e)\sin\left(\frac{\pi}{2}-\phi\right)\cos\lambda\\ (R+e)\sin\left(\frac{\pi}{2}-\phi\right)\sin\lambda\\ (R+e)\cos\left(\frac{\pi}{2}-\phi\right)\end{bmatrix} \tag{7.2}$$

where the reference surface is centred on the Cartesian origin, the x-axis intersects the surface at $[\phi \quad \lambda] = [0 \quad 0]$, and the y-axis also intersects the equator. A DTM's grid of data points, mapped to 3D space, describes a regular triangular mesh (where each triangle is a primitive) which approximates the shape of the associated planetary surface (as is described in section 7.1.2.1).

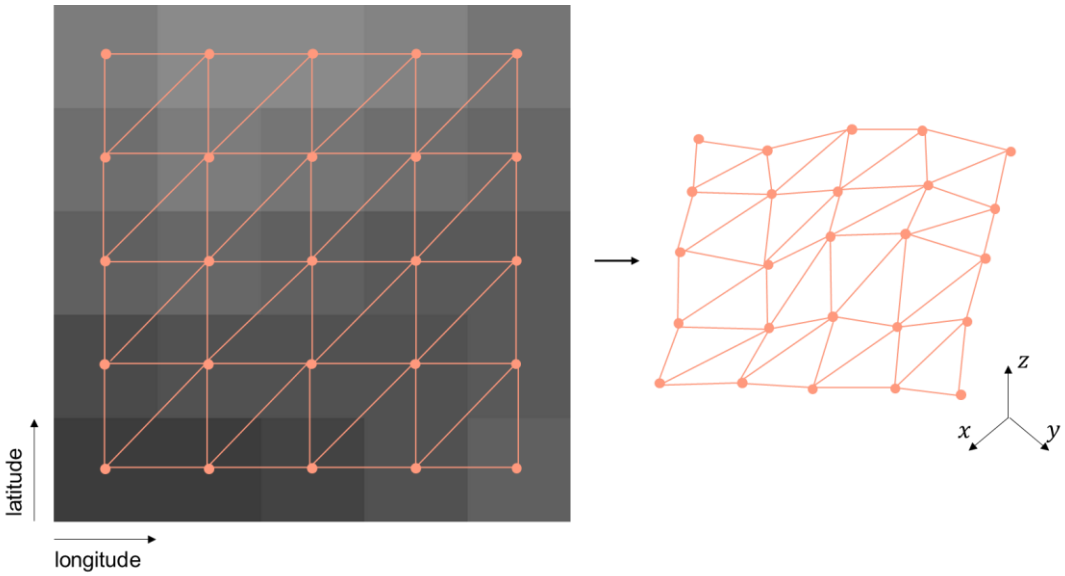

Figure 7.5: An illustration of the formation of a regular triangular mesh in 3D space from a 5x5 pixel DTM.

DTMs representing planetary surfaces can cover large areas, can have small physical grid spacing, and indeed can have a combination of the two. It is therefore common for DTMs to contain millions of vertices, and likewise their associated surface meshes to contain millions of triangles. Given the requirement to keep render times short in spite of

Python's relatively slow execution, rasterisation was therefore chosen over ray tracing as the method of generating images of surface models.

As described in section 7.1.2.2, rasterisation requires mapping 3D scene coordinates to 2D image coordinates, according to the simulated imager's camera model (see section 2.1). Because this step involves the repeated application of a mapping function to a large number of vertices, professional rasterisation software (e.g. the graphics engines of computer games) commonly performs this using a computer's GPU (Graphics Processing Unit), whose hardware is optimised for parallel computations (Laine & Karras, 2011). Software exists for accelerating Python code by using GPU computing, but it is not widely accessible or compatible with the majority of computers, and using it would have increased the complexity and reduced the reusability of the rendering software. Instead, accelerating the image simulation and achieving acceptably short render times can be achieved by utilising the NumPy library for Python (Harris, et al., 2020).

NumPy is designed to speed up Python code that handles element-wise calculations with sets of numbers. Consider a set of scalars $a_1, a_2, \dots, a_n$, a set of identically sized vectors $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_m$, and a vector function $\vec{f}(a_1, a_2, \dots, a_n, \vec{x}_1, \vec{x}_2, \dots, \vec{x}_m)$. The calculation of function $\vec{f}$ is element-wise if

$$f(a_1, a_2, \dots, a_n, \vec{x}_1, \vec{x}_2, \dots, \vec{x}_m)_i = \vec{f}(a_1, a_2, \dots, a_n, x_{1i}, x_{2i}, \dots, x_{mi}) \tag{7.3}$$

A simple example of an element-wise calculation is

$$a\vec{x} + \vec{y} = \begin{bmatrix} ax_1 + y_1 \\ ax_2 + y_2 \end{bmatrix} \tag{7.4}$$

Performing an element-wise calculation in Python requires iterating through all vector indices i, and calculating f using the associated values of the input vectors (i.e. the right hand side of equation 7.3). This iteration, potentially through millions of values, is inefficient in Python, and the source of the speed bottleneck (Jun & Ling, 2010). NumPy addresses this by performing the iteration in a faster language, C, and then returning the calculated values to Python, resulting in significant speed increases. NumPy automatically handles the conversion of values to C, the iteration, and the returning of values to Python, keeping code simple. Effectively, in the Python environment the multiple calculations can be considered as being performed in parallel, and the process is referred to as being vectorised. Figure 7.6 gives an example of the code difference between pure Python and NumPy vectorised calculations. It is important to note that the 'vectors' used in vectorised NumPy calculations need not strictly be vectors (i.e. one-

dimensional lists of values) but in fact can be arrays of values with any shape (e.g. 2x2 or 3x4x100), as long as all ‘vectors’ involved in a calculation are of the same shape and size.

```
def f(a, x, y):
    return a * x + y


# Iteration in pure Python
data = [0, 1, 2, 3, 4, 5]
for datum in data:
    print(f(2, datum, datum))
# 0
# 3
# 6
# 9
# 12
# 15

# Iteration using NumPy
data = numpy.array([0, 1, 2, 3, 4, 5])
print(f(2, data, data))
# [0 3 6 9 12 15]
```

Figure 7.6: The implementation of equation 7.4 in Python using both pure Python and NumPy.

The vectorised rasterisation code is capable of mapping surface meshes to image coordinates significantly faster than can be achieved with pure Python, as shown in Figure 7.7. The absolute time taken to render an image, and the speed increases yielded by vectorisation, vary significantly depending on the nature of the scene and the camera, but use of vectorisation generally increases rendering speed by a factor of 20-100. Without the use of NumPy, render times become too large as vertex number approaches 10^7.

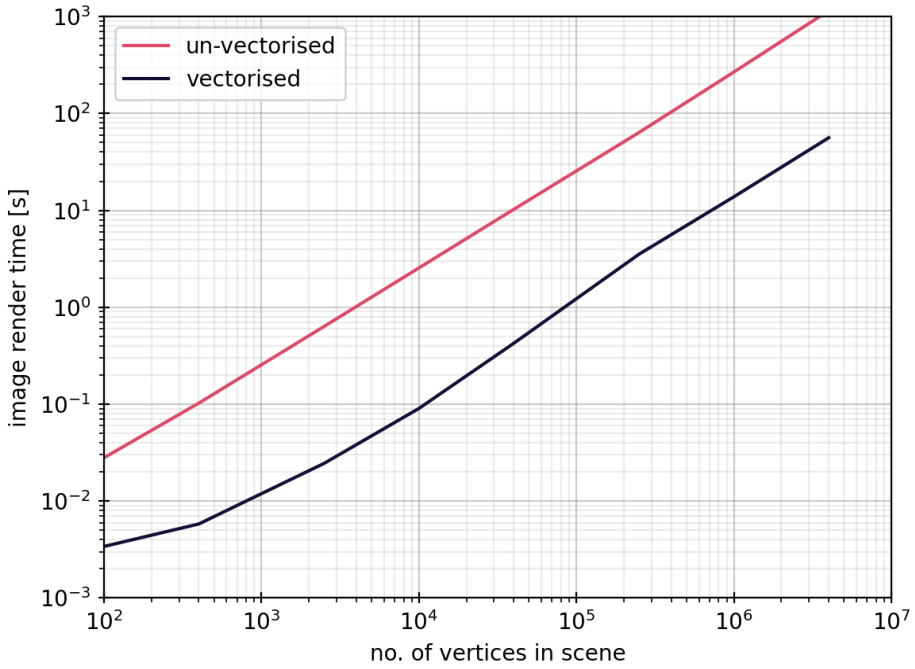

Figure 7.7: Time taken to render images as a function of the number of vertices in the scene, for both un-vectorised (pure Python) and vectorised (NumPy) methods.

After the mapping of the triangles of a surface mesh to image coordinates has been calculated, it is still necessary to determine the resulting signals that the pixels record (a step referred to as ‘shading’ in image rendering literature). For the simulated surface images used for this thesis (for example the descent images used in the stereophotogrammetry of section 5.6 and Brydon et al. (2021)), image-based shading is used.

Image-based shading maps surface colours directly from a texture image to the rendered image. For this thesis’ work, ORIs are used as the texture images from which to determine surface appearance. An ORI uses the same grid layout as DTMs, but its pixels store the appearance, rather than the elevation, of surface points. They are generally produced from one or more surface images captured from a spacecraft, which are reprojected from the camera’s image coordinates to a global frame (removing viewing and perspective effects). They are often greyscale, but multispectral products are also used (e.g. Edwards et al. (2011)). An ORI’s colour values may be in units related to the response of the camera that captured them, or in calibrated physical units such as radiance or reflectance. The effect of shadows, phase angle and atmospheres can all be accounted for in the production of an ORI to remove as much as possible any

peculiarities of the specific imaging conditions used to capture it (Edmundson, et al., 2012).

Using an ORI for image-based shading results in an approximation of what an image captured through the simulated camera would look like if it were to be calibrated similarly to the ORI. The advantage of this approach is its simplicity, whilst it also means that surface features in the simulated images have an appearance that is representative of a real image. For image-based shading, it is necessary that the ORI covers the same extent as the DTM, and shares the same coordinate system. When a triangle – defined by three adjacent DTM vertices – is drawn on the rendered image, it is filled with the colour of the ORI at the corresponding position of those vertices. Triangles are drawn in decreasing depth order, so that occlusions of the scene are properly represented.

The datasets used for producing the penetrator descent camera images in chapter 3 are first detailed in section 5.4, Table 5.1. For convenience, that table is repeated here (Table 7.1), as several of the datasets are displayed throughout this section also. Figure 7.8 shows an example DTM-ORI pair covering Victoria crater on Mars (dataset 5.2), and a rendered image of their resulting surface model, using image-based shading.

Table 7.1: Summary of the data used for simulation of surface images presented in this chapter (reproduction of Table 5.1).

Dataset ID	Constituents	Source
5.1	Gale Crater CTX[a] and HRSC[b] blended DTM[f] mosaic	Persaud et al. (2019)
	Gale Crater CTX[a] greyscale ORI[g] mosaic	
5.2	Victoria Crater HiRISE[c] DTM[f]	Paar et al. (2014)
	Victoria Crater HiRISE[c] SRR ORI[g]	Tao and Muller (2016)
5.3	Victoria Crater CTX[a] DTM[f] and ORI[g]	Persaud (2018)
5.4	Mars HRSC[b] and MOLA[d] blended global DTM[f] mosaic	Fergason et al. (2018)
	Mars MOC[e] global ORI[g] mosaic	Caplinger (2002)

[a]Context Camera
[b]High Resolution Stereo Camera
[c]High Resolution Imaging Science Experiment
[d]Mars Orbiter Laser Altimeter
[e]Mars Orbiter Camera
[f]Digital Terrain Model
[g]Orthorectified Image

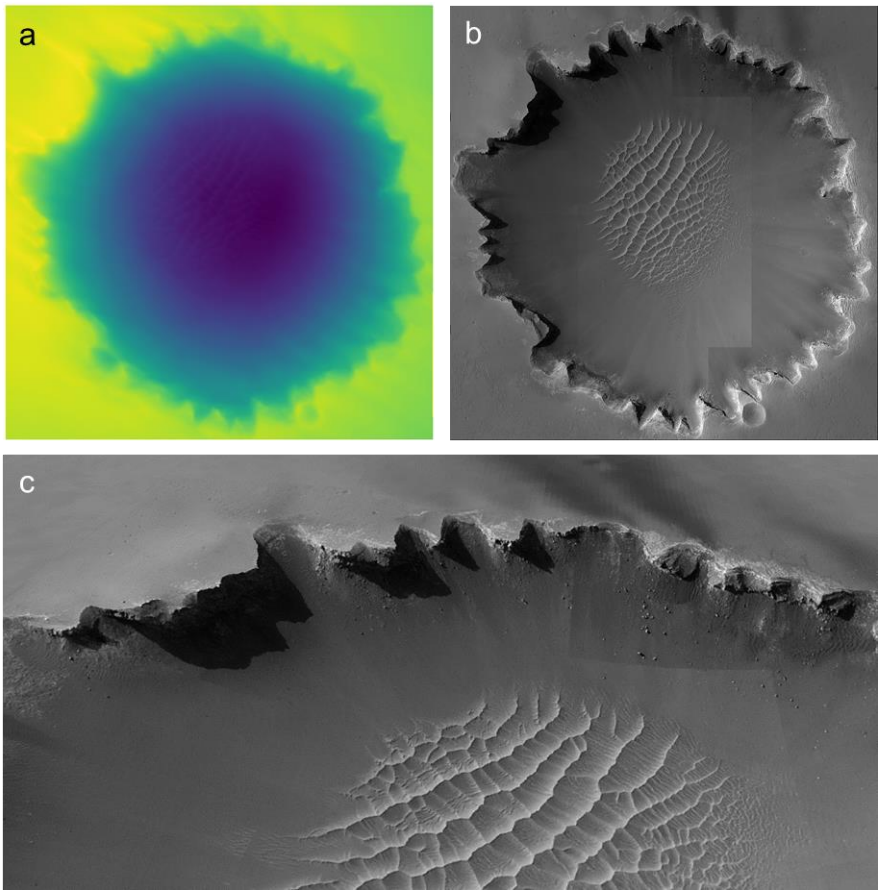

Figure 7.8: DTM (a) and ORI (b) of matching extent over Mars' Victoria crater, and an example view of a surface model constructed from the two (c). The DTM is colour coded, with dark (light) colours indicating low (high) terrain. Simulated with dataset 5.2.

Relying on ORIs to determine surface appearance in simulated images does have its limitations. However well-calibrated an ORI is, because it is constructed from real images captured from a specific perspective at a specific time, it will inevitably contain features peculiar to that imaging scenario. As a demonstration of this, Figure 7.9 (a1 and b1) shows two simulated images of the same region of Martian surface, with an identical imaging setup, using two different surface models (datasets 5.1 and 5.4 for a1 and b1 respectively), both containing ORIs constructed from satellite observations. Whilst the same surface features are seen (e.g. the elevated peak of Mount Sharp, small craters spanning the bottom third of each image, a curved channel cutting through Mount Sharp on the right hand side of the image), there are significant differences in their appearance (aside from the lower grid spacing of b1's surface mesh). Solar illumination is different in each image, resulting in markedly different shadows and highlights, and giving a different impression of the topography.

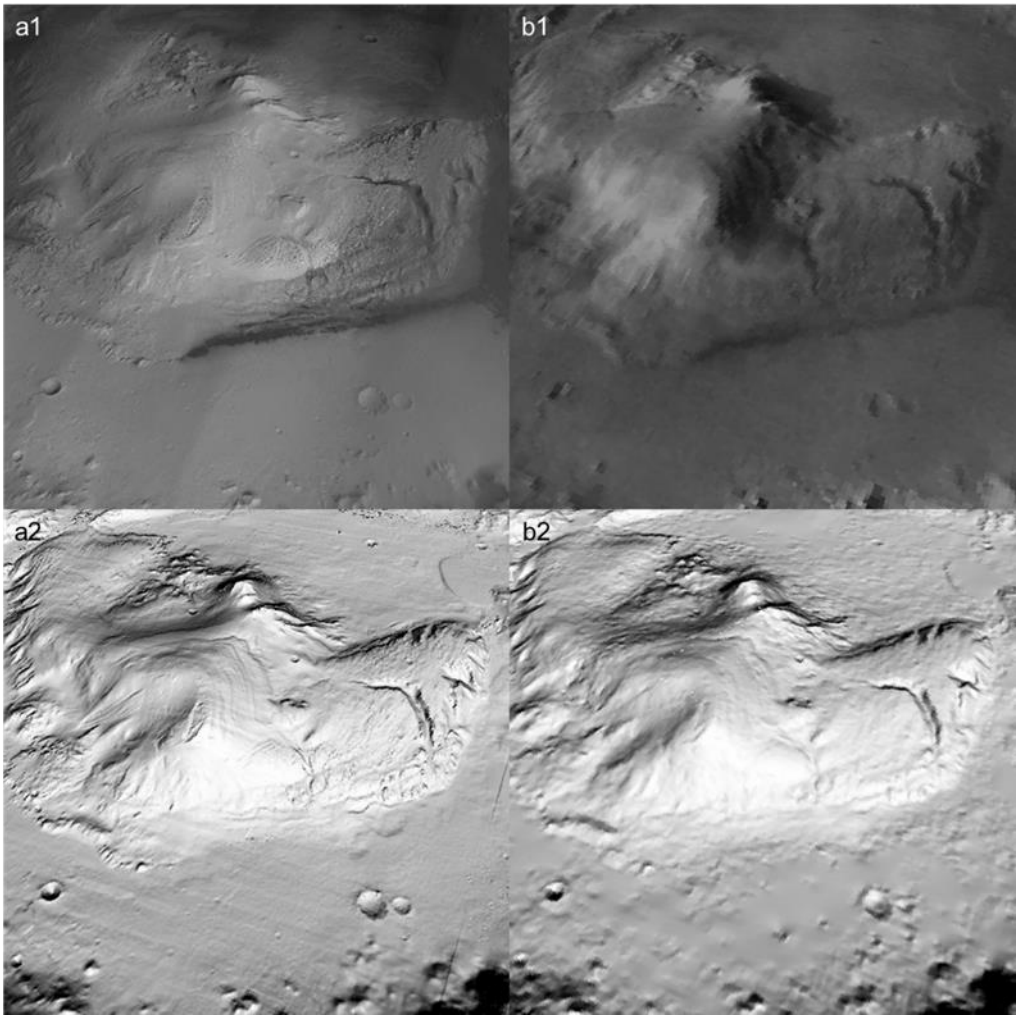

Figure 7.9: Four simulated images of an identical view of Gale Crater's Mount Sharp, Mars, using two datasets (datasets 5.1 and 5.4 for a and b respectively). Images a1 and b1 are simulated using ORIs to determine surface appearance. Surface appearance in images a2 and b2 is calculated based on Lambertian reflection.

To avoid the inclusion of ORIs' observation-specific peculiarities, surface appearance can be physically modelled based on the surface shape (described solely by the DTM), desired lighting conditions, viewing geometry and a physical model of the reflection/scattering properties of the surface. Panels a2 and b2 of Figure 7.9 show the same view as a1 and b1, and the same datasets 5.1 and 5.4 for a2 and b2 respectively, but surface appearance is calculated according to Lambertian reflection (Lester, et al., 1979). With this physically based (albeit very simplified) model of reflection, the two datasets result in almost identical images. When simulated images are required to provide a radiometrically accurate representation of what a camera will measure under specific imaging conditions, this approach is preferable to mapping surface appearance from ORIs. Whilst Lambertian reflection is used for this example, more rigorous reflection

models can readily be applied to increase the radiometric accuracy of the image simulation (e.g. Shepard & Helfenstein (2007)). A downside to this approach is that image contrast and detail are generally lower in a DTM than an ORI. Additionally, spatial variations in the reflection properties, such as albedo, are not replicated in this simulation approach.

As previously mentioned, the motivation for using rasterisation rather than ray tracing is the higher achievable rendering speeds. In particular, rasterisation doesn't require per-ray calculations, and in principle the time taken to render an image is independent of the number of image pixels. This is important for simulation of the penetrator descent camera's images, but also for simulating other scientific cameras, as high resolution sensors are common on modern imaging instruments. Figure 7.10 plots the times taken to render the image of Figure 7.8 (c) as a function of the number of vertices in the scene (where vertex number is varied by down sampling the DTM and ORI), and the number of pixels in the rendered image. The absolute rendering time of a specific image will depend on the surface model and imaging geometry, but Figure 7.10 demonstrates the minimal effect that image size has on render times.

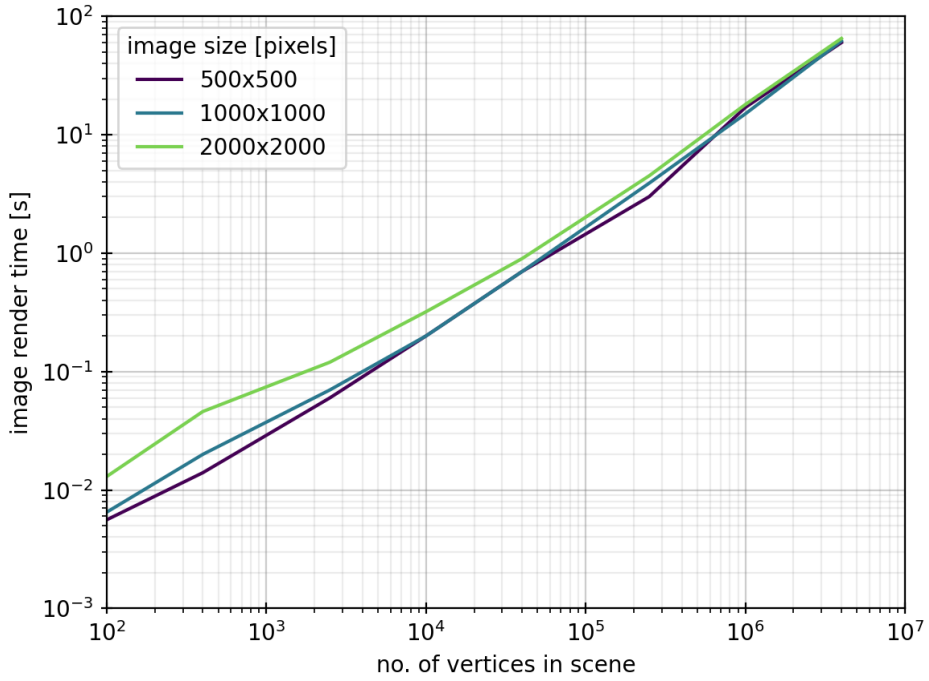

Figure 7.10: Time taken to render images as a function of the number of vertices in the scene, and the number of pixels in the rendered image.

Whilst this rendering method allows image resolution to be increased without penalising rendering time, there is still a need for caution when simulating high resolution images. Any surface mesh constructed from tessellated triangles is limited to representing surface details (i.e. topographic features, colour variation) at scales larger than the triangles themselves (i.e. it is not possible to replicate surface features at scales smaller than the inter-vertex spacing of the surface's DTM). As the simulated image's resolution is increased, each pixel samples a smaller region of the surface. When the spatial scale of a pixel matches, or is smaller than that of mesh triangles, the image is over-sampling the surface model. This can result in low image contrast at small scales, and the appearance of individual mesh triangles in the simulated image (Figure 7.11 (a)). Whilst a simulated image with these issues may be acceptable for illustrative purposes, it would not be suitable for testing analysis techniques such as the photogrammetry described in section 5.6.

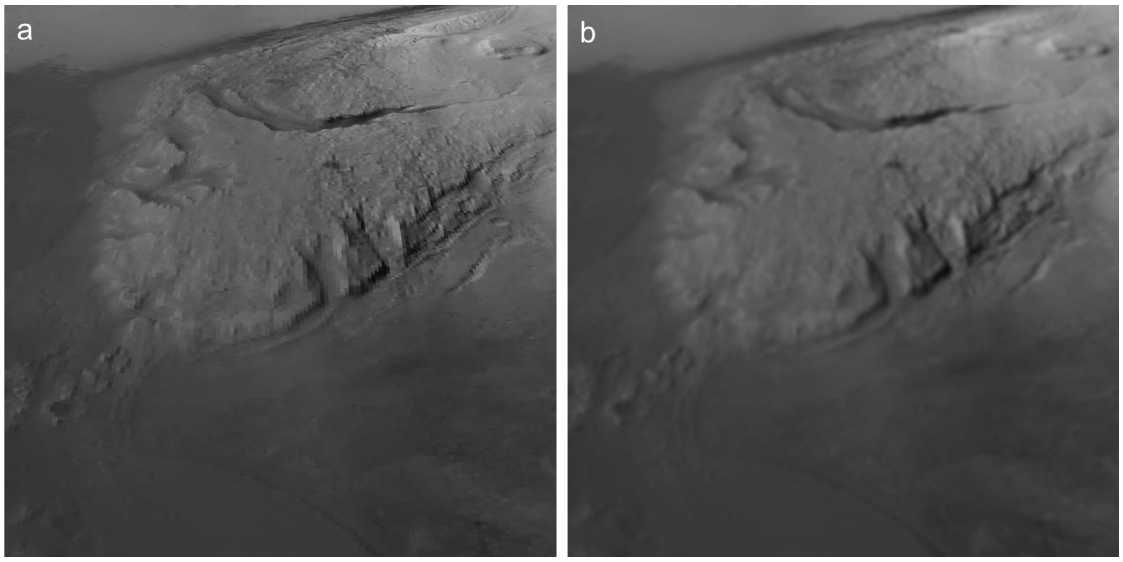

Figure 7.11: (a) A simulated image where the camera oversamples the surface mesh. Individual mesh triangles and their sharp boundaries can be seen, and the image is clearly synthetic. (b) The same image as (a), but with a lower resolution camera such that oversampling does not occur. The image cannot be easily identified as synthetic.

To avoid oversampling of the surface model, the spatial scale of the simulated camera's pixels should not be less than the DTM grid spacing, and should ideally by at least twice as large. When a real camera forms an image, each of its pixels collects signal from an area whose appearance is likely not uniform, and may contain significant variation. For a particular imager, the extent of the region contributing to a pixel's signal is characterised by its point spread function (PSF) and modular transfer function (MTF)

(Rossmann, 1969). To account for pixel sampling in the image simulation pipeline, rasterisation is performed on an image a factor of S larger than the desired final rendering. A Gaussian blur with a radius set to match that of the camera's PSF is then applied to the image, before it is downsampled by a factor S to produce the final image (e.g. Figure 7.11 (b)). The process is illustrated in Figure 7.12.

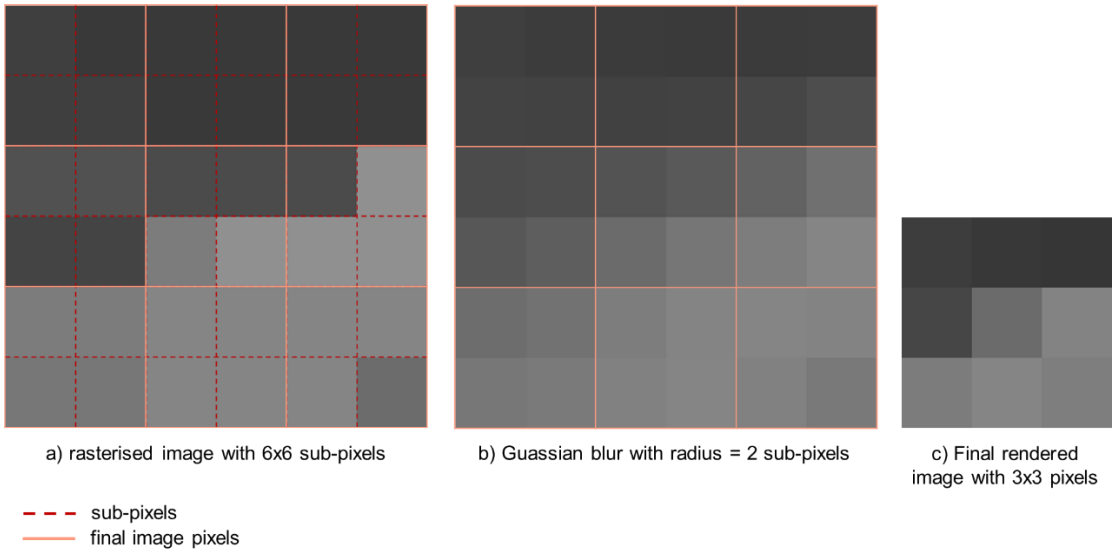

Figure 7.12: Illustration of accounting for sub-pixel detail in a 3x3 pixel rendered image. (a) Rasterisation is performed to a larger, 6x6 sub-pixel image. (b) A Gaussian blur is applied, with a radius matching the camera's PSF. (c) The image is downsampled, combining sub-pixels to produce the final 3x3 pixel image.

7.4 Simulating Images of Cometary Dust

Whilst the observation technique of EnVisS is very similar to that of the penetrator descent camera concept, the images it will observe will be very different. Rather than resolving surface features, EnVisS will image extended, diffuse structures which generally vary on large scales. The signals measured from these structures will be low, and will accumulate from volumes of material, rather than points on a surface. In order to simulate images of cometary material, an entirely different rendering approach is required than was used for the penetrator descent camera. Rather than attempt to simulate the appearance of fine scale structures, it is more important to broadly simulate the signal levels (and their large scale spatial variation) that the camera will measure. To

support the study of the EnVisS camera, image simulation had to meet the following objectives:

- Replicate the geometry, coverage and viewing perspectives of EnVisS.
- Simulate EnVisS images with radiometric accuracy.
- Permit the modelling of simple dust structures, representative of those seen in real comets.

Whilst rendering surface images is largely a matter of determining the location and geometry of each pixel's viewpoint, the simulation of cometary dust images is more a question of how much material lies along each pixel's LOS, and how much signal that material produces. Because signal is produced by extended volumes (often called 'participating media' in rendering literature (Max, 1995)) along a LOS, rather than a fixed point on a defined surface, rasterisation is not a suitable rendering technique. Instead, a simple form of ray tracing is utilised, whereby a ray is traced out of the camera for each of its pixels (using the camera's specific outward mapping function, and the signal directed back to the camera along that ray as a result of any material through which the ray passes is calculated.

7.4.1 Describing Radiative Transfer

Producing radiometrically accurate simulated images requires determining the radiance arriving at the camera along each of its pixels' LOSs. To achieve this, it is essential to incorporate the physics governing the propagation of light through participating media into the simulation code. Here, a description of the radiative transfer model employed in the rendering code will be described (a more general and thorough description of radiative transfer is given by Chandrasekhar (1960)). In an attempt to be consistent with planetary photometry, notation of some quantities will differ slightly from Chandrasekhar's, and will instead follow the conventions used in Lester et al. (1979) (see their Table 1).

Consider an extended volume of medium comprising identical particles with number density n and geometrical cross section σ_{geo}. Consider also a beam of light with spectral radiance L, defined by the energy E transferred per unit time dt through a unit surface da, into a unit solid angle $d\Omega$ (Figure 7.13). Strictly speaking, in the general case, radiance and many of the other variables that will be introduced below are wavelength dependent. However, it will be assumed here that they can be considered independent of wavelength over the spectral range imaged by the camera.

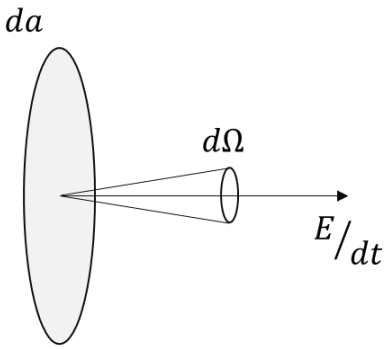

Figure 7.13: Diagram of the quantities used in the main text to describe radiative transfer.

In propagating a distance ds through the medium, the beam will encounter $n\,ds$ particles per unit area, with a total geometrical cross section of $n\,\sigma_{geo}\,ds$ per unit area. The interaction of the light with these particles can modify the propagating radiance by three processes: absorption, scattering and emission.

In the presence of absorption, a fraction of the beam's radiation is permanently lost from the radiation field as a result of its interaction with the particles. The resulting change in the light beam's radiance is

$$dL = -L\,n\,\sigma_{ab}\,ds \tag{7.5}$$

where σ_{ab}, the absorption cross section of a particle, is related to its geometrical cross section via its absorption efficiency, q_{abs} according to

$$\sigma_{abs} = q_{abs}\,\sigma_{geo} \tag{7.6}$$

Some additional radiant energy will be lost from the propagating beam due to scattering (the distinction being that scattered energy remains in the same form, and is simply diverted to new propagation directions, whilst absorbed energy is lost to other forms). The change in radiance due to scattering losses is

$$dL = -L\,n\,\sigma_{sc}ds \tag{7.7}$$

where the scattering cross section, σ_{sc} is the product of the scattering efficiency q_{sc} and the geometrical cross section, i.e.

$$\sigma_{sc} = q_{sc}\,\sigma_{geo} \tag{7.8}$$

The total extinction of radiant energy from the propagating beam due to both absorption and scattering is therefore

$$dL = -L\, n\, \sigma_{ext}\, ds \tag{7.9}$$

where the extinction cross section σ_{ext} is the sum of the absorption and scattering cross sections, and is also defined by the extinction efficiency q_{ext} and geometrical cross section:

$$\sigma_{ext} = \sigma_{sc} + \sigma_{abs} = q_{ext}\, \sigma_{geo} \tag{7.10}$$

The portioning of lost energy between absorption and scattering can be represented by the single-scattering albedo

$$\alpha_{sc} = \frac{\sigma_{sc}}{\sigma_{ext}} \tag{7.11}$$

i.e. the single scattering albedo is the fraction of the total extinct energy that is lost to scattering, whilst the fraction lost to absorption is $1 - \alpha_{sc}$. The scattered radiant energy is redistributed over a full 4π steradians. This redistribution is, in general, not isotropic, and is described by the scattering phase function $p(\phi)$, which defines what portion of the scattered radiation is directed into a unit solid angle along a phase angle ϕ. The phase function is normalised such that

$$\int \frac{p(\phi)}{4\pi} d\Omega = 1 \tag{7.12}$$

If scattering can direct radiant energy out of the propagating beam, then it can also direct energy into the propagating beam. The radiance directed into the propagating beam from a phase angle ϕ is

$$dL(\phi) = L(\phi)\, n\, \sigma_{sc} \frac{p(\phi)}{4\pi} ds \tag{7.13}$$

The total radiant energy gained by the beam due to scattering is found by integrating equation 7.13 over a full 4π steradians, thus accounting for light scattered from all directions into the beam:

$$dL = \iint_{\Omega} L(\phi)\, n\, \sigma_{sc} \frac{p(\phi)}{4\pi} d\Omega \,\mathrm{ds} = j_{sc}\, ds \tag{7.14}$$

where the scattered-emission coefficient,

$$j_{sc} = \iint_{\Omega} L(\phi)\, n\, \sigma_{sc} \frac{p(\phi)}{4\pi} d\Omega \tag{7.15}$$

is the total energy scattered by a unit volume of medium per unit time into a unit solid angle along the beam's direction of propagation.

It is worth noting that the absorption and scattering properties of the medium can also be, and often are, represented by absorption and scattering coefficients μ_{abs} and μ_{sc} respectively, and alternatively by the medium's mass density ρ and the mass absorption and mass scattering coefficients κ_{abs} and κ_{sc} respectively (and likewise for extinction):

$$\mu_{abs} = n\,\sigma_{abs} = \rho\,\kappa_{abs}$$

$$\mu_{sc} = n\,\sigma_{sc} = \rho\,\kappa_{sc}$$

$$\mu_{ext} = n\,\sigma_{ext} = \rho\,\kappa_{ext} \tag{7.16}$$

Returning to the medium's modification of the beam's radiance, energy can also be introduced to the beam due to emission of radiation particles. If the energy emitted by a unit volume of medium per unit time into a unit solid angle along the beam's direction of propagation is described by the intrinsic-emission coefficient j_e, the increase of the beam's radiance is

$$dL = j_e\,ds \tag{7.17}$$

The contributions from both scattering and intrinsic-emission combine to give a total emission coefficient, j:

$$j = j_s + j_e \tag{7.18}$$

The ratio of this emission coefficient to the extinction coefficient is called the source function, J:

$$J = \frac{j}{\mu_{ext}} \tag{7.19}$$

Combining the subtractive (absorption and outward scattering, equation 7.9) and additive (inward scattering and emission, equations 7.14 and 7.17) contributions gives the total change in the beam's radiance over its traversal of distance ds:

$$dL = -L\,\mu_{ext}\,ds + j\,ds \tag{7.20}$$

Rearranging and combining with the source function (equation 7.19) yields the radiative transfer equation:

$$-\frac{1}{\mu_{ext}}\frac{dL}{ds} = L - J \tag{7.21}$$

The radiative transfer equation is a first order linear differential equation, and its general solution is

$$L(s_2) = L(s_1)e^{-\tau(s_1,s_2)} + \int_{s_1}^{s_2} \mu_{ext}(s)\, J(s)\, e^{-\tau(s,s_2)}\, ds \tag{7.22}$$

where τ is the optical thickness of a path through the medium, and is given by

$$\tau(s_1, s_2) = \int_{s_1}^{s_2} \mu_{ext}(s)\ ds \tag{7.23}$$

The exponential $e^{-\tau(s_1,s_2)}$ gives the fraction of light that is not extincted over the path $s_1 \rightarrow s_2$. Equation 7.22 gives the radiance of light at s_2 after traversing the path $s_1 \rightarrow s_2$, having arrived at s_1 with a radiance $L(s_1)$.

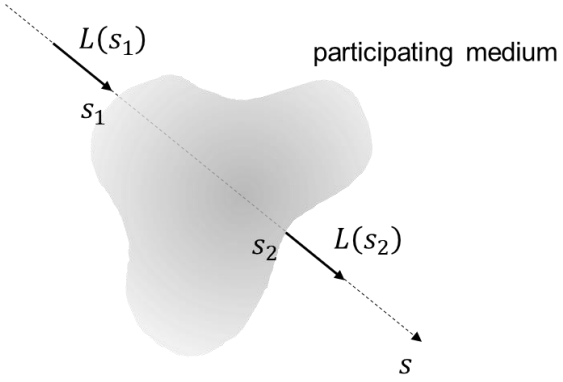

Figure 7.14: Radiances along a path, as in equation 7.22.

Equation 7.22 describes radiative transfer in the presence of absorption, scattering and emission, and can be used to simulate highly accurate physically based images of a wide range of scenes (Max, 1995). For the simulation of large-scale cometary dust environments, a number of approximations and assumptions can be made. These assumptions simplify equation 7.22 and the radiative transfer model. This makes the code implementation of the model simpler, but the main motivation for invoking these assumptions is to reduce the number of computations required to calculate observed radiances, and therefore keep image rendering times acceptably low.

- **Assumption 1**: The scene is illuminated by a finite number of discrete light sources, all of which can be approximated as infinitely distant point sources each illuminating the scene with a uniform flux density. As a result, performing the integration of equation 7.15 yields a scattered-emission coefficient of

  $$j_{sc} = \mu_{sc} \sum_i F_i \frac{p(\phi_i)}{4\pi} \tag{7.24}$$

  where F_i is the flux density of a light source's radiation, ϕ_i is the light source's phase angle with the observer, and the sum is performed over all light sources.

In other words, the total scattered contribution is the superposition of each individual light source's scattered contribution. For a cometary image, the Sun is the sole light source:

$$j_{sc} = F_{\odot}\,\mu_{sc}\frac{p(\phi_{\odot})}{4\pi} \tag{7.25}$$

- **Assumption 2**: The comet's dust is purely scattering, i.e. $\mu_{ext} = \mu_{sc}$ and $j = j_{sc}$.
- **Assumption 3**: The cometary dust is assumed to always be optically thin, and it is assumed that any attenuation of the light's intensity is negligible compared to its total intensity, i.e. optical thickness $\tau = 0$.

Applying assumptions 1-3 to equation 7.22 yields a simplified form:

$$L(s_2) = L(s_1) + F_{\odot}\frac{p(\phi_{\odot})}{4\pi}\int_{s_1}^{s_2}\mu_{sc}(s)\;ds \tag{7.26}$$

If the scattering properties of individual particles (i.e. σ_{sc}) are assumed to be spatially uniform, equation 7.26 can be further simplified:

$$L(s_2) = L(s_1) + \sigma_{sc}F_{\odot}\frac{p(\phi_{\odot})}{4\pi}\int_{s_1}^{s_2}n(s)\;ds \tag{7.27}$$

The integral of $n(s)$ between s_1 and s_2 gives the total number of particles within a column of unit cross-sectional area, extending from s_1 to s_2. This is known as the column density, given by:

$$n_{col} = \int_{s_1}^{s_2}n(s)\;ds \tag{7.28}$$

The radiance observed along a path can therefore be written in terms of the path's column density:

$$L(s_2) = L(s_1) + \sigma_{sc}F_{\odot}\frac{p(\phi_{\odot})}{4\pi}n_{col} \tag{7.29}$$

The assumptions applied here, and the resulting equation 7.29, are consistent with the techniques used in a number of published cometary studies, such as by Marschall et al. (2019; 2016) to model observations of comet 67P's inner coma.

Whilst equations 7.27 and 7.29 treat the dust as uniform in all respects other than density, it is desirable and necessary to be able to account for variation of other physical properties too. Comets can exhibit a number of structures each potentially with different optical properties. Moreover, a single structure will comprise particles with a distribution of sizes, and each particle size will interact differently with light. These differences can

be represented by considering multiple distinct participating media, each with its own optical properties. The overall observed radiance along a LOS is the sum of all the contributions from every participating medium:

$$L = \sum_j \sigma_{sc,j} F_\odot \frac{p_j(\phi_\odot)}{4\pi} n_{col,j} \tag{7.30}$$

For example, for the coma image simulations used in the detailed study of EnVisS' expected signals (see chapter 6, section 6.3), the dust was treated as comprising four populations of different particle radius: $1x10^{-6}$, $3x10^{-6}$, $1x10^{-5}$ and $3x10^{-5}$ m (taken from the Comet Interceptor dust coma model introduced in that same section). Dust scattering efficiencies, geometrical cross sections and phase functions for each of these particle sizes were taken from Marschall et al. (2020).

It is worth noting that an image of a scene is not just a function of the radiance observed, but also the particular response of the camera to that radiance. Particularly relevant for studies of EnVisS is the addition of noise to the pixels' measurements. The noise model used to simulate EnVisS signals is described in chapter 6, section 6.3.3, so will not be addressed here. But it is important to state that simulation of a radiometrically accurate image requires that noise be included according to a physically based model.

7.4.2 Simulating Cometary Structures

With a radiometric model established, it is necessary to implement ways of modelling real cometary structures, and calculating the column densities that would be observed by LOSs through these structures in order to simulate the radiances that cameras would observe. In order to produce simulated images with radiometrically accurate signals, it is important that modelled cometary environments are physically representative of real comets.

7.4.2.1 The Isotropic Outflow Coma

A simple but commonly used model for a comet's coma assumes a constant isotropic outflow of uniform dust particles at a rate Q_d with constant velocity v_d from a point-source nucleus. It is acknowledged that this model is not sufficient for describing the immediate surroundings of a nucleus (where the dust is undergoing acceleration, and the nucleus' extended surface cannot be treated as a point source). However, it is generally observed to be a good approximation of comae beyond 10 nuclear radii from the nucleus' surface (Gerig, et al., 2018). Given that EnVisS will be imaging the large scale coma, and will not be resolving the near-nucleus regions, the isotropic outflow model is well suited for modelling coma brightnesses.

Under the assumption of constant isotropic outflow, the dust particle number density follows a r^{-2} distribution, where r is the distance from nucleus centre:

$$n(r) = \frac{n(r_0)\, {r_0}^2}{r^2} = \frac{k}{r^2} \tag{7.31}$$

The constant k relates physically to the dust outflow by

$$k = \frac{Q_d}{4\pi\, v_d} \tag{7.32}$$

The observed brightness when viewing such a coma from Earth or large distance is often stated in the literature as being proportional to b^{-1}, where b is the impact parameter of the observation's LOS (Brandt, 1968). But a more general expression for the observed brightness from an arbitrary perspective (such as a spacecraft within the coma) is generally not addressed in the literature. Consider an observer at a distance R from the nucleus, observing a LOS with an angle θ from the observer-nucleus line, as illustrated in Figure 7.15.

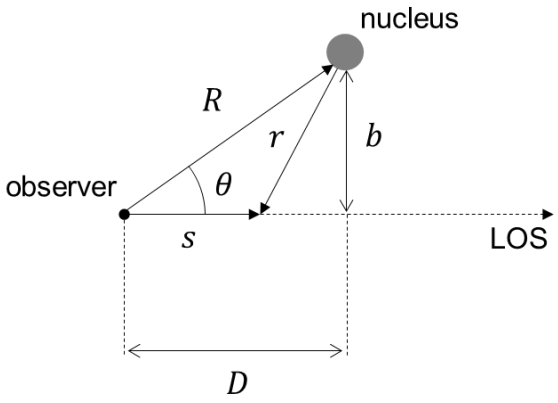

Figure 7.15: The geometry of a coma observation.

A point along the LOS with a distance (or depth) s from the observer has a distance r from the nucleus given by

$$r^2 = b^2 + (D - s)^2 \tag{7.33}$$

where the LOS's impact parameter b and distance to closest approach D are given by

$$b = R \sin\theta$$

$$D = R \cos\theta \tag{7.34}$$

Combining equations 7.28, 7.31 and 7.33, the observed column density is given by

$$n_{col} = \int_{s_1}^{s_2} \frac{k}{b^2 + (D-s)^2}\, ds = \left[-\frac{k \tan^{-1}\left(\frac{D-s}{b}\right)}{b} \right]_{s_1}^{s_2} \tag{7.35}$$

Equations 7.34 and 7.35 can together be used to calculate the column density of any LOS through a spherically symmetric coma with r^{-2} density distribution. For an observation obtained from a spacecraft, the observed column density is found by integrating from a depth of 0 to ∞:

$$n_{col} = \frac{k}{b}\left(\frac{\pi}{2} + \tan^{-1}\frac{D}{b}\right) \tag{7.36}$$

An observation from a very distant point, such as the Earth, can be treated as viewing from an infinite distance (i.e. $D = \infty$), reducing equation 7.36 to the well known b^{-1} relationship:

$$n_{col} = \frac{k\,\pi}{b} \tag{7.37}$$

Equations 7.29 and 7.36 facilitate the calculation of the radiances observed by EnVisS' pixels when viewing a spherical coma with isotropic outflow of dust particles. A major benefit of the r^{-2} distribution is that the column density integral has an algebraic solution, allowing its fast calculation. Further to this, in Python, equation 7.35 can be evaluated using NumPy (introduced in section 7.3), allowing the parallel calculation of all pixels' observed column densities. This helps to keep render times short, even when simulating high resolution images, as shown in Figure 7.16.

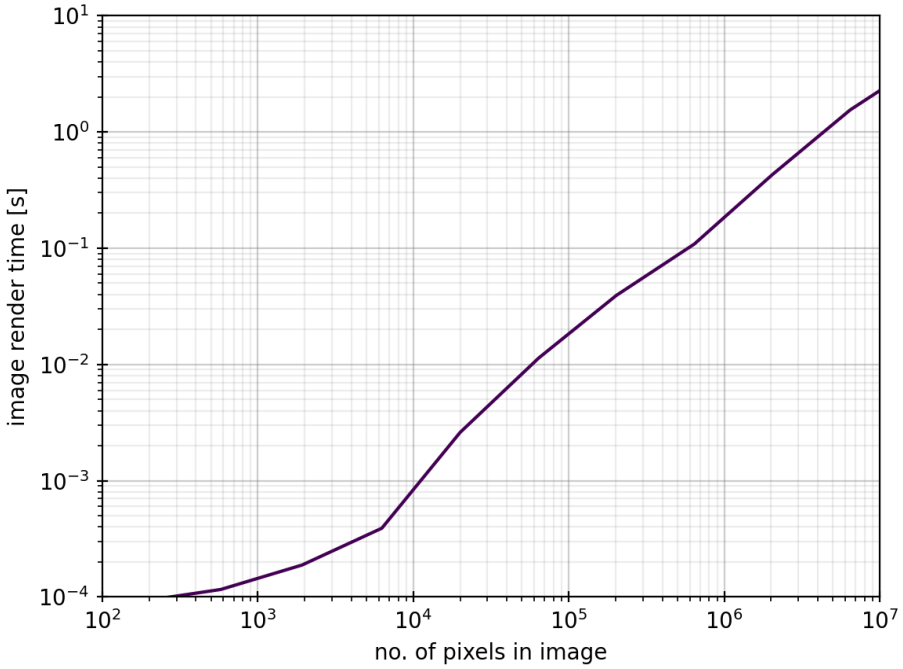

Figure 7.16: Time taken to render images of a spherical coma as a function of the number of pixels in the rendered image.

Whilst the isotropic outflow coma is simple and lacks localised features, it is still a useful model for simulating EnVisS images. It can be used not only to calculate order-of-magnitude values for the radiances that EnVisS will observe, but also reveal the distribution of high- and low-signal regions across the sky, both of which are valuable for understanding and planning the best approach to capturing images.

As mentioned in section 7.4.1, coma images were simulated by modelling the coma as a combination of four dust density distributions, each associated with a different dust particle radius. To achieve radiometric accuracy, the absolute densities of each of these populations were scaled to match the Comet Interceptor Engineering Dust Coma Model (EDCM), as described in chapter 4, section 6.3.2.2. The dust particle radii and their associated values of k (as in equation 7.31) are listed in Table 7.2. Figure 7.17 shows three examples of radiance images simulated from the isotropic outflow coma model, using the dust parameters of Table 7.2.

Table 7.2: Dust particle radii and associated k values (equation 7.32) used for coma simulation.

Dust Radius [m]	k [m^{-1}]
$1x10^{-6}$	$8.8x10^{-12}$
$3x10^{-6}$	$3.5x10^{-11}$
$1x10^{-5}$	$1.5x10^{-10}$
$3x10^{-5}$	$6.6x10^{-8}$

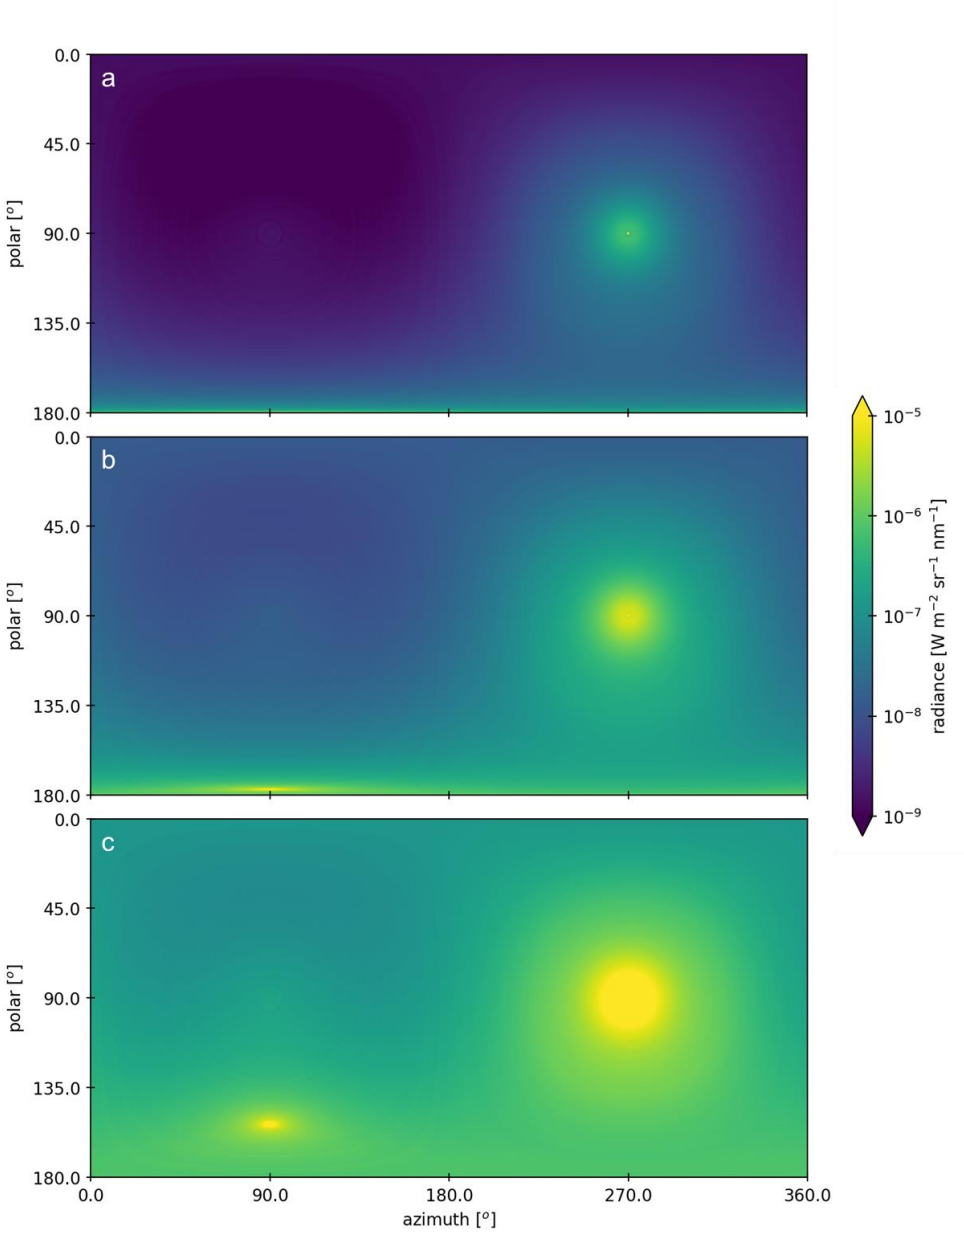

Figure 7.17: Three examples of simulated radiance images from a single flyby at spacecraft-nucleus distances of 10^5, 10^4 and 10^3 km (a, b and c respectively), demonstrating the change in scene shape and brightness as the flyby progresses.

7.4.2.2 Simple Dust Jets

The spherically symmetric coma is a useful description of a comet's broad dust environment, and valuable for modelling the levels of signal that are expected within an image, but it lacks smaller scale and localised variations in dust density which are known to be common in comets. Simulating the nature of more diverse structures is valuable for demonstrating the expected form of images, and potentially also for developing image

processing pipelines that are robust to, for example, deviations in EnVisS' spin axis orientation.

One of the most common features that is observed in cometary comae is dust jets (e.g. Sekanina & Larson (1986)). In contrast to the coma. dust jets are outbursts of material with significant directionality, but they can still be modelled by equation 7.35. Huebner et al. (1988) model dust jets as resulting from point sources on the nucleus surface, which eject dust uniformly into a cone. In this model, dust density is independent of angle and follows equation 7.31 within the cone, whilst being zero everywhere outside the cone. The column density observed by a LOS is therefore found by applying equation 7.35 with limits s_1 and s_2 dictated by the LOS's intersection with the dust jet's cone (Figure 7.18).

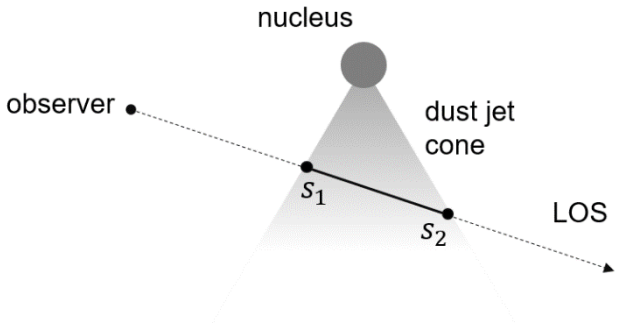

Figure 7.18: Geometry of the path integration used to determine dust jet column densities.

A cone is a quadric surface, whose points of intersection with a straight line in 3D space are algebraically calculable. This permits the fast calculation of s_1 and s_2, helping to keep render times short. Solving the intersection of straight lines with quadric surfaces is trivial and thoroughly described in the literature (e.g. Sarraga (1983)), so it will not be addressed here. Note that the bounding of a r^{-2} dust distribution can equally be described by other surfaces (e.g. sphere, cylinder), allowing the visualisation of a range of dust structures. In Python, the calculation of pixels' intersection points is readily vectorised, allowing significant speed improvement through the use of NumPy.

Figure 7.19 shows an example of a simulated radiance image of a comet with an isotropic outflow coma and three dust jets (labelled by the red arrows) originating from the nucleus. The rightmost dust jet is directed at the Sun, and appears, due to EnVisS' imaging geometry, in a different part of the image to the nucleus. This demonstrates the significant effect the imaging geometry of a camera can have on the appearance of the scene, and that features within images could be missed or misinterpreted if the geometry is not properly accounted for in analysis.

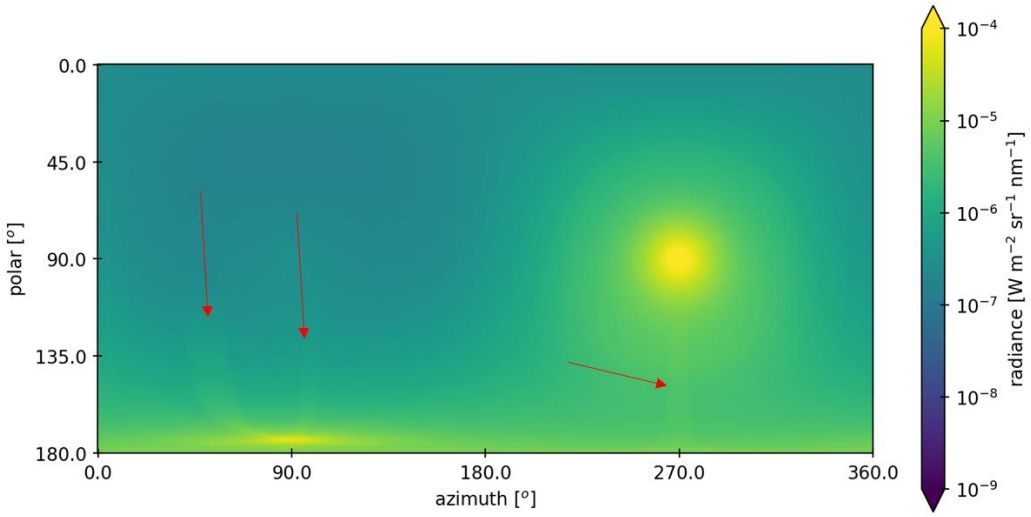

Figure 7.19: An example of a simulated radiance image of an isotropic outflow coma with three conic dust jets (indicated by red arrows). The right-most jet is directed at the Sun.

7.4.2.3 Simulating Arbitrary Dust Structures

The r^{-2} density distribution arises under, and can only approximate, a limited set of physical conditions. The large and varied range of dust structures that have been observed in cometary environments can, more generally, be described by any dust particle number density function $n(\vec{r})$, where $\vec{r}$ is a position vector. Commonly, $n(\vec{r})$ will be such that the integral of equation 7.28 cannot be solved symbolically, and the observed column densities instead must be numerically approximated.

Consider an arbitrary quadric surface that defines the limits of a dust structure (just as with the above discussion of conic dust jets) with an arbitrary density distribution $n(\vec{r})$, and an observer at position $\vec{o}$ viewing this dust structure along a LOS with direction vector $\vec{l}$.

The limits of integration s_1 and s_N are first determined by an intersection test with a bounding surface, as described above (section 7.4.2.2) for the conical dust jets. The column density is found by sampling and summing the density at N locations along the path $s_1 \rightarrow s_N$ (i.e. a Riemann sum):

$$n_{col} = \Delta s \sum_{i=1}^{N} n(\vec{r}_i) \tag{7.38}$$

where the inter-sample spacing Δs is given by

$$\Delta s = \frac{s_N - s_1}{N - 1} \tag{7.39}$$

The sampling position $\vec{r}_i$ is found from the observer position, LOS direction vector and sample depth s_i by

$$\vec{r}_i = \vec{o} + s_i \vec{l} \tag{7.40}$$

whilst sample depth is given by

$$s_i = s_1 + (i - 1)\Delta s \tag{7.41}$$

The method of sampling a dust structure's density is illustrated in Figure 7.20.

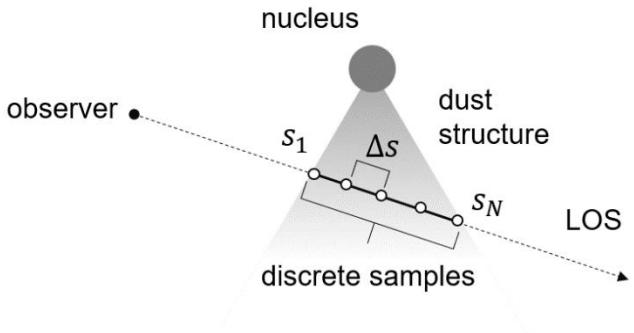

Figure 7.20: The method of numerically determining column density along a LOS.

Because this approach requires the dust density of each dust distribution in the scene be sampled N times per pixel, the number of computations required to calculate the column density viewed by every pixel of an image is larger than when an algebraically integral dust distribution, such as the r^{-2} of sections 7.4.2.1 and 7.4.2.2, is used. With this in mind, NumPy can be utilised to significantly increase the speed of the calculations, and keep rendering times short.

- For a single dust structure, s_1 and s_N are calculated by ray-surface intersection for all pixels in parallel, with the results stored in two $x \times y$ arrays.
- Δs is calculated for all pixels in parallel using equation 7.39.
- All sampling depths s_i for all pixels are generated, and stored in a $x \times y \times N$ array. This is the functionality of NumPy's built-in linspace() function, resulting in fast calculation of all sampling depths.
- Densities $n(\vec{r}_i)$ are calculated from the sampling depths for all pixels in parallel, using equations 7.40 and 7.41.

- Column densities for each pixel are calculated with equation 7.38. The summation of each pixel's N sampled densities is performed with NumPy's built in sum() function, minimising the calculation time.

Selecting the value of N is a trade-off between greater accuracy (larger N) and shorter render times (smaller N). Figure 7.21 plots the time taken to render an image comprising a single dust structure as a function of samples per pixel N and number of pixels in the image. Render times are approximately a factor N longer than in the case of the isotropic outflow coma (Figure 7.16).

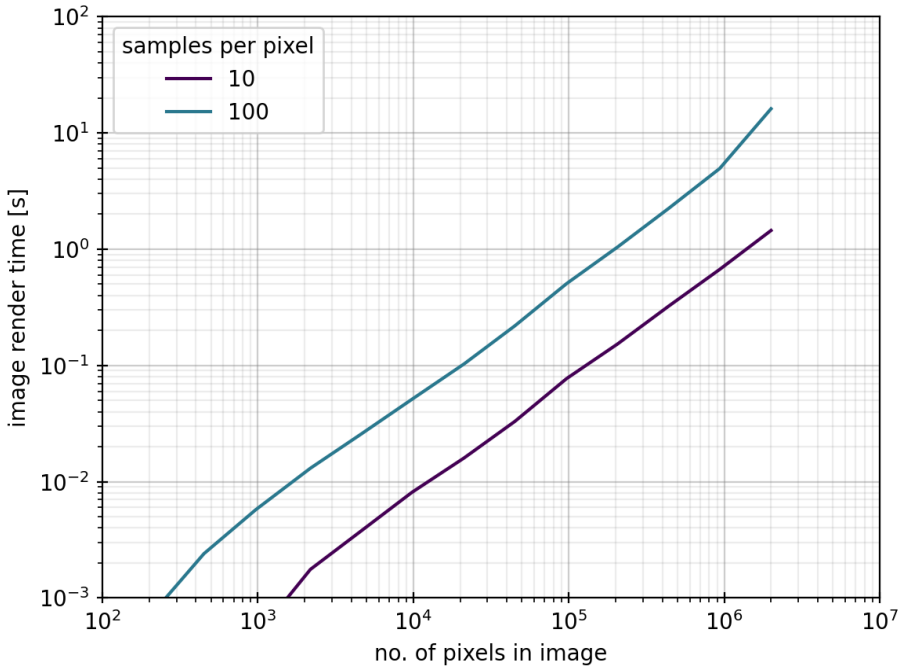

Figure 7.21: The time taken to render an image of a single dust structure using the numeric integration method, as a function of samples per pixel N and number of pixels in the image.

It can be seen from Figure 7.21 that high resolution images (10^6-10^7 pixels) approach render times as large as 100 seconds if ~100 samples per pixel are used. Without the use of NumPy's acceleration, render times would be even larger (by a factor ~100). This highlights the challenge posed by rendering with Python, and the necessity of techniques such as using NumPy to keep render times acceptably short.

At the cost of its increased render time, numerically integrating column densities offers the flexibility to simulate a wider variety of dust structures and environments. The physical accuracy of images simulated through this method is dependent on the

accuracy of the dust density distribution $n(\vec{r})$ used. Dust density distributions can be derived from physical processes (the r^{-2} distribution being a simple example of this), comet observations, or tailored to produce images with desired appearance (useful for images intended for illustration or outreach, rather than quantitative analysis).

As a simple example, consider a spinning nucleus which constantly ejects dust isotropically, resulting in a r^{-2} dust distribution. In addition to this, a single small, isolated active region on the nucleus ejects dust at a rate proportional to its solar irradiance. For simplicity, let the nucleus spin axis orientation be at a constant right angle to the Sun-nucleus line, and the active region be positioned on the nucleus' equator such that it passes through the sub-solar point. Assume also that any material ejected from the active region is evenly distributed over the dayside hemisphere, and none of its material is ejected into the nightside hemisphere (Figure 7.22).

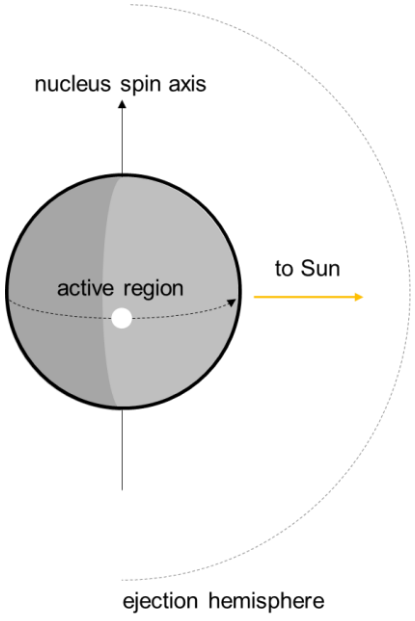

Figure 7.22: A spinning nucleus with isolated active region, whose ejection rate is proportional to its solar insolation, and isotropically distributed over the ejection (dayside) hemisphere.

Assuming the active region to be small compared to the nucleus, the dust particle ejection rate from the active region as a function of time t is

$$q_a(t) = \max\left[Q_a \cos\left(2\pi\frac{t}{T}\right), 0\right] \tag{7.42}$$

where T is the nucleus rotation period and Q_a is the maximum ejection rate of the active region (where it has been arbitrarily chosen that $t = 0$ corresponds to the active region

crossing the sub-solar point). A negative cosine term corresponds to when the active region is on the nightside, during which irradiance is zero. The time since ejection, Δt_e of dust at a distance r from the nucleus is

$$\Delta t_e(r) = \frac{r}{v_a} \tag{7.43}$$

where v_a is the outflow velocity of the active region's ejected dust. Therefore, dust at a radius r at a time t was ejected from the nucleus at a rate

$$q_{ae}(r,t) = q_a(t - \Delta t_e) = \max\left[Q_a \cos\left(2\pi \frac{t - \frac{r}{v_a}}{T}\right), 0\right] \tag{7.44}$$

From equations 7.31 and 7.32, the density of the active region's dust in the ejection hemisphere at a given time and radius is therefore given by

$$n_a(r,t) = \frac{q_{ae}(r,t)}{2\pi\, v_a\, r^2} \tag{7.45}$$

whilst the density of the active region's dust in the nightside hemisphere is always zero. Total dust density is the sum of the densities due to isotropic outflow (equation 7.31) and the active region (equation 7.45).

Figure 7.23 (a) shows an example of a simulated column density image of a coma comprising both an isotropic outflow dust source and an active region dust source as described above. Dust arcs can be seen in the ejection hemisphere due to the periodic variation of dust emission. For comparison, Figure 7.23 (b) shows a real image of comet Hale-Bopp (reproduced from Woodney et al. (2002)) in which real dust arcs of similar appearance can be seen.

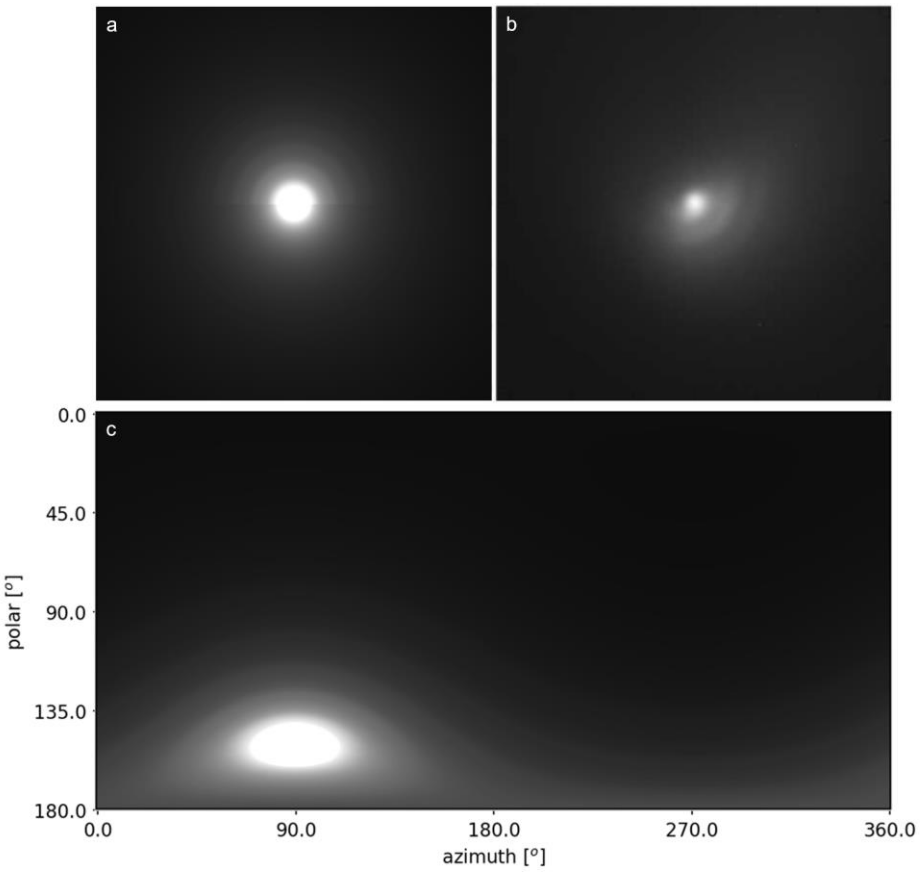

Figure 7.23: (a) Simulated column density image of a dust coma with two dust sources: constant isotropic outflow and a single active region, as described in the main text. (b) A real image of comet Hale-Bopp showing comparable arc features, reproduced from Woodney et al. (2002). (c) The same scene as panel (a), but imaged by EnVisS from within the coma.

The comparison of Figure 7.23 demonstrates that even a very simplified model of an active region on the nucleus is capable of simulating images with features closely resembling those seen in real observations. A sharp boundary is visible in the simulated image between the dayside and nightside hemispheres due to the oversimplification that dust density due to the active region is zero everywhere in the latter. Figure 7.23 (c) shows a simulated image of the same scene as (a), but captured by EnVisS from within the coma. It illustrates the form that dust arcs would be expected to take if present in EnVisS observations. This simulation method could be used to model the way in which a wide range of cometary features would manifest in EnVisS images. Samarasinha et al. (1999) demonstrate how the parameters and rigour of dust density models can be varied to produce a range of images with different dust features.

7.5 Discussion

Image simulation is a valuable tool to assist the study of imaging systems and analysis. The research of this thesis was largely facilitated by the ability to representatively simulate the data of the novel camera concepts investigated. Pre-existing image rendering software and tools are not always available or compatible with the requirements of a specific image simulation task, especially if the requirements, as was the case for this thesis, are quite niche. This necessitated the development of simple bespoke techniques to simulate images of both planetary surfaces and cometary atmospheres. This chapter has described the approach taken to simulating such images.

For easy implementation, flexibility and speed, the image simulation code was designed to be as simple as possible. For compatibility and continuity with all other code used for this thesis' research, the simulation software was written in Python. Whilst a large number of well-developed libraries exist for Python (e.g. NumPy (Harris, et al., 2020), SciPy (Virtanen, et al., 2020), Astropy (Robitaille, et al., 2013)), providing scientists with a ride range of valuable functionalities, no package exists that provides the necessary image simulation capabilities for the work of this thesis. Furthermore, the slow execution speed of Python generally precludes its use in image simulation. However, this chapter demonstrates that relatively simple rendering pipelines can be implemented in Python with the capability to simulate planetary images that can be used to assist with the study and development of scientific cameras.

Simulated high-resolution surface images can be used to demonstrate the achievable perspectives and coverage of a specific camera and spacecraft trajectory, or perform rigorous studies of photogrammetric analysis, such as stereophotogrammetry, in order to quantify the capabilities of an imaging system. This can help to inform the development of a camera and its associated data processing techniques, as well as assist with observation planning. Meanwhile, implementing the radiometric model of equation 7.22 in Python allows the physical simulation of cometary dust images. Indeed, simulation of images of a range of planetary atmospheres would be possible with the appropriate application of approximations. The approximations described in this chapter permit images to be simulated from a physically accurate model (which is consistent with independent studies, such as Marschall et al. (2019; 2016)), whilst keeping rendering times acceptably low. Even with the relatively simple physical models used in this chapter, a wide range of cometary images can be simulated.

For both surface and cometary images, rendering times range from milliseconds for low detail images (suitable for rapidly assessing a certain perspective, for example) to tens

of seconds for the highest detail, highest resolution images. These rendering times are achievable only through the use of Python's NumPy library, without which images would take up to hours to produce. It is worth noting that the image simulation approach described in this chapter could be readily implemented in a range of languages.

Simulating the data of planetary imagers is a valuable tool in their development, especially at their concept and early stages. Image simulation can produce a camera's data before any hardware exists (as is the case for the cameras studied in this thesis). Compared to constructing prototype or development hardware, simulation of cameras' images is significantly cheaper, faster and lower risk. Modern computing makes image simulation accessible to anyone with a computer, and the achievable realism will only increase as computing power continues to grow.

8 Conclusion

Spaceborne imaging has been, is, and will continue to be a crucial part of planetary science. This thesis studied a novel spin-scanning approach to imaging, with the aim of assessing its capabilities, demonstrating its feasibility and laying the groundwork for the technique's use on future missions. The fundamental concept entails performing scan imaging from a spinning platform, drawing on instruments such as Giotto's Halley Multicolour Camera (Keller, et al., 1987), and using multiple optical filters to facilitate multichannel imaging, similarly to instruments such as JunoCam (Hansen, et al., 2017). Also key is the use of wide-angle optics to achieve maximum image coverage.

The work focussed on two different applications of this spin-scanning concept, studying two distinct instruments: a descent camera for a spinning planetary penetrator, and an all-sky imager for a cometary flyby probe. The former of these is a concept, whilst the latter is an instrument under development for ESA's Comet Interceptor mission, and is intended to launch to space in 2029.

In the course of studying these cameras, a simple approach to simulating planetary imaging data was developed. It was purposefully designed to be accessible, requiring no specialised or commercial software, and easy to implement and modify.

8.1 A Scanning Descent Camera for a Planetary Penetrator

The work focusses on a camera mounted to a rotating planetary penetrator with a vertical spin axis and free fall descent trajectory. Optics with a 90° vertical FOV image a narrow strip of surface extending from the penetrator's nadir to above the horizon. With the probe's rotation, the camera is able to image the entire visible surface. Descent imaging is a valuable method of studying planetary surfaces. It can be used to accurately locate a landing site, and characterise the geological processes present in the local area. In some cases, descent imaging may provide the only means by which to obtain surface images.

The penetrator descent camera concept observes the surface with a unique imaging geometry, distinct from any that has previously been employed (it most closely resembles the combined performance of the Huygens probe's three DISR cameras (Tomasko, et al., 1999)). Imaging characteristics, such as spatial resolution, vary

significantly over the instrument's large FOV, with resolution highest at the penetrator's nadir, and decreasing monotonically as surface distance from nadir increases. Pixel footprints are elongated in the radial direction (away from nadir), and this radial component is the limiting factor in the camera's spatial resolution. A given point on the surface is imaged with maximum radial resolution when the camera is at an altitude approximately equal to the point's radius from nadir. This dictates an optimal imaging sequence, but the required increasing imaging frequency as altitude decreases can likely not be sustained all the way to impact. Discarding regions of surface that are imaged with greater resolution in other images of the sequence could be used to reduce the instrument's data volume. A radiometric study of the penetrator camera's observations, similar to that performed for EnVisS, would shed light on what exposure times, resolutions and SNRs can feasibly be achieved.

The changing altitude of the camera provides the opportunity to observe the surface from a range of perspectives. In this thesis, and also presented in Brydon et al. (2021), it is shown that this provides the opportunity to measure surface topography. The vertical baseline between two images of different altitude permits the measurement of surface elevation by two-image stereophotogrammetry. The achievable accuracies of elevation measurements are dependent on viewing geometry, and are lowest at the nadir and horizon, and highest in the mid region of the camera's FOV. Measurement of the elevation of a given point on the surface is found to have a maximum accuracy when the lower stereo image has an altitude ~0.2-0.5 times the point's radius from nadir, and altitude difference between the stereo images is ~0.2-0.4 times the lower image's altitude.

A key finding of the descent stereophotogrammetry study is that measurement accuracy is largely linked to the performance of the stereo matching algorithm. The camera's wide-angle imaging geometry and significant change of perspective results in scale and shape changes and occlusions within images that are not seen in more traditional stereo data sets (e.g. orbital). The semi-global block matching algorithm used in this study is not ideally suited to significant shape or scale changes. A quantitative assessment of stereo matching algorithms' performance with the penetrator images would facilitate higher accuracy topography measurements, and permit the development of a more robust stereophotogrammetry pipeline for applying to wide-angle descent images.

Use of penetrator spin motion to form images can potentially lead to variable imaging geometry. Using simulated data obtained from a drone-mounted camera, this thesis demonstrates a simple correction for non-uniform spin rate. On a real penetrator,

complementary spin rate data would be required, possibly obtained from ancillary sensors such as gyroscopes or sun sensors, or from the images themselves. Rectification and co-registration of different imaging channels would be required if multichannel products such as colour images were to be derived from the camera's data. This thesis shows that this is relatively straight forward when the penetrator's spin state is well-constrained and slowly-varying, but it may be a significant challenge in the case of more complex motion of the probe.

Relatedly, the early-stage study of the descent camera concept carried out in this thesis demonstrates the feasibility and capabilities of the instrument in principle. A small-scale hardware mock-up of the instrument, capable of replicating the camera's spin-scanning mode of imaging with optics and a sensor representative of the concept's FOV, would be a valuable avenue by which to further this research and test the camera's capabilities in practice. Such research could assist with elevation of the instrument concept to a more mature design, potentially suitable for inclusion in future mission proposals.

And with that in mind, it is worth noting that many of the techniques used in the study of the penetrator descent camera, and many of the study's findings, have scope beyond this specific instrument concept. The theory of section 5.3 can be applied, with modification where necessary, to any descent imager, even one of significantly different imaging geometry. Additionally, cameras with similar imaging geometry, and therefore similar spatial resolution profiles and optimum imaging altitudes, could be achieved with significantly different setups to that of the penetrator descent camera. A downward looking camera with hemispherical fisheye optics on a non-spinning probe, for example. Finally, the camera's multi-altitude image sequence, which permits the surface to be imaged with a range of resolutions and perspectives, and facilitates stereophotogrammetry, could be achieved with a variety of craft (e.g. such as the Mars 2020 Ingenuity helicopter (Balaram, et al., 2021) and the DAVINCI+ Venus probe (Garvin, et al., 2020)).

8.2 EnVisS: The Entire Visible Sky Camera for the Comet Interceptor Mission

The EnVisS instrument is under development for the European Space Agency's Comet Interceptor mission, which is scheduled to launch in 2029. The mission aims to perform a multi-probe flyby of a long period comet, ideally targeting one that is also dynamically new. Performing push-frame imaging with a 180° FOV from a spinning spacecraft, the

camera is designed to capture images that cover almost the whole sky, allowing the study of the comet's large-scale dust structures.

This thesis describes work that has made a valuable contribution to the development of the EnVisS instrument in its early stages, and will help to inform its continued advancement towards its flight on Comet Interceptor. In this thesis, the camera's imaging geometry, angular resolution, filter FOVs and coverage of the sky are calculated. At full resolution, the camera will image with an angular scale of ~0.09° per pixel. Rotation of the spacecraft, expected to occur with a fixed period in the range 4-15 s, imposes a maximum possible blur-free exposure time of ~1-4 ms, depending on spin period. With a provisional focal plane design incorporating optical filters each with a FOV spanning ~5° and 180° in the along-track and across-track directions respectively, achieving complete coverage of the sky will require a minimum framelet imaging frequency of ~18 Hz for a 4 s spin period, or 4.8 Hz for a 15 s spin period. The camera's filters do not have uniform along-track FOVs, and oversampling of the sky occurs toward the two poles of the spacecraft's spin axis.

The limitation on exposure time imposed by spacecraft spin restricts the instrument's signal, and leads to low SNRs. Comparison of EnVisS' imaging characteristics with those of previous cometary cameras, a review of previous spaceborne cometary observations, and physical models of the expected dust environment all indicate that EnVisS is likely to observe spectral radiances in the range 10^{-9} to 10^{-4} $Wm^{-2}sr^{-1}nm^{-1}$ (at the wavelengths to which the instrument is sensitive), and that detecting brightnesses at least as low as 10^{-8} $Wm^{-2}sr^{-1}nm^{-1}$ will be essential for observing large scale dust structures.

To estimate the signals EnVisS will measure, a radiometrically accurate model of the instrument was developed, with which image SNR can be calculated from an input radiance. Imaging at full resolution, it will not be possible to achieve acceptable SNRs for the above expected scene brightnesses. This thesis therefore investigated a number of techniques for increasing the SNR of EnVisS images. Pixel binning is the most effective method available for increasing SNR, thanks to the longer exposure times permitted by the reduction in angular resolution. Binning up to a factor of 100x50 pixels was considered, making SNR increases greater than 1000 possible, but at the cost of reduced angular resolution. SNRs remain low – on the order of 5 – for the lowest expected brightnesses even with this maximum binning. As the flyby progresses, the average and peak brightnesses observed by EnVisS are likely to change, and the level of pixel binning should vary accordingly throughout the flyby. Broadly speaking, the instrument should aim to obtain low-resolution but high-signal images at the beginning

of the flyby, and progressively higher-resolution images as the spacecraft nears closest approach.

Co-addition of consecutive all-sky images also increases image SNR. Given that the factor by which SNR is increased is proportional to the square root of the number of co-added images, the possible SNR gains are relatively low compared to binning. Co-addition requires good alignment between images, which itself relies on stability of the spacecraft spin state, otherwise angular resolution can be degraded. For this reason, co-addition will likely be most effective when used on binned images, whose angular resolution is lower from the outset. Co-addition itself will likely only increase SNR by a factor of 2-5.

Many spaceborne scanning cameras utilise TDI to increase their effective exposure time and maximise SNR. This thesis shows that, with compatible sensor hardware, this technique could be used with EnVisS, but over a significantly reduced across-track FOV (due to the wide-angle optics, and curvature of scene motion). Performing 8 TDI steps (increasing SNR by a factor of approximately 8) would reduce the instrument's usable FOV from ~180° to 80°. This low improvement in SNR would only be functional in relatively bright EnVisS images in which it was desirable to retain high resolution, and such images will likely only be obtained near to closest approach.

The signal analysis in this thesis models EnVisS noise as a combination of photon shot noise, thermal electron shot noise and Gaussian read noise. As the instrument continues its development, it will be necessary to characterise its noise more precisely, accounting for additional contributions such as stray light contamination (particularly given its wide FOV), in order to increase the confidence with which its SNRs can be predicted. EnVisS signals are on a knife's edge, and in the instrument's ongoing development, it will be necessary to mitigate the camera's low SNRs wherever possible.

Polarimetry is a valuable and often-used technique in cometary science to study properties – particularly the nature of dust grains and their evolution – of cometary atmospheres that cannot be revealed by photometry alone. EnVisS aims to perform passive imaging polarimetry of the comet's coma, with its all-sky FOR allowing close to the full 180° phase function of the comet to be characterised. Following the methodology of Chen et al. (2014), this thesis modelled the accuracy of EnVisS polarimetry measurements (DOLP and AOP) by Monte Carlo simulation, in order to compare the performance of two possible approaches: 3-filter and 4-filter polarimetry. The measurement accuracies achievable with each method are comparable, with neither performing markedly better than the other. This helps to inform focal plane design, where

it is expected that minimising the number of filters on the detector will be crucial. The error modelling indicates that images intended for polarimetry will require larger SNRs than the instrument's other images. To achieve a 10% error in the measurement of peak polarisation will likely require an image SNR >50, whilst detection of a negative polarisation branch will likely require SNRs exceeding 120. Binning will be essential for the instrument's imaging polarimetry.

The low expected signals in EnVisS images ultimately stem from the instrument being mounted on a spin stabilised spacecraft, limiting the exposure times with which it can image. But the image products it aims to acquire intend to provide scientific insight that is no doubt novel and valuable. Although not possible for EnVisS and the Comet Interceptor mission (whose designs are too mature), the key capabilities of the instrument – all-sky coverage, imaging polarimetry, multispectral imaging – could potentially be achieved with a different instrument design, that is also capable of imaging with significantly higher SNRs. A fisheye lens with a hemispherical FOV could image half the sky instantaneously. Mounted on a 3-axis stabilised spacecraft, exposure times of at least 100 ms would be possible, even at full imaging resolution and closest approach. Multispectral and polarimetric imaging would require filter wheels, whilst full-sky coverage would require two hemispherical FOVs with opposite view directions. A detailed study of such an instrument's capabilities, in particular its images' SNRs, would give insight into whether the improved performance outweighs the increased complexity and mass, and help to inform how and if future cometary missions could pursue all-sky imaging of large-scale cometary environments.

8.3 Supporting Camera Design with Image Simulation

The aims of this thesis at the outset were to study methods of employing spin-scanning imaging on planetary probes, characterise the capabilities of such methods, develop approaches to both capturing and analysing their data, and demonstrate the feasibility of such imagers. But it became evident early on in the research that these aims couldn't be fully met without access to the data that such cameras would capture. Neither of the two cameras that this thesis studied existed in a form that could actually capture images (the penetrator camera is a concept, EnVisS is in development and no test models yet exist), and this therefore necessitated a way of simulating the instruments' data.

Software and libraries already exist for simulating images, but their suitability for the early research of scientific camera concepts is variable. The work of this thesis therefore developed a simple approach to image rendering, that allows geometrically and

radiometrically representative images to be simulated, for demonstrating the nature of conceptual cameras' data. Specifically, both planetary surfaces and cometary dust were simulated, and the equations presented in chapter 7 could be readily adapted to simulate a wider range of planetary scenes such as atmospheres. Planetary surfaces are constructed from DTMs (optionally with an accompanying ORI for surface appearance), allowing images of a wide range of planetary bodies to be simulated. Cometary dust is described by a density distribution, with which complex, physically accurate scenes can be represented. The code's simulated surface images were used, for example, to demonstrate and study descent stereophotogrammetry (section 5.6 and Brydon et al. (Brydon, et al., 2021)), whilst its cometary images were used to inform decisions on approaches to SNR mitigation in EnVisS development.

The simulation approach developed for this thesis can be implemented entirely in Python, relying on only a small number of freely accessible libraries which are in common use for scientific computing, including in the planetary sciences. Keeping the simulation code as simple as possible, particularly by utilising NumPy, means that rendering times of high-resolution images can be kept relatively low (10s of seconds and below), in spite of Python's low execution speed when performing large numbers of calculations.

The description of the rendering approach in this thesis was deliberately kept general, because it is designed to be easily implemented and adapted for a user's specific needs. No code is presented, because the simulation methodology can readily be applied to languages other than Python. However, the raw Python code developed and used for this thesis is available to access online (see the Appendix). Formalising the rendering software into a distributable and well-documented Python library was unfortunately beyond the scope and time constraints of this thesis, but would be a valuable future project for facilitating the research of other camera concepts.

9 References

Abbey, W. et al., 2019. A look back: the drilling campaing of the Curiosity rover during the Mars Science Laboratory's Prime Mission. *Icarus,* Volume 319, pp. 1-13.

Aboudan, A., Colombatti, G., Ferri, F. & Angrilli, F., 2008. Huygens probe entry trajectory and attitude estimated simultaneously with Titan atmospheric structure by Kalman filtering. *Plan. Space Sci.,* Volume 56, pp. 573-585.

Accomazzo, A. et al., 2017. The final year of the Rosetta mission. *Acta Astronautica,* Volume 136, pp. 354-359.

A'Hearn, M. F. et al., 2011. EPOXI at comet Hartley 2. *Science,* Volume 332, pp. 1396-1400.

Albee, A. et al., 2000. *Report on the loss of the Mars Polar Lander and Deep Space 2 missions,* s.l.: JPL.

Anderson, J. A. & Robinson, M. S., 2009. *Challenges utilizing pushframe camera images.* s.l., s.n.

Anderson, J. D., 1986. The Bayeux tapestry: a 900-year-old Latin cartoon. *The Classical Journal,* Volume 81, pp. 253-257.

Anderson, J. L., Collins, S. A., Klaasen, K. P. & Reilly, T. H., 1991. Planetary Cameras. *International Journal of Imaging Systems and Technology,* Volume 3, pp. 67-75.

Aplin, K. L. et al., 2020. Atmospheric electricity at the ice giants. *Space Sci. Rev.,* Volume 216.

Arago, F., 1858. *Astronomie populaire.* Paris: Gide et Baudry.

Ardalan, A. A., Karimi, R. & Grafarend, E. W., 2010. A new reference equipotential surface, and reference ellipsoid for the planet Mars. *Earth, Moon and Planets,* Volume 106.

Argudo-Fernandez, M. et al., 2015. Catalogues of isolated galaxies, isolated pais, and isolated triplets in the local Universe. *Astronomy & Astrophysics,* Volume 578.

Artal, P., 2015. Image formation in the living human eye. *Annu. Rev. Vis. Sci.,* Volume 1, pp. 1-17.

Balaram, J., Aung, M. & Golombek, M. P., 2021. The Ingenuity Helicopter on the Perseverance Rover. *Space Sci. Rev.,* Volume 217.

Balsiger, H. et al., 1986. Ion composition and dynamics at comet Halley. *Nature,* Volume 321, pp. 330-334.

Barber, S. J. et al., 2018. L-DART: a penetrator mission for lunar permanently shadowed regions. *49th Lunar and Planetary Science Conference.*

Barnes, R. et al., 2018. Geological analysis of Martian rover-derived digital outcrop models using the 3-D visualization tool, Planetary Robotics 3-D Viewer—PRo3D. *Earth and Space Science,* Volume 5, pp. 285-307.

Barnouin, O. S. et al., 2020. Digital terrain mapping by the OSIRIS-REx mission. *Plan. Space Sci.,* Volume 180, p. 104764.

Basilevsky, A. T., 2012. Analysis of suspicious objects in TV panorama taken by Venera-9 spacecraft. *Solar System Research,* Volume 46, pp. 404-408.

Basilevsky, A. T. & Head, J. W., 2003. The surface of Venus. *Rep. Prog. Phys.,* Volume 66.

Beauchamp, P. et al., 2000. *Miniature Integrated Camera Spectrometer (MICAS) validation report,* s.l.: NASA Technology Validation Symposium.

Behounkova, M., Soucek, O., Hron, J. & Cadek, O., 2017. Plume activity and tidal deformation on Enceladus influenced by faults and variable ice shell thickness. *Astrobiology,* Volume 17, pp. 941-954.

Bell III, J. F. et al., 2021. The Mars 2020 Perseverance Rover Mast Camera Zoom (Mastcam-Z) multispectral, stereoscopic imaging investigation. *Space Sci. Rev.,* Volume 217.

Bell III, J. F. et al., 2004. Pancam multispectral imaging results from the Opportunity rover at Meridiani Planum. *Science,* Volume 306, pp. 1703-1709.

Bell III, J. F. et al., 2004. Pancam multispectral imaging results from the Spirit rover at Gusev crater. *Science,* Volume 305, pp. 800-806.

Bell III, J. F. et al., 2003. Mars Exploration Rover Athena Panoramic Camera (Pancam) investigation. *J. Geophys. Res.,* Volume 108.

Belton, M. J. S. et al., 1996a. Galileo's encoutner with 243 Ida: an overview of the imaging experiment. *Icarus,* Volume 120, pp. 1-19.

Belton, M. J. S. et al., 1996b. Galileo's first images of Jupiter and the Galilean satellites. *Science,* Volume 274, pp. 377-385.

Bertini, I. et al., 2017. The scattering phase function of comet 67P/Churyumov-Gerasimenko coma as seen from the Rosetta/OSIRIS instrument. *MNRAS,* Volume 469, pp. 404-415.

Bhat, D. N. & Nayar, S. K., 1998. Ordinal Measures for Image Correspondence. *IEEE Trans. Pattern Anal. Mach. Intell,* Volume 20, pp. 415-423.

Bibring, J. -P. et al., 2015. 67P/Churyumov-Gerasimenko surface properties as derived from CIVA panoramic images. *Science,* p. 349.

Bibring, J. P. et al., 2007. CIVA. *Space Science Reviews,* Volume 128, pp. 397-412.

Bibring, J. -P. et al., 2007. The Rosetta lander ("Philae") investigations. *Space. Sci. Rev.,* Volume 128, pp. 205-220.

Biele, J. et al., 2015. The landing(s) of Philae and inferences about comet surface mechanical properties. *Science,* Volume 349.

Bierhaus, E. B., Chapman, C. R. & Merline, W. J., 2005. Secondary craters on Europa and implications for cratered surfaces. *Nature,* Volume 437, pp. 1125-1127.

Blume, W. H., 2005. Deep Impact Mission Design. *Space. Sci. Rev.,* Volume 117, pp. 23-42.

Bockelée-Morvan, D., 2011. An overview of comet composition. *Proceedings of the International Astronomical Union,* Volume 7, pp. 261-274.

Boehnhardt, H. et al., 2017. The Philae lander mission and science overview. *Philos. Trans. R. Soc. A.,* Volume 375.

Boice, D. C. et al., 2002. The Deep Space 1 encounter with comet 19p/Borrelly. *Earth, Moon and Planets,* Volume 89, pp. 301-324.

Bolton, S. J. et al., 2017. The Juno mission. *Space Sci. Rev.,* Volume 213, pp. 5-37.

Bos, B. J. et al., 2018. Touch and Go Camera System (TAGCAMS) for the OSIRIS-REx asteroid sample return mission. *Space Sci. Rev,* Volume 214.

Botley, F. V., 1951. A new use for the Plate Carree projection. *Geopgraphical Review,* pp. 640-644.

Boulos, S. et al., 2007. Packet-based whitted and distribution ray tracing. *Proc. Graphics Interface 2007,* pp. 177-184.

Boyle, W. S. & Smith, G. E., 1970. Charge coupled semiconductor devices. *Bell Sys. Tech. J.,* Volume 49, pp. 587-593.

Boynton, W. V. & Reinert, R. P., 1995. The cryo-penetrator: an approach to exploration of icy bodies in the solar system. *Acta Astronautica,* Volume 35, pp. 59-68.

Bradski, G., 2000. The OpenCV library. *Dr. Dobb's Journal of Software Tools.*

Brandt, J. C., 1968. The physics of comet tails. *Ann. Rev. Astron. Astrophys,* Volume 6, pp. 267-286.

Brydon, G. & Jones, G. H., 2018. Rotational push-broom imaging from a planetary penetrator. *EPSC meeting, EPSC2018-994,* Volume 12.

Brydon, G., Persaud, D. M. & Jones, G. H., 2021. Planetary topography measurement by descent stereophotogrammetry. *Plan. Space Sci.,* Volume 202.

Burke, B. E. et al., 1992. High-performance visible/UV CCD imagers for space-based applications. *Proc. SPIE ,* Volume 1693, pp. 86-100.

Burrows, C. J. et al., 1994. *Hubble Space Telescope Wide Field and Planetary Camera 2 instrument handbook,* s.l.: Space Telescope Science Insititue.

Burrows, D. N. et al., 2005. The Swift X-ray telescope. *Space Science Reviews,* Volume 120, pp. 165-195.

Caplinger, M., 2014. Juno JunoCam EDR V1.0. *NASA Planetary Data System, JUNO-J-JUNOCAM-2-EDR-L0-V1.0.*

Caplinger, M. A., 2002. Mars Orbiter Camera global mosaic. *Lunar and Planetary Science XXXIII abstract,* p. 1405.

Carcich, B. e. a., 2014. *Stardust NAVCAM calibrated images of 81P/WILD 2 - VERS 3.0, SDU-C-NAVCAM-3-RDR-WILD2-V3.0,* s.l.: NASA Planetary Data System.

Carr, M. H. et al., 1995. The Galileo Imaging Team plan for observing the satellites of Jupiter. *Journal of Geophysical Research: Planets,* Volume 100, pp. 18935-18955.

Carr, M. H. et al., 1976. Preliminary results from the Viking Orbiter Imaging Experiment. *Science,* Volume 193, pp. 766-776.

Chandrasekhar, S., 1960. *Radiative Transfer.* New York: Dover.

Chapman, C. R., Pollack, J. B. & Sagan, C., 1968. A analysis of the Mariner 4 photography of Mars. *SAO Special Report,* Volume 268.

Chassefiere, E. et al., 2009. European Venus Explorer: an in-situ mission to Venus using a balloon platform. *Adv. Space Res,* Volume 44, pp. 106-115.

Cheng, A. F. et al., 2002. Small-scale topography of 433 Eros from laster altimetry and imaging. *Icarus,* Volume 155, pp. 51-74.

Chen, Z., Wang, X., Pacheco, S. & Liang, R., 2014. Impact of CCD camera SNR on polarimetric accuracy. *Applied Optics,* Volume 53, pp. 7649-7656.

Christensen, P. R. et al., 2004. The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey mission. *Space Sci. Rev.,* Volume 110, pp. 85-130.

Chun, J.-B., Jung, H. & Kyung, C.-M., 2008. Suppressing rolling-shutter distortion of CMOS image sensors by motion vector detection. *IEEE Transactions on Consumer Electronics,* Volume 54, pp. 1479-1487.

Clark, B. C. et al., 2004. Release and fragmentation of aggregates to produce heterogeneous, lumpy coma streams. *J. Geophys. Res.,* Volume 109.

Clark, B. E. et al., 1998. Multispectral terrain analysis of Europa from Galileo images. *Icarus,* Volume 135, pp. 95-106.

Coates, A. J., 2004. Ion pickup at comets. *Advances in Space Research,* Volume 33, pp. 1977-1988.

Coates, A. J. et al., 2017. The PanCam instrument for the ExoMars rover. *Astrobiology,* Volume 17, pp. 511-541.

Coates, A. J., Mazelle, C. & Neubauer, F. M., 1997. Bow shock analysis at comets Halley and Grigg-Skjellerup. *JGR,* Volume 102, pp. 7105-7113.

Combi, M. R. et al., 2012. Narrow dust jets in a diffuse gas coma: a natural product of small active regions on comets. *Astrophys. J.,* Volume 749.

Cook, N. V., Ragozzine, D., Granvik, M. & Stephens, D. C., 2016. Realistic detectability of close interstellar comets. *Astrophys. J.,* Volume 825.

Cravens, T. E. & Gombosi, T. I., 2004. Cometary magnetospheres: a tutorial. *Advances in Space Research,* Volume 33, pp. 1968-1976.

Cremonese, G. et al., 1997. Neutral sodium from coemt Hale-Bopp: A third type of tail. *Astrophys. J,* Volume 490.

Crisp, J. A. et al., 2003. Mars Exploration Rover mission. *J. Geophys. Res.,* Volume 108.

D'Amario, L. A., Bright, L. E. & Wolf, A. A., 1992. Galileo trajectory design. *Space Sci. Rev.,* Volume 60, pp. 23-78.

Datta, A., Kim, J. S. & Kanade, T., 2009. Accurate camera calibration using iterative refinement of control points. *IEEE 12th International Conference on Computer Vision Workshops,* pp. 1201-1208.

Di, K. et al., 2014. Mars surface change detection from multi-temporal orbital images. *IOP Conference Series: Earth and Environmental Science,* Volume 17.

Di, K. et al., 2013. A self-calibration bundle adjustment method for photogrammetric processing of Chang'E-2 stereo Lunar imagery. *IEEE Transactions on Geoscience and Remote Sensing,* Volume 52, pp. 5432-5442.

Dobrovolsky, O. V., Kiselev, N. N. & Chernova, G. P., 1986. Polarimetry of comets: a review. *Earth, Moon and Planets,* Volume 34, pp. 189-200.

Donati, S., Martini, G. & Norgia, M., 2007. Microconcentrators to recover fill-factor in image photodetectors with pixel on-board processing circuits. *Optics Express,* Volume 15, pp. 18066-18075.

Dougherty, L. M. & Dollfus, A., 1989. Arago's polarimeter and his original observation of extraterrestrial polarisation in 1811. *Journal of the British Astronomical Association,* Volume 99, pp. 183-186.

Douxchamps, D. & Chihara, K., 2008. High-accuracy and robust localization of large control markers for geometric camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell,* Volume 31, pp. 376-383.

Duncan, M., Quinn, T. H. & Tremaine, S., 1988. The origin of short-period comets. *Astrophys. J.,* Volume 328, pp. 69-73.

Duxbury, T. C. et al., 2004. Asteroid 5535 Annefrank size, shape, and orientation: Stardust first results. *J. Geophys. Res.,* Volume 109.

Dyudina, U. A. et al., 2004. Lightning on Jupiter observed in the Hα line by the Cassini imaging science subsystem. *Icarus,* Volume 172, pp. 24-36.

Edmundson, K. L. et al., 2012. Jigsaw: the ISIS3 bundle adjustment for extraterrestrial photogrammetry. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci,* Volume 1, pp. 203-208.

Edwards, C. S. et al., 2011. Mosaicking of global planetary image datasets: 1. Techniques and data processing for Thermal Emission Imaging System (THEMIS) multi-spectral data. *J. Geophys. Res.,* Volume 116.

Elle, B. L., Heinmiller, C. S., Fromme, P. J. & Neumer, A. E., 1967. The Lunar Orbiter photographic system. *Journal of the SMPTE,* Volume 76, pp. 733-739.

El-Maarry, M. R. et al., 2019. Surface morphology of comets and associated evolutionary processes: A review of Rosetta's observations of 67P/Churyumov-Gerasimenko. *Space. Sci. Rev.,* Volume 215.

El-Maarry, M. R. et al., 2015. Fractures on comet 67P/Churyumov-Gerasimenko observed by Rosetta/OSIRIS. *Geophys. Res. Lett.,* Volume 42, pp. 5170-5178.

Elsila, J. E., Glavin, D. P. & Dworkin, J. P., 2009. Cometary glycine detected in samples returned by Stardust. *Meteoritics & Planetary Science,* Volume 44, pp. 1323-1330.

Emerich, C. et al., 1988. Temperature and size of the nucleus of comet P/Halley deduced from IKS infrared Vega 1 measurements.. In: M. Grewing, F. Praderie & R. Reinhard, eds. *Exploration of Halley's Comet.* Berlin: Springer, pp. 839-842.

Enoch, J. M., 2006. History of mirrors dating back 8000 years. *Optometry and Vision Science,* Volume 83, pp. 775-781.

Farnham, T., 2021. Comet C/2014 UN271 (Bernardinelli-Bernstein) exhibited activity at 23.8 AU. *The Astronomer's Telegram,* p. ATel #14759.

Farnham, T. L., 2009. Coma morphology of Jupiter-family comets. *Plan. Space Sci.,* Volume 57, pp. 1192-1217.

Farnham, T. L. et al., 2013. Connections between the jet activity and surface features on Comet 9P/Tempel 1. *Icarus,* Volume 222, pp. 540-549.

Farnham, T. L. et al., 2007. Dust coma morphology in the Deep Impact images of Comet9P/Tempel 1. *Icarus,* Volume 187, pp. 26-40.

Felipe-Garcia, B., Hernandez-Lopez, D. & Lerma, J. L., 2012. Analysis of the ground sample distance on large photogrammetric surveys. *Appl Geomat,* Volume 4, pp. 231-244.

Feng, W., Ji, Y. & Chen, L., 2013. The impact of signal-noise ratio on degree of linear polarization measurement. *Optik,* Volume 124, pp. 192-194.

Fergason, R. L., Hare, T. M. & Laura, J., 2018. HRSC and MOLA blended digital elevation model at 200m v2. *Astrogeology PDS Annex, US Geological Survey.*

Fernández, Y. R. et al., 2003. The nucleus of Deep Impact target Comet 9P/Tempel 1. *Icarus,* Volume 164, pp. 481-491.

Florensky, K. P. et al., 1982. Analysis of the panoramas of the Venera 13 and Venera 14 landing sites. *Sov. Astron. Lett.,* Volume 8, pp. 233-234.

Flynn, G. J., 2008. Physical, chemical, and mineralogical properties of Comet 81P/Wild 2 particles collected by Stardust. *Earth Moon Planet,* Volume 102, pp. 447-459.

Folkner, W. M. et al., 2006. Winds on Titan from ground-based tracking of the Huygens probe. *J. Geophys. Res.,* Volume 111.

Fossum, E. R., 1993. Active pixel sensors: Are CCDs dinosaurs?. *Proc. SPIE,* Volume 1900, pp. 2-14.

Fossum, E. R., 1997. CMOS image sensors: Electronic camera-on-a-chip. *IEEE Trans. Electron Devices,* Volume 44, pp. 1689-1698.

Fountain, J. & Gehrels, T., 1977. Pioneer images of Jupiter. *Highlights of Astronomy,* Volume 4, pp. 233-241.

Francis, P. J., 2005. The demographics of long-period comets. *Astrophys. J.,* Volume 635, pp. 1348-1361.

Frauenholz, R. B. et al., 2008. Deep Impact navigation system performance. *Journal of Spacecraft and Rockets,* Volume 45, pp. 39-56.

Fujiwara, A. et al., 2006. The rubble-pile asteroid Itokawa as observed by Hayabusa. *Science,* Volume 312, pp. 1330-1334.

Fulle, M., 2004. Motion of cometary dust. In: M. C. Festou, H. U. Keller & H. A. Weaver Jr., eds. *Comets II.* Tucson: University of Arizona Press, pp. 565-575.

Fulle, M. et al., 2007. Discovery of the atomic iron tail of comet McNaught using the Heliospheric Imager on STEREO. *Astrophys. J.,* Volume 661.

Furukawa, Y. & Ponce, J., 2009. Accurate camera calibration from multi-view stereo and bundle adjustment. *International Journal of Computer Vision,* Volume 84, pp. 257-268.

Gao, Y. et al., 2008. Lunar science with affordable small spacecraft technologies: MoonLITE and Moonraker. *Plan. Space Sci.,* Volume 56, pp. 368-377.

Garachorloo, N., Gupta, S., Sproull, R. F. & Sutherland, I. E., 1989. A characterization of ten rasterization techniques. *Proc. 16th Annual Conf. Computer Graphics and Interactive Techniques,* pp. 355-368.

Garvin, J. B. et al., 2020. DAVINCI+: deep atmosphere of Venus investigation of noble gases, chemistry, and imaging plus. *51st Lunar and Planetary Science Conference March 2020,* Volume 2326, p. 2599.

Gehrels, T. et al., 1980. Imaging Photopolarimeter on Pioneer Saturn. *Science,* Volume 207, pp. 434-439.

Gerig, S. -B. et al., 2018. On deviations from free-radical outflow in the inner coma of comet 67P/Churyumov-Gerasimenko. *Icarus,* Volume 311, pp. 1-22.

Gianinetto, M. & Scaioni, M., 2008. Automated geometric correction of high-resolution pushbroom satellite data. *PE&RS,* Volume 1, pp. 107-116.

Gibson, B. K. & Hickson, P., 1992. Time-delay integration CCD read-out technique: image deformation. *Mon. Not. R. astr. Soc.,* Volume 258, pp. 543-551.

Giese, B. et al., 2006. Topographic modeling of Phoebe using Cassini images. *Plan. Space Sci.,* Volume 54, pp. 1156-1166.

Gilard, O. et al., 2008. New approach for the prediction of CCD dark current distribution in a spaec adiation environment. *IEEE Transactions on Nuclear Science,* Volume 55, pp. 3626-3632.

Giovane, F. et al., 1991. Photomultiplier for optically probing Comet Halley. *Applied Optics,* Volume 30, pp. 2579-2591.

Gladstone, G. R. et al., 2016. The atmosphere of Pluto as observed by New Horizons. *Science,* Volume 351.

Glassmeier, K. -H. et al., 2007. The Rosetta mission: Flying towards the origin of the Solar System. *Space Sci. Rev.,* Volume 128, pp. 1-21.

Gloeckler, G. et al., 2000. Interception of comet Hyakutake's ion tail at a distance of 500 million kilometres. *Nature,* Volume 404, pp. 576-578.

Godet, O. et al., 2009. Modellng the spectral response of the Swift-XRT CCD camera: experience learnt from in-flight calibration. *Astron. Astrophys.,* Volume 494, pp. 775-797.

Gowen, R. A. et al., 2011. Penetrators for in situ subsurface investigations of Europa. *Adv. Space Res.,* Volume 48, pp. 725-742.

Grasset, O. et al., 2013. JUpiter ICy moons Explorer (JUICE): an ESA mission to orbit Ganymede and to characterise the Jupiter system. *Plan. Space Sci.,* Volume 78, pp. 1-21.

Greeley, R., 1994. *Planetary Landscapes.* s.l.:Springer US.

Greeley, R., Klemaszewski, J. E., Wagner, R. & Galileo Imagin Team, 2000. Galileo views of the geology of Callisto. *Plan. Space Sci.,* Volume 48, pp. 829-853.

Greeley, R. et al., 1998. Europa: initial Galileo geological observations. *Icarus,* Volume 135, pp. 4-24.

Grensemann, M. G. & Schwehm, G., 1993. Giotto's second encounter: the mission to comet P/Grigg-Skjellerup. *JGR,* Volume 98, pp. 20907-20910.

Griffiths, M., 2014. A brief history of astronomical Imaging. In: *Choosing and Using Astronomical Filters.* New York: Springer, pp. 1-8.

Grotzinger, J. P. et al., 2012. Mars Science Laboratory mission and science investigation. *Space Science Reviews,* Volume 170, pp. 5-56.

Gupta, R. & Hartley, R. I., 1997. Linear pushbroom cameras. *IEEE Trans. Pattern Anal. Mach. Intell.,* Volume 19, pp. 963-975.

Gwyn, S. D., 2008. MegaPipe: The MegaCam image stacking pipeline at the Canadian astronomical data centre. *Publications of the Astronomical Society of the Pacific,* Volume 120.

Hadamcik, E. & Levasseur-Regourd, A. C., 2003. Imaging polarimetry of cometary dust: different comets and phase angles. *J. Quant. Spectrosc. Radiat. Transf.,* Volume 79, pp. 661-678.

Hadamcik, E., Levasseur-Regourd, A. C., Leroi, V. & Bardin, D., 2007. Imaging polarimetry of the dust coma of Comet Tempel 1 before and after Deep Impact at Haute-Provence Observatory. *Icarus,* Volume 190, pp. 459-468.

Hallert, B., 1960. *Photogrammetry basic principles and general survey.* 1 ed. New York: McGraw-Hill.

Hampton, D. L. et al., 2005. An overview of the instrument suite for the Deep Impact mission. *Space Sci. Rev,* Volume 117, pp. 43-93.

Hansen, C. J. et al., 2017. Junocam: Juno's outreach camera. *Space Science Reviews,* Volume 213, pp. 475-506.

Harris, C. R. et al., 2020. Array programming with NumPy. *Nature,* Volume 585, pp. 357-362.

Hasegawa, I., 1979. Orbits of Ancient and Medieval Comets. *Publ. Astron. Soc. Japan,* Volume 31, pp. 257-270.

Heacock, R. L., 1968. Lunar photograpghy: techniques and results. *Space Science Reviews,* Volume 8, pp. 214-257.

Healey, G. E. & Kondepudy, R., 1994. Radiometric CCD camera calibration and noise estimation. *IEEE Trans. Pattern Anal. Mach. Intell,* Volume 16, pp. 267-276.

Heck, A., 1978. Some methods of determining the stellar absolute magnitude. *Vistas in Astronomy,* Volume 22, pp. 221-264.

Heidarzadeh, T., 2008. *A history of physical theories of comets, from Aristotle to Whipple.* s.l.:Springer.

Heipke, C. et al., 2007. Evaluating planetary digital terrain models—The HRSC DTM test. *Plan. and Space Sci,* Volume 55, pp. 2173-2191.

Henriksen, M. R. et al., 2017. Extracting accurate and precise topography from LROC narrow angle camera stereo observations. *Icarus,* Volume 283, pp. 122-137.

Heubner, W. F. et al., 1988. A model for intensity profiles of dust jets near the nucleus of comet Halley. *Icarus,* Volume 76, pp. 78-88.

Hirschmuller, H., 2008. Stereo processing by semiglobal matching and mutual information. *Pattern Analysis and Machine Learning, IEEE Transactions on,* Volume 30, pp. 328-341.

Hirt, C., Claessens, S. J., Kuhn, M. & Featherstone, W. E., 2012. Kilometer-resolution gravity field of Mars: MGM2011. *Plan. and Space Sci.,* Volume 67, pp. 147-154.

Hobson, S. M., Trundle, K. C. & Sackes, M., 2010. Using a planetarium software program to promote conceptual change with young children. *J. Sci. Educ. Technol.,* Volume 19, pp. 165-176.

Hopkinson, G. R., Mohammadzadeh, A. & Harboe-Sorensen, R., 2004. Radiation effects on a radiation-tolerant CMOS active pixel sensor. *IEEE Transactions on Nuclear Science,* Volume 51, pp. 2753-2762.

Huggins, W., 1882. Preliminary note on the photographic spectrum of Comet b 1881. *Proceedings of the Royal Society of London,* Volume 33, pp. 216-219.

Hughes, D. W., 1987. The history of Halley's Comet. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences,* Volume 323, pp. 349-367.

Hui, M. -T., 2013. Observations of Comet P/2003 T12 = 2012 A3 (SOHO) at large phase angle in STEREO-B. *Mon. Not. R. Astron. Soc,* Volume 436, pp. 1564-1575.

Hui, M. T. & Ye, Q. Z., 2020. Observations of disintegrating long-period comet C/2019 Y4 (ATLAS): A sibling of C/1844 Y1 (Great Comet). *The Astronomical Journal,* Volume 160.

Huntress, W. T., Moroz, V. I. & Shevalev, I. L., 2003. Lunar and planetary robotic exploration missions in the 20th century. *Space Science Reviews,* Volume 107, pp. 541-649.

Jackson, B. V. et al., 2013. Using comet plasma tails to study the solar wind. *AIP Conference Proceedings,* Volume 1539, pp. 364-369.

Jaffe, L. D., 1969. Recent observations of the Moon by spacecraft. *Space Science Reviews,* Volume 9, pp. 491-616.

Jaumann, R. et al., 2007. The High Resolution Stereo Camera (HRSC) experiment on Mars Express: Instrument aspects and experiment conduct from interplanetary cruise through nominal mission. *Plan. and Space Sci.,* Volume 55, pp. 928-957.

Jaumann, R. et al., 2016. The camera of the MASCOT asteroid lander on board Hayabusa 2. *Space Sci. Rev.,* Volume 208.

Jia, Y. et al., 2018. The scientific objectives and payloads of Chang'E-4 mission. *Plan. and Space Sci.,* Volume 162, pp. 207-215.

Jingjing, W., Zhenchuan, Z. & Yong, T., 2007. A method of CCD noise removal in digital images. *2nd IEEE Conference on Industrial Electronics and Applications,* pp. 2571-2574.

Jockers, K. et al., 2005. CCD imaging and aperture polarimetry of comet 2P/Encke: are there two polarimetric classes of comets?. *A&A,* Volume 441, pp. 773-782.

Johnsson, K., 1960. On the accuracy of stereo-plotting of convergent aerial photographs. *Photogrammetria,* Volume 17, pp. 83-98.

Jolliff, B. L., 1999. Clementine UVVIS multispectral data and the APollo 17 landing site: What can we tell and how well?. *J. Geohpys. Res.,* Volume 104, pp. 14,123-14,148.

Jones, G. H., Balogh, A. & Horbury, T. S., 2000. Identification of comet Hyakutake's extremely long ion tail from magnetic field signatures. *Nature,* Volume 404, pp. 574-576.

Jones, G. H. et al., 2017. The Akon Europa penetrator. *14th International Planetary Probe Workshop, Abstract.*

Josset, J. L. et al., 2017. The Close-Up Imager onboard the ESA ExoMars rover: pbjectives, description, operations, and science validation activities. *Astrobiology,* Volume 17, pp. 595-611.

Jun, L. & Ling, L., 2010. Comparative research on Python speed optimization strategies. *International Conference on Intelligent Computing and Integrated Systems IEEE,* pp. 57-59.

Kang, M., Wang, M. & Du, Q., 2015. A method of DTM construction based on Quadrangular Irregular Networks and related error analysis. *PLoS ONE,* Volume 10.

Karkoschka, E. et al., 2007. DISR imaging and the geometry of the descent of the Huygens probe within Titan's atmosphere. *Plan. and Space Sci,* Volume 55, pp. 1896-1935.

Keldysh, M. V., 1977. Venus explortion with the Venera 9 and Venera 10 spacecraft. *Icarus,* Volume 30, pp. 605-625.

Keller, H. U. et al., 1986. First Halley Multicolour Camera imaging results from Giotto. *Nature,* Volume 321, pp. 321-326.

Keller, H. U. et al., 2007. OSIRIS - The scientific camera system onboard Rosetta. *Space Sci. Rev,* Volume 128, pp. 433-506.

Keller, H. U. & Kührt , E., 2020. Cometary nuclei-from Giotto to Rosetta. *Space Sci. Rev.,* Volume 216.

Keller, H. U. et al., 2017. Seasonal mass transfer on the nucleus of comet 67P/Churyumov-Gerasimenko. *Monthly Notices of the Royal Astronomical Society,* Volume 469, pp. S357-S371.

Keller, H. U. et al., 1987. The Halley multicolour camera. *J. Phys. E: Sci.,* Volume 20, pp. 807-820.

Kelley, M. S. et al., 2013. A disribution of large particles in the coma of Comet 103P/Hartley 2. *Icarus,* Volume 222, pp. 634-652.

Keszthelyi, L. et al., 2008. High Resolution Imaging Science Experiment (HiRISE) images of volcanic terrains from the first 6 months of the Mars Reconnaissance Orbiter primary science phase. *JGR Plan,* Volume 113.

Keszthelyi, L. et al., 2001. Imaging volacanic activity on Jupiter's moon Io by Galileo during the Galileo Europa mission and the Galileo Millenium mission. *J. Geophys. Res.,* Volume 106, pp. 33,025-33,052.

Kirk, R. L., Barret, J. M. & Soderblom, L. A., 2003. Photoclinometry made simple. *Proc. ISPRS Commission IV Symposium,* Volume 34.

Kirk, R. L., Howington-Kraus, E., Hare, T. M. & Jorda, L., 2016. The effect of illumination on stereo DTM quality: simulations in support of Europa exploration. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences,* Volume III-4, pp. 103-110.

Kiselev, N. N., Jockers, K. & Bonev, T., 2004. CCD imaging polariemtry of Comet 2P/Encke. *Icarus,* Volume 168, pp. 395-391.

Kiselev, N. & Rosenbush, V., 2004. Polarimetry of comets: Progress and problems. *Photopolarimetry in Remote Sensing. NATO Science Series II: Mathematics, Physics and Chemsitry,* Volume 161, pp. 411-430.

Kiselev, N. et al., 2005. Database of comet polarimetry: analysis and some results. *Earth, Moon and Planets,* Volume 97, pp. 365-378.

Kissel, J. et al., 1986. Composition of comet Halley dust particles from Vega observations. *Nature,* Volume 321, pp. 280-282.

Knight, M. & Battams, K., 2014. Preliminary analysis of SOHO/STEREO observastions of sungrazing Comet ISON (C/2012 S1) around perihelion. *Astrophys. J Letters,* Volume 782.

Knollenberg, J. et al., 2016. A mini outburst from the nightside of comet 67P/Churyumov-Gerasimenko observed by the OSIRIS camera on Rosetta. *A&A,* Volume 596.

Kramer, T. & Noack, M., 2016. On the origin of inner coma structures observed by Rosetta during a diurnal rotation of comet 67P/Churyumov-Gerasimenko. *Astophys. J. Lett.,* Volume 823.

Ksanfomality, L. V., 2012. Results of the new processing of images obtained from the surface of Venus in a TV experiment onboard the VENERA-9 lander (1975). *Solar System Research,* Volume 46, pp. 364-373.

Kugler, A., 1996. The setup for triangle rasterization. *Workshop on Graphics Hardware,* pp. 49-58.

Läbe, T. & Förstner, W., 2006. Automatic relative orientation of images. *Proceedings of the 5th Turkish-German Joint Geodetic Days,* Volume 29.

Laine, S. & Karras, T., 2011. High-performance software rasterization on GPUs. *Proc. ACM SIGGRAPH Symposium on High Performance Graphics,* pp. 79-88.

Lallement, R., Bertaux, J. -L., Szego, K. & Nemeth, S., 2002. The shadow of comet Hale-Bopp in Lyman-Alpha. *Earth, Moon and Planets,* Volume 90, pp. 67-76.

Lamy, P. & Toth, I., 2009. The colors of cometary nuclei-comparison with other primitive bodies of the Solar System and implications for their origins. *Icarus,* Volume 201, pp. 674-713.

Lara, L. M. et al., 2015. Large-scale dust jets in the coma of 67P/Churyumov-Gerasimenko as seen by the OSIRIS instrument onboard Rosetta. *A&A,* Volume 583.

Le Roy, L. et al., 2015. Inventory of the volatiles on comet 67P/Churyumov-Gerasimenko from Rosetta/ROSINA. *A&A,* Volume 583.

Lebreton, J. P. et al., 2005. An overview of the descent and landing of the Huygens probe on Titan. *Nature,* Volume 438, pp. 758-764.

Leighton, R. B. et al., 1965. Mariner IV photography of Mars: inital results. *Science,* Volume 149, pp. 627-630.

Leprince, S., Barbot, S., Ayoub, F. & Avouac, J.-P., 2007. Automatic and precise orthorectification, coregistration, and subpixel correlation of satellite images, application to ground deformation measurements.. *IEEE Transactions on Geoscience and Remote Sensing,* Volume 45, pp. 1529-1558.

Lester, T. P., McCall, M. L. & Tatum, J. B., 1979. Theory of planetary photometry. *The Journal of the Royal Astronomical Society of Canada,* Volume 73, pp. 233-257.

Levasseur-Regourd, A. C. et al., 1984. A spectral photopolarimeter for Giotto: Halley Optical Probe Experiment. *Adv. Space Res.,* Volume 4, pp. 287-290.

Levasseur-Regourd, A. C. et al., 1993. Short Communication: Optical probing of the dust in comet Grigg-Skjellerup from the Giotto spacecraft. *Plan. Space Sci.,* Volume 41, pp. 167-169.

Levasseur-Regourd, A. C., Hadamcik, E. & Renard, J. B., 1996. Evidence for two classes of comets from their polarimetric properties at large phase angles. *Astron. Astrophys,* Volume 313, pp. 327-333.

Levasseur-Regourd, A. C., McBride, N., Hadamcik, E. & Fulle, M., 1999. Similarities between in situ measurements of local dust light scattering and dust flux impace data within the coma of 1P/Halley. *Astron. Astrophys.,* Volume 348, pp. 636-641.

Levison, H. F., 1996. Comet Taxonomy. *Astronomical Society of the Pacific Conference Proc.,* Volume 107, pp. 173-191.

Li, C. et al., 2010. The global image of the Moon obtained by the Chang'E-1: Data processing and lunar cartography. *Science China Earth Sciences,* Volume 53, pp. 1091-1102.

Li, C. et al., 2015. The Chang'e 3 mission overview. *Space Sci. Rev.,* Volume 190, pp. 85-101.

Lindberg, D. C., 1968. The theory of pinhole images from antiquity to the thirteenth century. *Archive for History of Exact Sciences,* Volume 5, pp. 154-176.

Linder, D. R. et al., 1998. The Cassini CAPS electron spectrometer. *Geophsyical Monograph-American Geophysical Union,* Volume 102, pp. 257-262.

Lindler, D. J. et al., 2013. Interpretation of results of deconvolved images from the Deep Impact spacecraft High Resolution Instrument. *Icarus,* Volume 222, pp. 571-579.

Lindler, D. J., A'Hearn, M. F. & McLaughlin, S. A., 2012. *EPOXI 103P/Hartley 2 Encounter - HRIV Deconvolved IMGS V1.0, DIF-C-HRIV-5-EPOXI-HARTLEY2-DECONV-V1.0,* s.l.: NASA Planetary Data System.

Lin, Z. -Y., Lin, C. -S., Ip, W. -H. & Lara, L. M., 2009. The outburst of comet 17P/Holmes. *Astron. J,* Volume 138.

Li, R. et al., 2002. Localization of Mars rovers using descent and surface-based image data. *JGR,* Volume 107.

Little, B. et al., 1999. Galileo images of lightning on Jupiter. *Icarus,* Volume 142, pp. 306-323.

Litwiller, D., 2011. *CCD vs CMOS: Facts and fiction,* s.l.: DALSA.

Liu, J. et al., 2019. Descent trajectory reconstruction and landing site positioning of Chang'E-4 on the lunar farside. *Nat Commun,* Volume 10.

Liu, Z. Q. et al., 2015. High precision landing site mapping and rover localization for Chang'e-3 mission. *Science China Physics, Mechanics & Astronomy,* Volume 58, pp. 1-11.

Li, Z., 1988. On the measure of digtal terrain model accuracy. *Photogrammetric Record,* Volume 12, pp. 873-877.

Li, Z., 1991. Effects of check points on the reliability of DTM accuracy estimates obtained from experimental tests. *Photogrammetric Engineering & Remote Sensing,* Volume 57, pp. 1333-1340.

Lorenz, R. D., 2010. Attitude and angular rates of planetary probes during atmospheric descent: implications for imaging. *Plan. and Space Sci,* Volume 58, pp. 838-846.

Lorenz, R. D., 2011. Planetary penetrators: their origins, history and future. *Advances in Space Research,* Volume 48, pp. 403-431.

Lorenz, R. D. et al., 2021. Selection and characteristics of the Dragonfly landing site near Selk Crater, Titan. *Planet. Sci. J.,* Volume 2.

Malin, M. C. et al., 2007. Context Camera investigation on board the Mars Reconnaissance Orbiter. *J. Geophys. Res.,* Volume 112.

Malin, M. C. et al., 2001. Mars Descent Imager (MARDI) on the Mars Polar Lander. *J. Geophys. Res.,* Volume 106, pp. 17635-17650.

Malin, M. C. & Edgett, K. S., 2001. Mars Global Surveyor Mars Orbiter Camera: interplanetary cruise through primary mission. *JGR,* Volume 106, pp. 23429-23570.

Malin, M. C. et al., 2017. The Mars Science Laboratory (MSL) Mast cameras and Descent imager: Investigation and instrument descriptions. *Earth Space Sci.,* Volume 4, pp. 506-539.

Malling, L. R. & Allen, J. D., 1966. The slow-scan vidicon as an interplanetary imaging device. *Advances in Electronics and Electron Physics,* Volume 22, pp. 835-847.

Mann, A. W. & von Braun, K., 2015. Revised filter profiles and zero points for broadband photometry. *Astronomical Society of the Pacific,* Volume 127.

Mannel, T. et al., 2016. Fractal cometary dust - a window into the early Solar system. *Mon. Not. R. AStron. Soc.,* Volume 462, pp. 304-311.

Marschall, R. et al., 2020. The dust-to-gas ratio, size distribution, and dust fall-back fraction of comet 67P/Churyumov-Gerasimenko: Inferences from linking the optical and dynamical properties of the inner comae. *arXiv:2005.13700 [astro-ph.EP].*

Marschall, R. et al., 2019. A comparison of multiple Rosetta data sets and 3D model calculations of 67P/Churyumov-Gerasimenko coma around equinox (May 2015). *Icarus,* Volume 328, pp. 104-126.

Marschall, R. et al., 2016. Modelling observations of the inner gas and dust coma of comet 67P/Churyumov-Gerasimenko using ROSINA/COPS and OSIRIS data: First results. *Astron. Astrophys.,* Volume 589.

Matijevic, J. & Shirley, D., 1997. The mission and operation of the Mars pathfinder microrover. *Control Engineering Practice,* Volume 5, pp. 827-835.

Max, N., 1995. Optical models for direct volume rendering. *IEEE Trans. Visualization Comp. Graphics,* Volume 1, pp. 99-108.

McCauley, J. F. et al., 1972. Preliminary Mariner 9 report on the geology of Mars. *Icarus,* Volume 17, pp. 289-327.

McDonnell, J. A. M. et al., 1993. Dust particle impacts during the Giotto encounter with comet Grigg-Skjellerup. *Nature,* Volume 362, pp. 732-734.

McEwen, A. S. et al., 2000. Galileo at Io: results from high-resolution imaging. *Science,* Volume 288, pp. 1193-1198.

McEwen, A. S. et al., 2007. Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE). *J. Geophys. Res.,* Volume 112, pp. 2124-2139.

McLauchlan, P. F. & Jaenicke, A., 2002. Image mosaicing using sequential bundle adjustment. *Image and Vision Computing,* Volume 20, pp. 751-759.

McLaughlin, S. A. et al., 2014. *Deep Impact 9P/Tempel Encounter - reduced HRIV images V3.0, DIF-C-HRIV-3/4-9P-ENCOUNTER-V3.0,* s.l.: NASA Planetary Data System.

Meech, K. J., Hainaut, O. R. & Marsden, B. G., 2004. comet nucleus size distributions from HST and Keck telescopes. *Icarus,* Volume 170, pp. 463-491.

Mendis, D. A. & Ip, W. H., 1977. The ionospheres and plasma tails of comets. *Space Sci. Rev.,* Volume 20, pp. 145-190.

Millman, K. J. & Aivazis, M., 2011. Python for scientists and engineers. *Computing in Science & Engineering,* Volume 13, pp. 9-12.

Mobasser, S. & Liebe, C. C., 2004. Micro sun sensor for spacecraft attitude control. *IFAC Proceedings Volumes,* Volume 37, pp. 833-838.

Moore, J. M. et al., 2016. The geology of Pluto and Charon through the eyes of New Horizons. *Science,* Volume 1293, pp. 1284-.

Moore, J. M. et al., 2016. The geology of Pluto Charon through the eyes of New Horizons. *Science,* Volume 351.

Moratto, Z. M. et al., 2010. Ames Stereo Pipeline, NASA's open source automated stereogrammetry software. *Lunar and Planetary Science Conference, 41st,* p. 2364.

Moreno, F. et al., 2016. The dust environments of comet 67P/Churyumov-Gerasimenko from Rosetta OSIRIS and VLT observations in the 4.5 to 2.9 AU heliocentric distance range inbound. *Astronomy and Astrophysics,* Volume 587.

Morgan, G. L., Liu, J. G. & Yan, H., 2010. Precise subpixel disparity measurement from very narrow baseline stereo. *IEE Transactions on Geoscience and Remote Sensing,* Volume 48, pp. 3424-3433.

Mottola, S. et al., 2015. The structure of the regolith on 67P/Churyumov-Gerasimenko from ROLIS descent imaging. *Science,* Volume 349.

Mutch, T. A. et al., 1976. The surface of Mars: the view from the Viking 1 lander. *Science,* Volume 193, pp. 791-801.

Naletto, G. et al., 2002. Optical design of the Wide Angle Camera for the Rosetta mission. *Applied Optics,* Volume 41, pp. 1446-1453.

Needham, J., 1959. *Science and Civilisation in China: Volume 3, Mathematics and the Sciences of the Heavens and the Earth.* s.l.:Cambridge University Press.

Neishtadt, A. I., Scheeres, D. J., Sidorenko, V. V. & Vasiliev, A. A., 2002. Evolution of comet nucleus rotation. *Icarus,* Volume 157, pp. 205-218.

Neubauer, F. M. et al., 1986. First results from the Giotto magnetometer experiment at comet Halley. *Nature,* Volume 321, pp. 352-355.

Neugebauer, M. et al., 2007. Encounter of the Ulysses spacecraft with the ion tail of comet McNaught. *Astrophys. J.,* Volume 667, pp. 1262-1266.

Newburn, R. L. et al., 2003. Stardust Imaging Camera. *J. Geophys. Res.,* Volume 108.

Nieke, J. et al., 2012. The Ocean and Land Colour Imager (OLCI) for the Sentinel 3 GMES mission: status and first test results. *Proc. SPIE,* Volume 8528.

Nieke, J., Solbrig, M. & Neumann, A., 1998. Signal-to-noise ratio reduction due to image smear concerning spaceborne imaging spectrometers for remote sensing of the Earth. *Proceedings of SPIE,* Volume 3439, pp. 492-502.

Nordholt, J. E. et al., 2003. Deep Space 1 encounter with Comet 19P/Borrelly: Ion composition measurements by the PEPE mass spectrometer. *Geophys. Res. Lett.,* Volume 30.

Nozette, S. et al., 1994. The Clementine mission to the Moon: scientific overview. *Science,* Volume 266, pp. 1835-1839.

Oberst, J. et al., 2001. A model for rotation and shape of asteroid 9969 Braille from ground-based observations and images obtained during the Deep Space 1 (DS1) flyby. *Icarus,* Volume 153, pp. 16-23.

Ogilvie, K. W., Coplan, M. A., Bochsler, P. & Geiss, J., 1986. Ion composition results during the International Cometary Expolorer encounter with Giacobini-Zinner. *Science,* Volume 232, pp. 374-377.

Öhman, Y., 1941. Measurements of polarization in the spectra of Comet Cunningham (1940 C) and Comet Paraskevopoulos (1941 C). *Stockholms Observatoriums Annaler,* Volume 13.

Olson, G., 2002. Image motion compensation with frame transfer CCDs. *Proc. SPIE,* Volume 4567, pp. 153-160.

Oort, J. H., 1950. The structure of the cloud of comets surrounding the Solar System and a hypothesis concerning its origin. *Bulletin of the Astronomical Institutes of the Netherlands,* Volume 11, pp. 91-110.

Öpik, E., 1932. Note on stellar perturbations of nearly parabolic orbits. *Proceedings of the American Academy of Arts and Sciences,* Volume 67, pp. 169-183.

O'Rourke, L. et al., 2020. The Philae lander reveals low-strength primitive ice inside cometary boulders. *Nature,* Volume 586, pp. 697-701.

Orton, G. S. et al., 2017. The first close-up images of Jupiter's polar regions: results from the Juno mission JunoCam instrument. *Geophys. Res. Lett.,* Volume 44, pp. 4599-4606.

Paar, G. et al., 2014. Fusion and Visualization of HiRISE Super-Resolution, Shape-from-Shading DTM with MER Stereo 3D Reconstructions. *AGU Fall Meeting,* p. Abstract #9664.

Papadimitriou, D. V. & Dennis, T. J., 1996. Epipolar line estimation and rectification for stereo image pairs. *IEEE trans. Image Process,* Volume 5, pp. 672-676.

Patat, F. & Romaniello, M., 2006. Error analysis for dual-beam optical linear polarimetry. *Publ. Astron. Soc. Pac,* Volume 118, pp. 146-161.

Pellicori, S. F., Russell, E. E. & Watts, L. A., 1973. Pioneer Imaging Photopolarimeter optical system. *Applied Optics,* Volume 12, pp. 1246-1258.

Pérez-Ayúcar, M., Nespoli, F. & Van Der Plas, P., 2018. MAPPS 3D tool: use for Rosetta science operations. *2018 SpaceOps Conference,* p. 2581.

Perminov, V. G., 1999. *The difficult road to Mars: a brief history of Mars exploration in the Soviet Union.* s.l.:NASA Headquarters History Division.

Pernechele, C. et al., 2021. Telecentric F-theta fisheye lens for space applications. *OSA Continuum,* Volume 4, pp. 783-789.

Pernechele, C. et al., 2019. Comet Interceptor's EnVisS camera sky mapping function. *Proc. SPIE Advances in Optical Astronomical Instrumentation,* Volume 11203.

Persaud, D. M., 2018. *CTX digital terrain model produced from CTX image pair B02_010486_1779_XN_02S005W / G01_018490_1779_XN_02S005W.* s.l.:Personal Communication.

Persaud, D. M., Tao, Y. & Muller, J.-P., 2019. Multi-resolution, nested orbital 3d images of Gale Crater for fused MSL rover-orbital image simulations. *EPSC-DPS meeting, EPSC-DPS2019-1540-1,* Volume 13.

Petro, N. E., 2020. The next decade of Lunar Reconnaissance Orbiter observations of the Moon: science and exporation in support of Artemis. *LPI Contributions,* Volume 2241, p. 5078.

Petrova, E. V., Jockers, K. & Kiselev, N. N., 2001. A negative branch of polarization for comets and atmosphereless celestial bodies and the light scattering by aggregate particles. *Solar System Research,* Volume 35, pp. 390-399.

Pharr, M., Humphreys, G. & Jakob, W., 2016. *Physically based rendering: from theory to implementation.* 3rd ed. s.l.:Morgan Kaufmann Publishers.

Picardi, G. et al., 2005. Radar soundings of the subsurface of Mars. *Science,* Volume 310, pp. 1925-1928.

Pilcher, E. J., 1923. Portable sundial from Gezer. *Palestine Exploration Quarterly,* Volume 55, pp. 85-89.

Pineda, J., 1988. A parallel algorithm for polygon rasterization. *Proc. 15th Annual Conf. Computer Graphics and Interactive Techniques,* pp. 17-20.

Porco, C. C. et al., 2005a. Cassini imaging science: initial results on Saturn's atmosphere. *Science,* Volume 307, pp. 1243-1247.

Porco, C. C. et al., 2005b. Cassini imaging science: initial results on Saturn's rings and small satellites. *Science,* Volume 307, pp. 1226-1236.

Porco, C. C. et al., 2005. Imaging of Titan from the Cassini spacecraft. *Nature,* Volume 434, pp. 159-168.

Porco, C. C. et al., 2003. Cassini imaging of Jupiter's atmosphere, satellites and rings. *Science,* Volume 299, pp. 1541-1547.

Prakash, R. et al., 2008. Mars Science Laboratory entry, descent, and landing system overview. *IEEE Aerospace Conference Proceedings,* pp. 1-18.

Preusker, F. et al., 2017. The global meter-level shape model of comet 67P/Churyumov-Gerasimenko. *Astronomy and Astrophysics,* Volume 601.

Price, O. et al., 2019. Fine-scale structure in cometary dust tails I: Analysis of striae in Comet C/2006 P1 (McNaught) through temporal mapping. *Icarus,* Volume 319, pp. 540-557.

Rand, J. L. & Phillips, M. L., 2002. A superpressure balloon for Mars observations. *Adv. Space Res,* Volume 30, pp. 1245-1250.

Reich, R. K. et al., 1993. Integrated electronic shutter for back-illuminated charge-coupled devices. *IEEE Trans. Electron Devices,* Volume 40, pp. 1231-1237.

Reinhard, R., 1986. The Giotto encounter with comet Halley. *Nature,* Volume 321, pp. 313-318.

Remetean, E. et al., 2016. Philae locating and science support by robotic vision techniques. *Acta Astronautica,* Volume 125, pp. 161-173.

Richard Gott III, J. et al., 2005. A map of the Universe. *Astrophys. J,* Volume 624.

Richardson, J. E. & Melosh, H. J., 2013. An examination of the Deep Impact collision site on Comet Tempel 1 via Stardust-NExT: Placing further constraints on cometary surface properties. *Icarus,* Volume 222, pp. 492-501.

Richardson, J., Lorenz, R. D. & McEwen, A., 2004. Titan's surface and rotation: new results from Voyager 1 images. *Icarus,* Volume 170, pp. 113-124.

Rich, V., 1996. Space: Mars mission failure dents Russia's pride. *Physics World,* Volume 9, p. 8.

Rizk, B. et al., 2018. OCAMS: The OSIRIS-REx Camera Suite. *Space Sci. Rev.,* Volume 214.

Roatsch, T. et al., 2008. High-resolution Enceladus atlas derived from Cassini-ISS images. *Plan. and Space Sci,* Volume 56, pp. 109-116.

Robbins, J. S. & Hynek, B. M., 2012. A new global database of Mars impact craters ≥ 1 km: 1. Database creation, properties, and parameters. *J. Geophys. Res,* Volume 117.

Robinson, M. S. et al., 2010. Lunar Reconnaissance Orbiter Camera (LROC) Instrument Overview. *Space Sci. Rev.,* Volume 150, pp. 81-124.

Robitaille, T. P. et al., 2013. Astropy: a community Python package for astronomy. *A&A,* Volume 558.

Rochus, P. et al., 2020. The Solar Orbiter EUI instument: the Extreme Ultraviolet Imager. *Astronomy & Astrophysics,* Volume 642.

Rodgers, A. W. et al., 1980. SEC Vidicon photometry of the main sequence of omega centauri. *Publ. Astron. Soc. Pac.,* Volume 92, pp. 288-299.

Rossmann, K., 1969. Point spread-function, line spread-function, and modulation transfer-function: tools for the study of imaging systems. *Radiology,* Volume 93, pp. 257-272.

Roth, S. D., 1982. Ray casting for modeling solids. *Computer Graphics and Image Processing,* Volume 18, pp. 109-144.

Saad, E. & Hirakawa, K., 2020. Improved photometric acceptance testing in image feature extraction tasks. *J. Electron. Imag.,* Volume 29.

Sagan, C. & Druyan, A., 1997. *Comet.* 2nd ed. London: Headline.

Sagdeev, R. Z., Linkin, V. M., Blamont, J. E. & Preston, R. A., 1986b. The VEGA Venus balloon experiment. *Science,* Volume 231, pp. 1407-1408.

Sagdeev, R. Z. et al., 1986a. Television observations of comet Halley from Vega spacecraft. *Nature,* Volume 321, pp. 262-266.

Samarasinha, N. H., Mueller, B. E. A. & Belton, M. J. S., 1999. Coma morphology and constraints on the rotation of Comet Hale-Bopp (C/1995 O1). *Earth, Moon and Planets,* Volume 77, pp. 189-198.

Sánchez-Lavega, A. et al., 2018. The rich dynamics of Jupiter's Great Red Spot from JunoCam: Juno images. *AJ,* Volume 156.

Sarraga, R. F., 1983. Algebraic methods for intersections of quadric surfaces in GMSOLID. *Computer Vision, Graphics, and Image Processing,* Volume 22, pp. 222-238.

Schneider, D., Schwalbe, E. & Maas, H. -G., 2009. Validation of geometric models for fisheye lenses. *ISPRS Journal of Photogrammetry and Remote Sensing,* Volume 64, pp. 259-266.

Scholten, F. et al., 2012. GLD100: The near-global lunar 100 m raster DTM from LROC WAC stereo image data. *JGR,* Volume 117.

Scholten, F. et al., 2019. The Hayabusa2 lander MASCOT on the surface of asteroid (162173) Ryugu–Stereo-photogrammetric analysis of MASCam image data. *Astronomy & Astrophysics,* Volume 632.

Schultz, P. H. et al., 2007. The Deep Impact oblique impact cratering experiment. *Icarus,* Volume 190, pp. 295-333.

Schurmeier, H. M., Heacock, R. L. & Wolfe, A. E., 1966. The Ranger missions to the Moon. *Scientific American,* Volume 214, pp. 52-69.

Seiff, A., 1993. Mars atmospheric winds indicated by motion of the Viking landers during parachute descent. *J. Geophys. Res.,* Volume 98, pp. 7461-7474.

Sekanina, Z., 1996. Morphology of cometary dust coma and tail. *International Astronomical Union Colloquium,* Volume 150, pp. 377-382.

Sekanina, Z. et al., 2004. Modeling the Nucleus and Jets of Comet 81P/Wild 2 based on the Stardust encounter data. *Science,* Volume 304, pp. 1769-1774.

Sekanina, Z. & Larson, S. M., 1986. Dust jets in comet Halley observed by Giotto and form the ground. *Nature,* Volume 321, pp. 357-361.

Shepard, M. K. & Helfenstein, P., 2007. A test of the Hapke photometric model. *J. Geophys. Res.,* Volume 112.

Shiraishi, H., Tanaka, S., Fujimara, A. & Hayakawa, H., 2008. The present status of the Japanese penetrator mission: LUNAR-A. *Advances in Space Research,* Volume 42, pp. 386-393.

Shorthill, R. W. et al., 1976. The environs of Viking 2 lander. *Science,* Volume 194, pp. 1309-1318.

Shortis, M. R., Seager, J. W., Harvey, E. S. & Robson, S., 2005. The influence of Bayer filters on the quality of photogrammetric measurement. *Proceedings,* pp. 164-171.

Siddiqi, A. A., 2018. Beyond Earth: a chronicle of deep space exploration, 1958-2016. *National Aeronautics and Space Administration, Office of Communications, NASA History Division.*

Skinner, C. J. et al., 1998. On-orbit properties of the NICMOS detectors on HST. *Infrared Astronomical Instrumentation,* Volume 3354, pp. 2-13.

Smith, B. A. et al., 1977. Voyager imaging experiment. *Space Sci. Rev.,* Volume 21, pp. 103-127.

Smith, B. A. et al., 1979a. The Galilean satellites and Jupiter: Voyager 2 imaging science results. *Science,* Volume 206, pp. 927-950.

Smith, B. A. et al., 1979b. The Jupiter sysem through the eyes of Voyager 1. *Science,* Volume 204, pp. 951-972.

Smith, B. A. et al., 1982. A new look at the Saturn system: the Voyager 2 images. *Science,* Volume 215, pp. 504-537.

Smith, B. A. et al., 1981. Encounter with Saturn: Voyager 1 imaging science results. *Science,* Volume 212, pp. 163-191.

Smith, D. E. et al., 2001. Mars Orbiter Laser Altimeter: Experiment summary after the first year of global mapping of Mars. *J. Geophys. Res,* Volume 106, pp. 23689-23722.

Smith, G. E., 2009. The invention and early history of the CCD. *Nucl. Instrum. Methods Phys. Res. A,* Volume 607, pp. 1-6.

Smith, R. W. & Tatarewickz, J. N., 1985. Replacing a technology: the Large Space Telescope and CCDs. *Proc. IEEE,* Volume 73, pp. 1221-1235.

Smith, W. J., 2000. *Modern Optical Engineering.* 3rd ed. s.l.:McGraw-Hill.

Smits, B., 1998. Efficiency issues for ray tracing. *Journal of Graphics Tools,* Volume 3, pp. 1-14.

Smrekar, S. et al., 1999. Deep Space 2: The Mars microprobe mission. *Geophysical Research,* Volume 104, pp. 27013-27030.

Snodgrass, C. & Jones, G. H., 2019. The European Space Agency's Comet Interceptor lies in wait. *Nature Communications,* Volume 10, p. 5418.

Snyder, J. M., 1987. Ray tracing comples models containing surface tessellations. *Computer Graphics,* Volume 21, pp. 119-128.

Soderblom, L. A. et al., 2002. Observations of Comet 19P/Borelly by the Miniature Integrated Camera and Spectrometer aboard Deep Space 1. *Science,* Volume 296, pp. 1087-1091.

Soderblom, L. A. et al., 2007. Topography and geomorphology of the Huygens landing site on Titan. *Plan. and Space Sci,* Volume 55, pp. 2015-2024.

Soman, M. et al., 2014. Design and characterisation of the new CIS115 sensor for JANUS, the high resolution camera on JUICE. *Proc. of SPIE,* Volume 9154.

Sparks, W. B. & Axon, D. J., 1999. Panoramic Polarimetry Data Analysis. *Publ. Astron. Soc. Pac,* Volume 111, pp. 1298-1315.

Squyres, S. W. et al., 2006. Overview of the Opportunity Mars Exploration Rover mission to Meridiani Planum: Eagle Crater to Purgatory Ripple. *J. Geophys. Res.,* Volume 111.

Squyres, S. W. et al., 2009. Exporation of Victoria crater by the Mars Rover Opportunity. *Science,* Volume 324, pp. 1058-1061.

Steinbrugge, G. et al., 2020. The surface rougness of Europa derived from Galileo stereo images. *Icarus,* Volume 343.

Stelzried, C., Efron, L. & Ellis, J., 1986. Halley Comet Missions. *TDA Progress Report,* Volume 42-87, pp. 240-242.

Stern, A., 2018. NEW HORIZONS calibrated MVIC PLUTO ENCOUNTER V3.0, NH-P-MVIC-3-PLUTO-V3.0. *NASA Planetary Data System.*

Stern, S. A., 2003. The evolution of comets in the Oort cloud and Kuiper belt. *Nature,* Volume 424, pp. 639-642.

Stern, S. A. et al., 2015. The Pluto system: Initial results from its exploration by New Horizons. *Science,* Volume 350.

Stroebel, L. & Zakia, R. D., 1993. *The focal encyclopedia of photography.* 3rd ed. s.l.:Focal Press.

Strojnik, M. & Paez, G., 1997. Push-broom reconnaissance camera with time expansion for a (Martian) landing-site certification. *Optical Engineering,* Volume 36.

Sunshine, J. M. et al., 2006. Exposed water ice deposits on the surface of Comet 9P/Tempel 1. *Science,* Volume 311.

Surkov, Y. A. & Kremnev, R. S., 1998. Mars-96 mission: Mars exporation with the use of penetrators. *Planet. Space Sci.,* Volume 46, pp. 1689-1696.

Sutton, S. S. et al., 2017. *Correcting spacecraft jitter in HiRISE images.* s.l.:s.n.

Tabataba-Vakili, F. et al., 2020. Long-term tracking of circumpolar cyclones on Jupiter from polar observations with JunoCam. *Icarus,* Volume 335, p. 113405.

Tao, Y. & Muller, J.-P., 2016. A novel method for surface exploration: Super-resolution restoration of Mars repeat-pass orbital imagery. *Planet. Space Sci.,* Volume 121, pp. 103-114.

Tao, Y., Muller, J.-P. & Poole, W., 2016. Automated localisation of Mars rovers using co-registered HiRISE-CTX-HRSC orthorectified images and wide baseline Navcam orthorectified mosaics. *Icarus,* Volume 280, pp. 139-157.

Tao, Y. et al., 2018. Massive stereo-based DTM production for Mars on cloud computers. *Plan. Space Sci,* Volume 154, pp. 30-58.

Taylor, M. G. G. T., Altobelli, N., Buratti, B. J. & Choukroun, M., 2017. The Rosetta mission orbiter science overview: the comet phase. *Philos. Trans. R. Soc.,* Volume 375.

Thomas, N. et al., 2015. Redistribution of particles across the nucleus of comet 67P/Churyumov-Gerasimenko. *A&A,* Volume 583.

Thomas, N. & Keller, U., 1990. Photometric calibration of the Halley Multicolour Camera. *Applied Optics,* Volume 29, pp. 1503-1519.

Thomas, P. C. et al., 1999. Mathilde: size, shape, and geology. *Icarus,* Volume 140, pp. 17-27.

Thomas, P. C. et al., 2007. The shape, topography, and geology of Tempel 1 from Deep Impact observations. *Icarus,* Volume 187, pp. 4-15.

Thomas, P. C. et al., 1994. The shape of Gaspra. *Icarus,* Volume 107, pp. 23-36.

Thompson, W. T., 2015. Linear polarization measurements of Comet C/2011 W3 (Lovejoy) from STEREO. *Icarus,* Volume 261, pp. 122-132.

Thompson, W. T. & Reginald, N. L., 2008. The Radiometric and Pointing Calibration of SECCHI COR1 on STEREO. *Solar Phys,* Volume 250, pp. 443-454.

Thom, R. D., Koch, T. L., Langan, J. D. & Parrish, W. J., 1980. A fully monolithic InSb infared CCD array. *IEEE Trans. Electron Devices,* Volume 27, pp. 160-170.

Thorpe, T. E., 1973. Mriner 9 photometric observations of Mars from november 1971 through march 1972. *Icarus,* Volume 20, pp. 482-489.

Tomasko, M. G. et al., 1999. The Descent Imager/Spectral Radiometer (DISR) experiment on the Huygens entry probe of Titan. *Space Sci. Rev.,* Volume 104, pp. 469-551.

Trilling, D. E. et al., 2018. Spitzer observations of interstellar object 1I/'Oumuamua. *Astron. J.,* Volume 156.

Tsou, P. et al., 2004. Stardust encounters comet 81P/Wild 2. *J. Geophys. Res.,* Volume 109.

Tubiana, C. & ESA, 2017. *Determination of the absolute calibration coefficients to radiometrically calibrate OSIRIS images.* [Online] Available at: ftp://psa.esac.esa.int/pub/mirror/INTERNATIONAL-ROSETTA-MISSION/OSIWAC/RO-C-OSIWAC-3-PRL-67PCHURYUMOV-M05-V2.0/DOCUMENT/CALIB/RADIOMETRIC_CALIB_V01.PDF
[Accessed 09 2019].

Tubiana, C. et al., 2015. Scientific assesment of the quality of OSIRIS images. *Astron. Astrophys.,* Volume 583.

Tubiana, C. et al., 2015. 67P/Churyumov-Gerasimenko: activity between March and June 2014 as observed from Rosetta/OSIRIS. *Astronomy and Astrophysics,* Volume 573.

Tyo, J. S., Goldstein, D. L., Chenault, D. B. & Shaw, J. A., 2006. Review of passive imaging polarimetry for remote sensing applications. *Applied Optics,* Volume 45, pp. 5453-5469.

Vasavada, A. R. et al., 2006. Cassini imaging of Saturn: southern hemisphere winds and vortices. *JGR Plan.,* Volume 111.

Vernimmen, R., Hooijer, A. & Pronk, M., 2020. New ICESat-2 satellite LiDAR data allow first global lowland DTM suitable for accurate coastal flood risk assessment. *Remote Sens,* Volume 12.

Veverka, J., Belton, M., Klaasen, K. & Chapman, C., 1994. Galileo's encounter with 951 Gaspra: overview. *Icarus,* Volume 107, pp. 2-17.

Veverka, J. et al., 2013. Return to Comet Tempel 1: Overview of Stardust-NExT results. *Icarus,* Volume 222, pp. 424-435.

Vincent, J. B. et al., 2016. Are fractured cliffs the source of cometary dust jets? Insights from OSIRIS/Rosetta at 67P/Churyumov-Gerasimenko. *A&A,* Volume 587.

Virtanen, P. et al., 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods,* Volume 17, pp. 261-272.

Volk, K. & Malhotra, R., 2008. The scattered disk as the source of the Jupiter Family Comets. *Astrophys. J.,* Volume 687, pp. 714-725.

Vondrak, R., Keller, J., Chin, G. & Garvin, J., 2010. Lunar Reconnaissance Orbiter (LRO): observations for Lunar exploration and science. *Space Sci. Rev.,* Volume 150, pp. 7-22.

Walker, S., 2007. New features in SOCET SET. *Photogrammetric Week,* Volume 7.

Watanabe, S. et al., 2019. Hayabusa2 arrives at the cabonaceous asteroid 162173 Ryugu - a spinning top-shaped rubble pile. *Science,* Volume 364, pp. 268-272.

Weaver, H. A. et al., 1995. The Hubble Space Telescope (HST) observing campaign on comet Shoemaker-Levy 9. *Science,* Volume 267, pp. 1282-1288.

Weaver, H. A. et al., 2016. The small satellites of Pluto as observed by New Horizons. *Science,* Volume 351.

Weissman, P. R., 1990. The Oort cloud. *Nature,* Volume 344, pp. 825-830.

Wellman, J. B., Landauer, F. P., Norris, D. D. & Thorpe, T. E., 1976. The Viking Orbiter Visual Imaging Subsystem. *Journal of Spacecraft and Rockets,* Volume 13, pp. 660-666.

Whipple, F. L., 1950. A comet model. I. The acceleration of Comet Encke. *The Astrophysical Journal,* Volume 111, pp. 375-394.

Whitted, T., 1980. An improved illumination model for shaded display. *Communications of the ACM 23,* Volume 6, pp. 343-349.

Widemann, T., Ghail, R., Wilson, C. F. & Titov, D. V., 2020. EnVision: Europe's proposed mission to Venus. *AGU Fall Meeting abstract,* pp. P022-02.

Widenhorn, R., Dunlap, J. C. & Bodegom, E., 2010. Exposure time dependence of dark current in CCD imagers. *IEEE Trans. Electron Devices,* Volume 57, pp. 581-587.

Willner, K., Shi, X. & Oberst, J., 2014. Phobos' shape and topography models. *Plan. and Space Sci,* Volume 102, pp. 51-59.

Wilson, J. K., Baumgardner, J. & Mendillo, M., 1998. Three tails of comet Hale-Bopp. *Geophys. Res. Lett.,* Volume 25, pp. 225-228.

Wong, H., Yao, Y. L. & Schlig, E. S., 1992. TDI charge-coupled devices: Design and applications. *IBM Journal of Researh and Development,* Volume 36, pp. 83-106.

Woodney, L. M. et al., 2002. Morphology of HCN and CN in comet Hale-Bopp (1995 O1). *Icarus,* Volume 157, pp. 193-204.

Wray, J., 2013. Gale crater: The Mars Science Laboratory/Curiosity rover landing site. *International Journal of Astrobiology,* Volume 12, pp. 25-38.

Wu, W. et al., 2019. Lunar farside to be explored by Chang'e-4. *Nature Geoscience,* Volume 12, pp. 222-223.

Xiao, Z. & Werner, S. C., 2015. Size-frequency distribution of crater populations in equilibrium on the Moon. *J. Geophys. Res. Planets,* Volume 120, pp. 2277-2292.

Yeomans, D. K. et al., 2004. Cometary orbit determination and nongravitational forces. In: M. C. Festou, H. U. Keller & H. A. Weaver Jr., eds. *Comets II.* Tucson: University of Arizona Press, pp. 137-151.

Ye, Q. et al., 2020. Pre-discovery activity of new interstellar comet 2I/Borisov beyond 5 AU. *Astron. J.,* Volume 159.

Young, D. T. et al., 2004. Solar wind interactions with Comet 19P/Borrelly. *Icarus,* Volume 167, pp. 80-88.

Young, M., 1989. The pinhole camera: imaging without lenses or mirrors. *The Physics Teacher,* Volume 27, pp. 648-655.

Young, R. E., 2003. The Galileo probe: how it has changed our understanding of Jupiter. *New Astronomy Reviews,* Volume 47, pp. 1-51.

Zakharov, A. I., Prokhorov, M. E., Tuchin, M. S. & Zhukov, A. O., 2013. Minimum star tracker specifications required to achieve a given attitude accuracy. *Astrophys. Bull.,* Volume 68, pp. 481-493.

Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.,* Volume 22, pp. 1330-1334.

Zhong, X. & Jia, J. Q., 2009. Stary light removing design and simulation of spaceborne camera. *Opt. Precision Eng,* Volume 17, pp. 621-625.

Appendix: Image Simulation Code

The raw code developed and used throughout this thesis for simulation of planetary camera images is available from University College London's Research Data Depository, accessible at https://doi.org/10.5522/04/16943998.v1.

As discussed in section 8.3, developing the code into a well-documented, distributable and user-friendly library or piece of software was beyond the scope of this thesis. The code has therefore been made available in its raw form, with no accompanying documentation, and very limited commenting.