Planetary Topography Measurement by Descent Stereophotogrammetry

G. Brydon[a,b], D.M. Persaud[a,b], G.H. Jones[a,b]

[a]Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey RH5 6NT, UK

[b]The Centre for Planetary Sciences at UCL/Birkbeck, London, UK

Corresponding author:

G. Brydon

E-mail: george.brydon.15@ucl.ac.uk

Postal: Mullard Space Science Laboratory

Holmbury St. Mary, Dorking

Surrey

RH5 6NT

UK

Abstract

Digital Terrain Models (DTMs) provide valuable insights into the nature of solar system surfaces, facilitating geological analysis, landing site selection and characterisation, and contextualising in situ measurements. For missions or solar system bodies for which orbiters and soft landed platforms are technologically or financially challenging to achieve, low mass descent or ascent probes (e.g. planetary penetrators) provide an alternative means by which to access the atmosphere and/or surface, and a platform from which to image the surface from a range of altitudes and perspectives. This paper presents a study into the concept of large-coverage descent stereophotogrammetry, whereby the stereo geometry of vertically offset wide-angle descent images is used to measure surface topography over a region of large extent. To do this, we simulate images of Mars' Gale Crater using a large coverage, high resolution DTM of the area, and derive topographic measurements by stereo matching pairs of simulated images. These topographic measurements are compared directly with the original DTM to characterise their accuracy, and dependence of elevation measurement accuracy on stereo geometry is thus investigated. For a stereo pair with a given altitude (corresponding to the altitude of its lower image), error in elevation measurement is found to have its minimum value for surface at a horizontal distance between 1 and 3 times the altitude. For a point on the surface with given horizontal distance from the imaging location, a stereo imaging altitude between 0.2 and 0.5 times this distance is found to achieve best elevation measurement accuracy. Surface appearance, and its change between two images of a stereo pair, is found to have a significant impact on stereo matching performance, limiting stereo baseline length to an optimum value range of 0.2-0.4 times the lower image's altitude, and resulting in the occurrence of occlusions and blind spots, particularly at oblique viewing angles.

1 Introduction

Images of planetary surfaces reveal many properties, and are a valuable tool for their investigation. But a single image gives a flattened representation of the surface, and the possible analysis can be elevated if knowledge of the surface topography is also available. Digital Terrain Models (DTMs), a common end product of topographic mapping, are an often utilised and valuable product of spacecraft imagers. They store information on the 3D structure of a surface in a form that is both intuitive to inspect by eye and easy to represent in 3D visualisation software, and can be utilised for rigorous geological analysis and landing site characterisation. They are commonly produced from orbital imagery, allowing large surface coverage datasets (e.g. Tao et al. (2018)). Meanwhile, high resolution DTMs with small spatial extent have been derived from rover images (e.g. Barnes et al. (2018)). A crucial step in the production of a DTM is the triangulation of three-dimensional (3D) positions on the surface.

For missions or solar system bodies for which orbiters and soft landed platforms are technologically or financially challenging to achieve, such as the moons of Jupiter and Saturn, low mass descent platforms (e.g. penetrators/impactors) provide an alternative means by which to access the atmosphere and/or surface, and a platform from which to image (Gowen, et al., 2011; Lorenz, 2011). Topographic measurement of the probe's landing site and surrounding area can provide valuable geological context to any in situ measurements, reveal local surface processes, and constrain landing site location (Liu, et al., 2019). This paper focusses on the technique of utilising descent images to measure surface topography.

Descent imaging itself has been employed on a range of missions to solar system bodies. The majority of these missions went either to Mars or the Moon, and captured series of descent images containing tens to hundreds of images. Malin et al. (2001) describe the motivation for descent imaging, which derives largely from its ability to determine the location and geological characteristics of the landing or impact site. The nested, multi-scale nature of descent image sequences allows observations at a range of scales and accuracies to be tied together.

NASA's Mars Exploration Rovers (Crisp, et al., 2003) were each equipped with a lander-mounted descent camera, which viewed directly downward with a square 45° field of view (FOV). Images were captured during descent with the prime objective of assisting automated on-board estimation of the crafts' horizontal velocities (Maki, et al., 2003). The Mars Descent Imager (MARDI) of the Mars Science

Laboratory followed the same theme, looking directly downward with its 90° FOV and capturing sequential images (several hundred) from the lander's heat shield release to its final touchdown, with pixel scales ranging from 1.5 m to 1.5 mm (Grotzinger, et al., 2012; Malin, et al., 2017). The primary purpose of its images was to determine the rover's landing site location and characteristics.

Whilst not the objective for the above mentioned cameras, topographic mapping has been a key outcome of several descent imaging campaigns. Both the Chang'e 3 (Li, et al., 2015) and Chang'e 4 (Jia, et al., 2018) missions placed rovers on the Moon's surface, and acquired descent images during the landing phases. Liu et al. (2015) applied feature matching and bundle adjustment to 180 Chang'e 3 descent images to produce sub-metre precision landing site DTMs extending as far as 1800 m from the landing site. Liu et al. (2019) used a similar approach with Chang'e 4 descent images, producing a landing site DTM and precisely constraining the landing site location.

Huygens descent probe images (Karkoschka, et al., 2007) contrasted to those of the Martian and Lunar missions' downward looking cameras, in that their larger FOVs extended toward the horizon. Soderblom et al. (2007) performed stereophotogrammetry of multiple Huygens descent image pairs to derive topographic maps of two regions of Titan's surface, and the descent images were recently revisited to produce a new DTM of higher spatial accuracy (Daudon, et al., 2020). Descent imaging has additionally been used to derive the topography of small bodies. For example, Mottola et al. (2015) performed stereophotogrammetry on two images from the Rosetta mission's Philae lander to produce a DTM of a small region of comet 67P/Churyumov-Gerasimenko's surface.

In spite of descent imaging being a common and important feature of planetary landers, and the applicability of its images to topography measurement having been demonstrated by several missions, a comprehensive study of the general technique of descent stereophotogrammetry has not previously been conducted. This paper aims to address this by providing a general quantitative assessment of the technique's achievable accuracies and coverage. Additionally, whilst most descent images are focussed on a limited region below the spacecraft, we address the concept of large-coverage descent stereophotogrammetry, whereby the stereo geometry of vertically offset wide-angle descent images is used to measure surface topography over a region of large extent.

To achieve large and multi-directional coverage, we envisage that each image's extent covers a full hemisphere, centred on the camera's nadir, such that it captures the entire visible surface (described

further in section 2.2). The large field of view captures the surface with a range of ground sample distances (GSDs) and emission angles. This variation in imaging geometry will impact the efficacy of stereoscopic analysis, and this study focusses on characterising that dependency.

2 Method

The accuracy of surface topographic measurement is a function of many factors, and cannot be fully assessed without knowledge of the surface it represents. Often, 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 topography measurements 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 their accuracy (Li, 1991).

When the objective is not to map a specific surface, but to investigate a method of topography measurement, simulated images of a pre-existing surface model (such as a DTM) can be used. By utilising a well constrained surface model, images can be simulated and then themselves be used to measure the surface's topography. The measured topography can subsequently be directly compared with the original surface model 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. The model should however contain features that are representative of the surfaces on which the topographic measurement method is intended to be used (Kirk, et al., 2016).

Whilst simulated images of a surface model are convenient for performing direct quantitative comparisons of the derived and ground truth data, such an experimental setup introduces some simplifications versus the use of real data. When a surface's topography is mapped, differences between the resulting surface model and the real surface it represents arise from the presence of errors, resolution limits and noise in the imaging and stereophotogrammetry processes employed. In any simulation of imagers, their data and their use in topographic mapping, it's important that these effects be accurately represented.

An additional 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 replicate this in simulated images, it is important that the ground truth surface model have a higher spatial density than can be resolved by the simulated imaging system. The computational cost of this for studying high resolution cameras

or large-coverage surface models is significant, and has historically limited the usefulness of this technique.

Whilst simulated images are ideal for an early or general study (such as is presented in this manuscript), the development and design of an actual imager intended for a specific mission payload should also include the investigation of its performance capturing actual images of real scenes.

This study utilises simulated images of a high resolution surface model in order to derive topographic point clouds from vertically separated stereo pairs. The accuracy of these topographic point clouds is assessed and used to investigate the dependence of measured elevation accuracy on stereo viewing geometry.

2.1   Ground Truth Data

A set of images obtained throughout a vertical descent captures a surface in a wide range of perspectives and scales. At high altitudes the camera's FOV encompasses a very large area on the surface, whilst low altitude images achieve high spatial resolution. The nature of the wide-angle imaging investigated in this study exacerbates this, and leads to large images with significantly varying spatial resolution. To accommodate this, and facilitate a useful simulated sequence of descent images, this study required a ground truth of both high spatial resolution and large extent.

In this study we therefore utilise a DTM mosaic covering Gale Crater, Mars, the landing site of the Mars Science Laboratory (Wray, 2013). This dataset comprises Mars Reconnaissance Orbiter Context Camera (CTX) (Malin, et al., 2007) and High-Resolution Stereo Camera (HRSC) (Jaumann, et al., 2007) 3D imagery. The CTX DTMs were processed using the CASP-GO processing suite (Tao, et al., 2018), with 3% interpolation from HRSC DTMs processed using the KM09-VICAR pipeline (Kim & Muller, 2009), as described in Persaud et al. (2019).

This dataset was selected due to its large surface coverage (spanning a region of approximately 170 x 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. In addition to this, an accompanying 6 m per pixel greyscale CTX orthorectified image (ORI) of the same extent, draped over the DTM, provides surface detail. The DTM and ORI which together constitute the ground truth dataset are shown in Figure 1
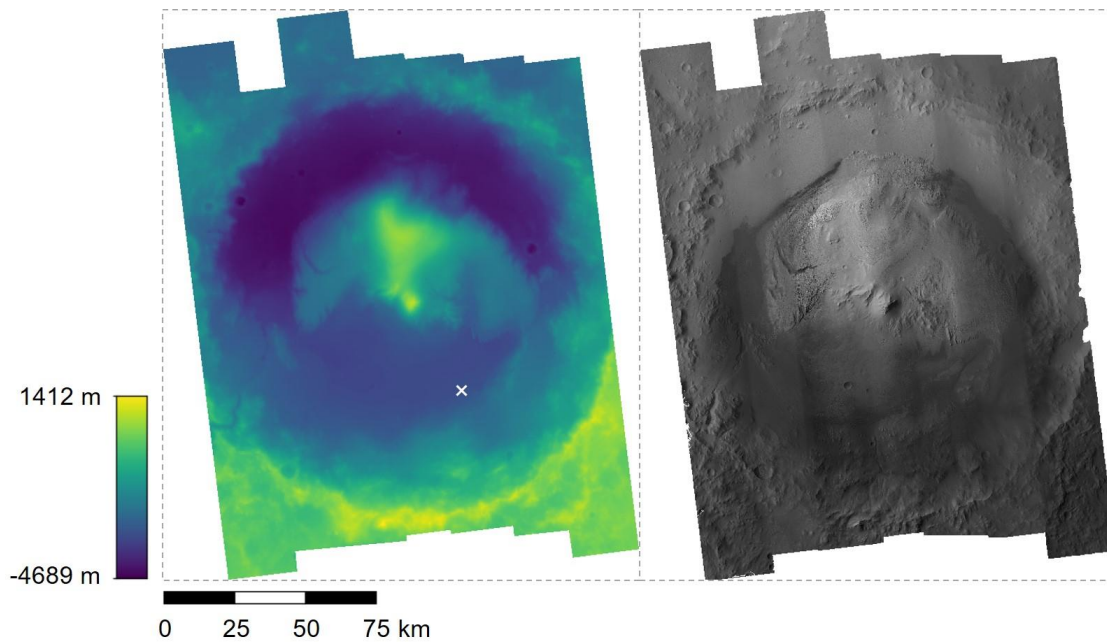
Figure 1: The ground truth data used for simulating this study's descent images, from Persaud et al. (2019). (left) DTM of Gale Crater with 18 m spatial resolution and a vertical extent of -4689 m (dark purple) to 1412 m (light yellow). (right) 6 m per pixel ORI of Gale Crater. Both images span 170 km 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 the Gale Crater data set and extend the area of simulated surface. 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.

2.2   Image Concept

The fundamental imaging concept on which this study is based is the acquisition of wide-angle images with coverage of the entire visible surface from the probe's perspective. These images cover a full hemisphere, imaging from the camera's nadir up to horizontal, over a full 360° of azimuth. An example is shown in Figure 2.
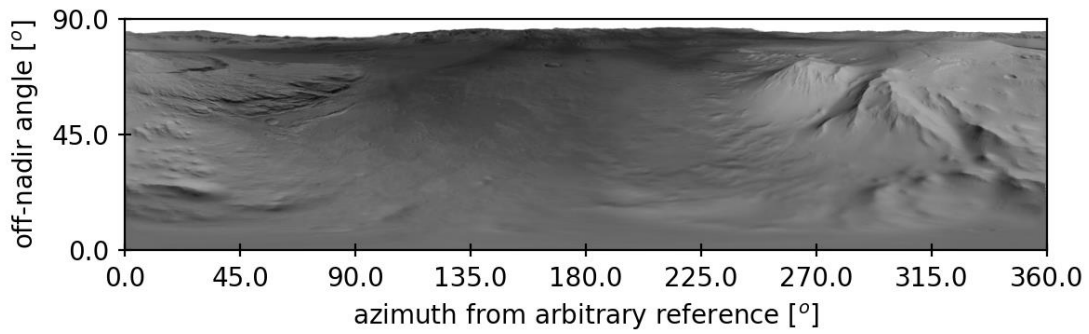
Figure 2: An example simulated hemispherical image of the Gale crater dataset.

We envisage three main ways by which images of this nature can be captured:

- Rotational push-broom or push-fame scanning with a ≥90° vertical FOV, mounted on a spinning probe such as a penetrator (Brydon & Jones, 2018; Schneider & Schwalbe, 2005).

- Stitching of few overlapping frame images, from a small number of individual frame cameras or a single frame camera aboard a slowly rotating platform (e.g. Tomasko et al. (2002)).

- A single downward looking camera with fisheye optics and a full hemispherical FOV.

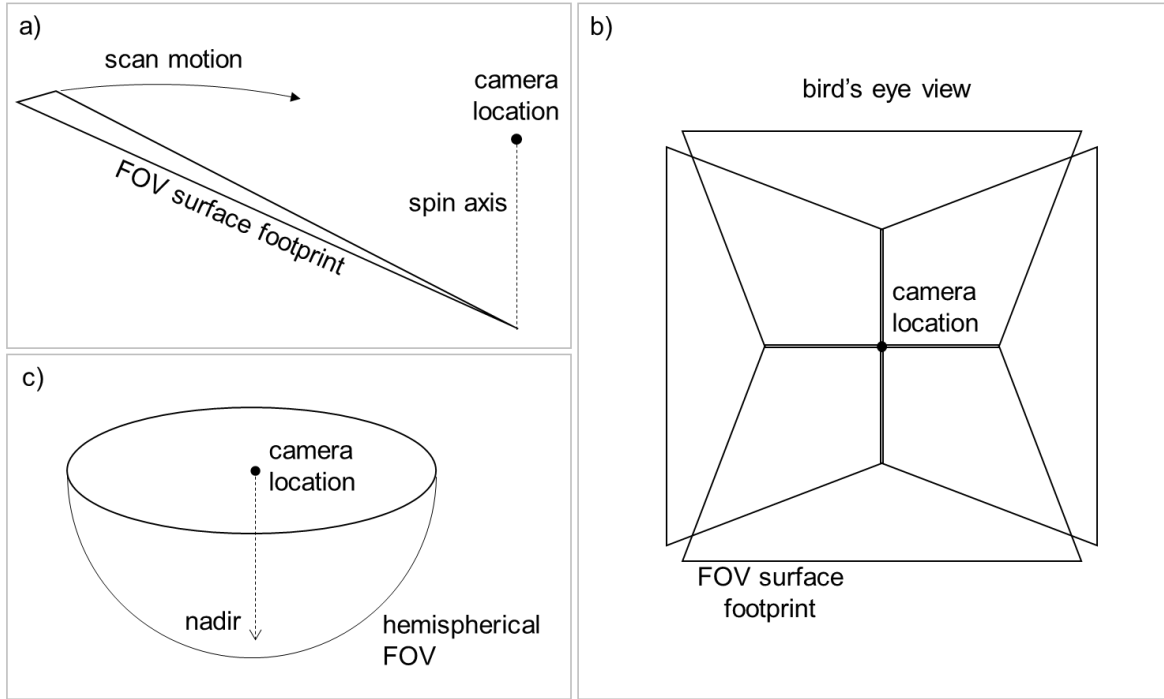These methods of surface imaging are illustrated in Figure 3.

Figure 3: Three broad modes of imaging by which images of the type shown in Figure 2 could be obtained. (a) rotational push-broom or push-frame camera with narrow 90° vertical FOV on a spinning probe. (b) overlapping frame images, acquired either from multiple cameras, or a single rotating frame camera. (c) single downward looking camera with full hemispherical fisheye FOV.

As the probe descends (or ascends), a series of nested images is captured with spatial resolution scaling as altitude changes. A sequence of images over the course of the descent provides a gradual and continuous transition from orbital to landed viewing perspective, and provides a stereo geometry (with vertical baseline) by which to triangulate surface points (Brydon & Jones, 2018).

For the sake of this study, images (of the type shown in Figure 2) were simulated to have a uniform vertical angular resolution, such as would be obtained by equidistant fisheye optics (Schneider, et al., 2009). The horizontal angular resolution of the images is also taken to be uniform, and equal to the vertical. The angle above nadir, a of a given pixel's line of sight (LOS) is therefore given by

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

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.

A camera with optics of a different geometry (e.g. perspective projection) would produce an image with a different relationship between a and N, in particular having a lower angle per pixel at the extremes of its FOV than at the centre. However, images captured by different optics could still be projected to the azimuth versus angle-from-nadir coordinates consistent with Figure 2 and equation 1, and treated in the same way as those used in this study. The geometry and accuracy of scanned and stitched images (i.e. cases (a) and (b) in Figure 3) would be dependent on the stability of the imaging platform, and may require more involved processing (e.g. bundle adjustment) to constrain.

2.3   Image Simulation

All images used in this study were simulated with our own bespoke software. It is built to be flexible to the mode of imaging, allowing control of the camera parameters, the imaging mechanism (e.g. scanning or frame) and the descent trajectory. This study focusses on 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 not 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.

All images were simulated as being captured during a single vertical descent, during which the camera orientation remains constant, whilst its altitude varies. For a real mission, variations in trajectory and pointing could, to some degree, 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)).

The camera was simulated to have a coarse horizontal and vertical resolution of 0.18° per pixel (corresponding to 500 pixels spanning the full vertical FOV of Figure 2) in order to keep its pixels' footprints larger than the ground truth's 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.

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 in order to investigate the effect of both imaging altitude and stereo baseline on the achievable elevation measurement accuracies. For any stereo pair, we choose to define its imaging altitude as that of the lower image. The height and baseline-to-height ratio of the 33 simulated stereo pairs used in this study are plotted in Figure 4.
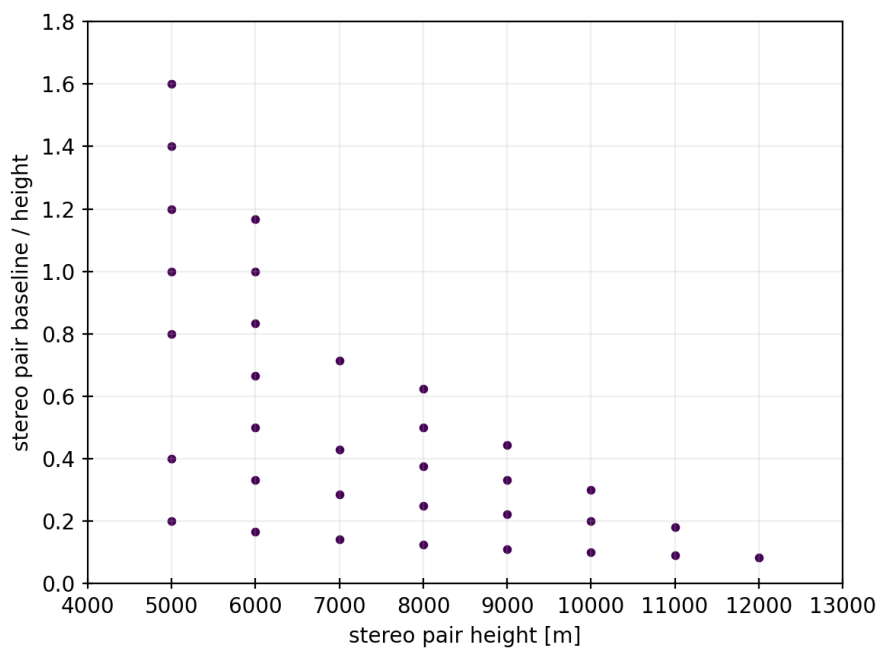

Figure 4: The heights and baselines of the 33 simulated stereo pairs used in this study.

2.4   Topography Measurement

Surface points were triangulated by applying stereo matching to two images of different altitude in order to produce a single surface point cloud. The simulated images are inherently rectified by their common azimuth coordinates (horizontal axis), and projection to common azimuth coordinates could be used as

the mechanism of rectification for images from a real descent. 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 of this study's simulated images. However, semi-global matching does not account for significant changes in the scale and shape of image features, both of which will occur in wide-angle descent images due to the surface's changing proximity to the camera and the distortion necessary for a large FOV (because our images are rectified by projecting them to common azimuth coordinates, and azimuth angles are independent of image altitude, the scale and shape changes occur only along the vertical axes of the images). We chose not to attempt to account for these scale and shape changes in the implementation of the stereo matching algorithm, and instead use an error model (see section 2.6) to identify the effect and presence of stereo matching errors within our simulated topography measurements.

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), resulting in a disparity map. For a stereo pair of images, the LOS angle above nadir a_1 of a pixel in the lower altitude image is given by equation 1, whilst the LOS angle above nadir a_2 of the matched pixel in the higher altitude image is given by

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

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.
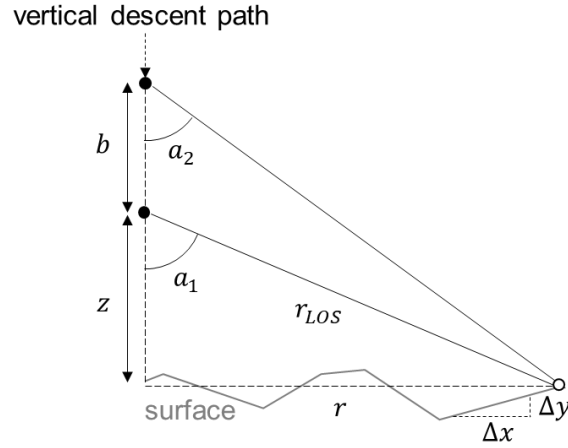
Figure 5: Stereo viewing geometry by which the physical positions of points on the surface are

triangulated.

The point's vertical displacement from the lower imaging location (whose altitude defines the stereo

pair's altitude), is given by

$$z = r_{LOS} \cos a_1 \#(3)$$

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

$$r = r_{LOS} \sin a_1 \#(4)$$

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

$$r_{LOS} = b \frac{\sin a_2}{\sin(a_1 - a_2)} \tag{5}$$

From z, r and viewing azimuth, the location of each pixel's viewed point is constrained, and readily

transformed back into the planetocentric cylindrical 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.

2.5   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.

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, the elevation accuracy was measured as a function of radius from nadir. This was achieved by segmenting the difference cloud into radial bins (measured perpendicular to the descent axis). The elevation root mean square error (RMSE) was calculated for each bin individually.

Whilst the ground truth used in this study represents the Martian surface in physical units, stereoscopy is scale invariant and the ground truth could just as reasonably span 1 m or 1000 km (with the caveat that surface curvature will manifest differently in each case, unless it too scales). Rather than presenting results in units tied to the physical dimensions of the ground truth dataset (i.e. metres), values are normalised relative to imaging geometry and camera resolution, in order to maintain their general applicability.

2.6 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, and this study, perhaps unusually, utilises a general purpose disparity algorithm to perform its stereo matching. Given that a large contributing factor in the accuracy of topography measurements is the stereo matching performance, an error model is here presented which assists with contextualising the results of this study, and drawing comparison with other stereo matching software.

As discussed in the previous section, Figure 5 and equations 3-5 describe the stereo viewing geometry by which surface points are triangulated. Performing error analysis yields

$$\sigma(z)^2 = \sigma(r_{LOS})^2 \left(\frac{\partial z}{\partial r_{LOS}}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial z}{\partial a_1}\right)^2 \tag{6}$$

$$\sigma(r)^2 = \sigma(r_{LOS})^2 \left(\frac{\partial r}{\partial r_{LOS}}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial r}{\partial a_1}\right)^2 \tag{7}$$

where

$$\frac{\partial z}{\partial r_{LOS}} = \cos a_1 \tag{8}$$

$$\frac{\partial z}{\partial a_1} = -r_{LOS} \sin a_1 \tag{9}$$

15

$$\frac{\partial r}{\partial r_{LOS}} = \sin a_1 \tag{10}$$

$$\frac{\partial r}{\partial a_1} = r_{LOS} \cos a_1 \tag{11}$$

and

$$\sigma(r_{LOS})^2 = \sigma(b)^2 \left(\frac{\partial r_{LOS}}{\partial b}\right)^2 + \sigma(a_1)^2 \left(\frac{\partial r_{LOS}}{\partial a_1}\right)^2 + \sigma(a_2)^2 \left(\frac{\partial r_{LOS}}{\partial a_2}\right)^2 \tag{12}$$

where

$$\frac{\partial r_{LOS}}{\partial b} = \frac{\sin a_2}{\sin(a_1 - a_2)} \tag{13}$$

$$\frac{\partial r_{LOS}}{\partial a_1} = -b \frac{\sin a_2 \cos(a_1 - a_2)}{sin^2(a_1 - a_2)} \tag{14}$$

$$\frac{\partial r_{LOS}}{\partial a_2} = b \frac{\sin a_1}{sin^2(a_1 - a_2)} \tag{15}$$

This study assumes exact knowledge of the stereo geometry, meaning $\sigma(b) = 0$. Because images were simulated with equidistant geometry (meaning uniform angular resolution), $\sigma(a_1)$ and $\sigma(a_2)$ are assumed to be constant over the camera's 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). The variance of e is therefore given by

$$\sigma(e)^2 = \sigma(z)^2 + s_{RMS}^2 \, \sigma(r)^2 \tag{16}$$

Where s_{RMS} is the ground truth's RMS slope. Figure 6 shows the RMS slope of the ground truth DTM, measured over a range of horizontal baselines. Also plotted is a histogram of the relative frequency of radial GSDs with which the surface is imaged in the full set of simulated images used in this study. The majority of pixels' radial GSDs are below 120 m. The RMS slope values corresponding to baselines <120 m are 0.18-0.19, and we therefore take 0.19 as the value of s_{RMS} throughout this paper.
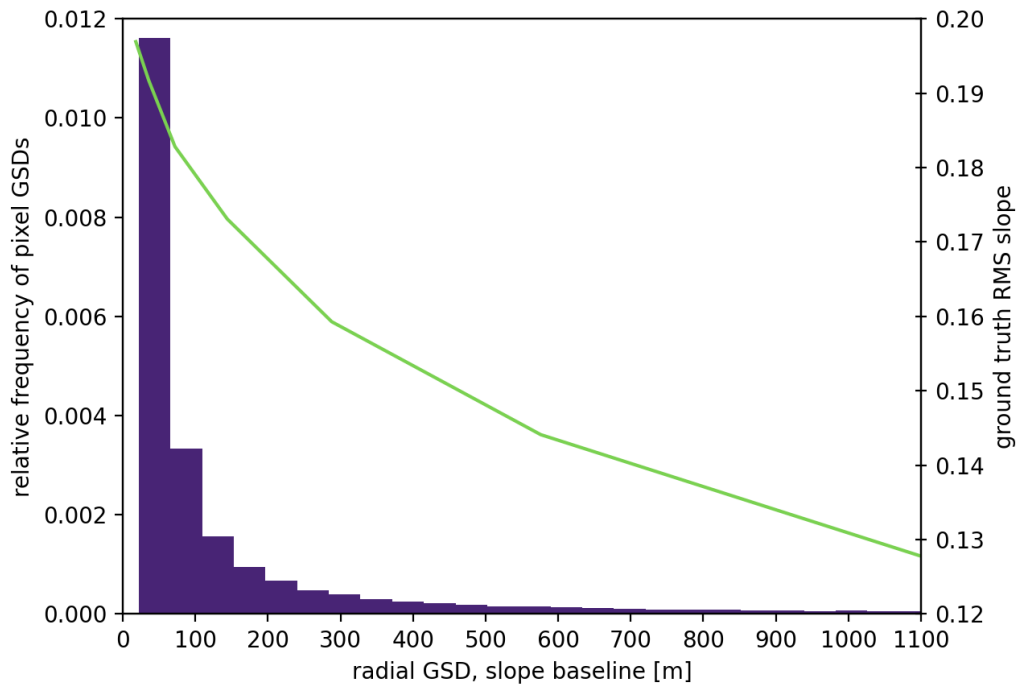
Figure 6: Histogram (purple) of relative frequency of radial GSDs of the pixels in the full set of simulated images used for topography measurement, plotted with the RMS slope (green) of the ground truth DTM measured over a range of baselines.

The following three plots illustrate the nature of elevation measurement error, as modelled by equations 6-15. Figure 7 shows the behaviour of elevation error as a function of ground radius, for three different stereo baselines, 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. Also plotted are the radial GSD (the camera's GSD measured radially away from its nadir) and the azimuthal GSD (the camera's GSD measured perpendicular to the radial GSD). Plotted values are normalised to the GSD of the lower camera's nadir-viewing pixel.
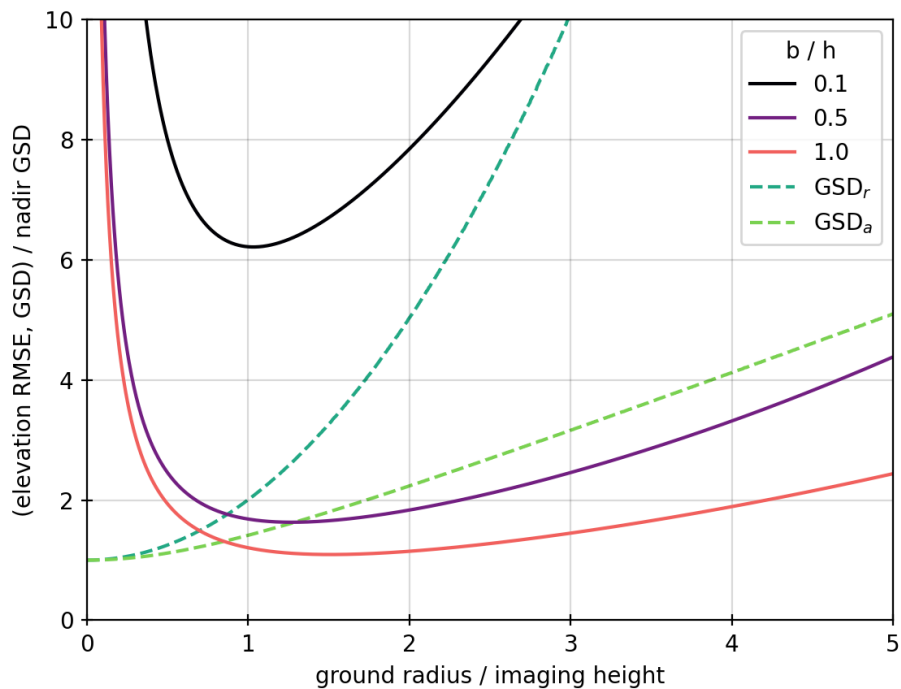
Figure 7: Modelled elevation RMSEs and GSDs as a function of ground radius, when imaging from a fixed height. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution.

It is evident from Figure 7 how significant the wide angle viewing geometry is in dictating topographic measurement accuracy. Error grows to infinity as ground radius reduces, because the surface viewing angles approach parallel with the stereo baseline, reducing the observed parallax. Large surface 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 the baseline. Larger baselines achieve greater elevation measurement accuracy, and exhibit a larger range of ground radii over which elevation measurement is optimal. Note also the pixel GSDs, which contribute to the horizontal grid spacing of topographic measurements (though this is not quantitatively assessed here). Radial GSD grows significantly faster than azimuthal GSD, due to increasing viewing obliquity.

Figure 8 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 7. Again, radial and azimuthal GSDs are also potted. Values are normalised to the GSD of the camera's nadir viewing pixel were the camera is at an altitude equal to the ground radius of the observed point.
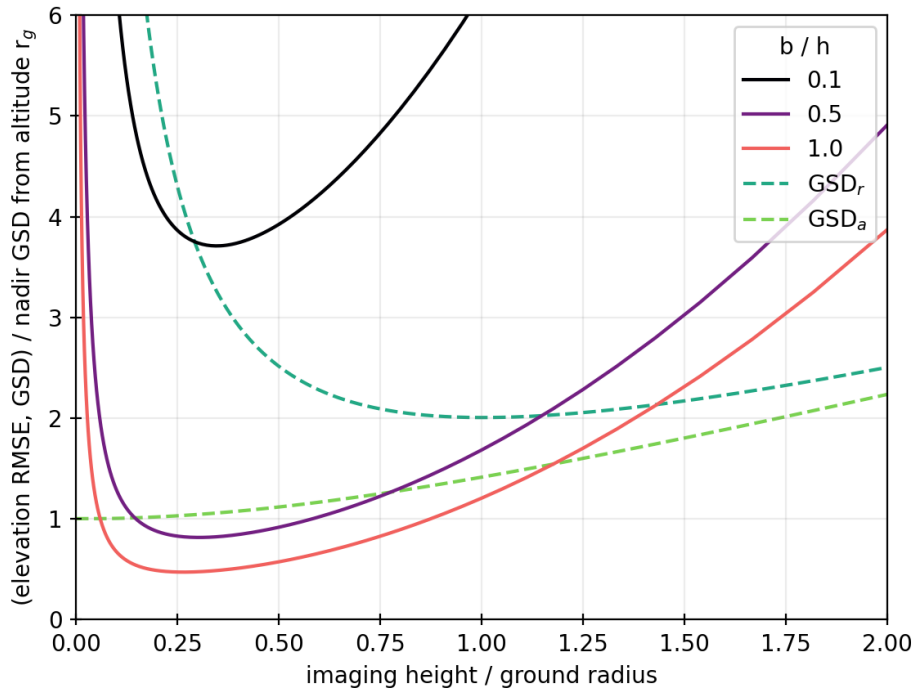
Figure 8: Modelled elevation RMSEs and GSDs as a function of imaging altitude, when viewing a point on the surface with fixed ground radius. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution.

Figure 8 reveals a clear minimum in elevation error as a function of height, meaning an optimum imaging height exists for any given point on the surface. For the plotted baselines, optimum height-to-ground-radius ratio occurs between 0.25 and 0.3. Radial GSD also exhibits a minimum, though at a significantly different height to those of the elevation error curves. It should be noted that the minima in elevation measurements only occur for non-zero values of s_{RMS} in equation 16, otherwise the elevation RMSE decreases monotonically with decreasing imaging height (similarly to azimuthal GSD).

Finally, Figure 9 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 ground radius to imaging altitude, are plotted. Modelled slope and angular errors remain unchanged from Figure 7.
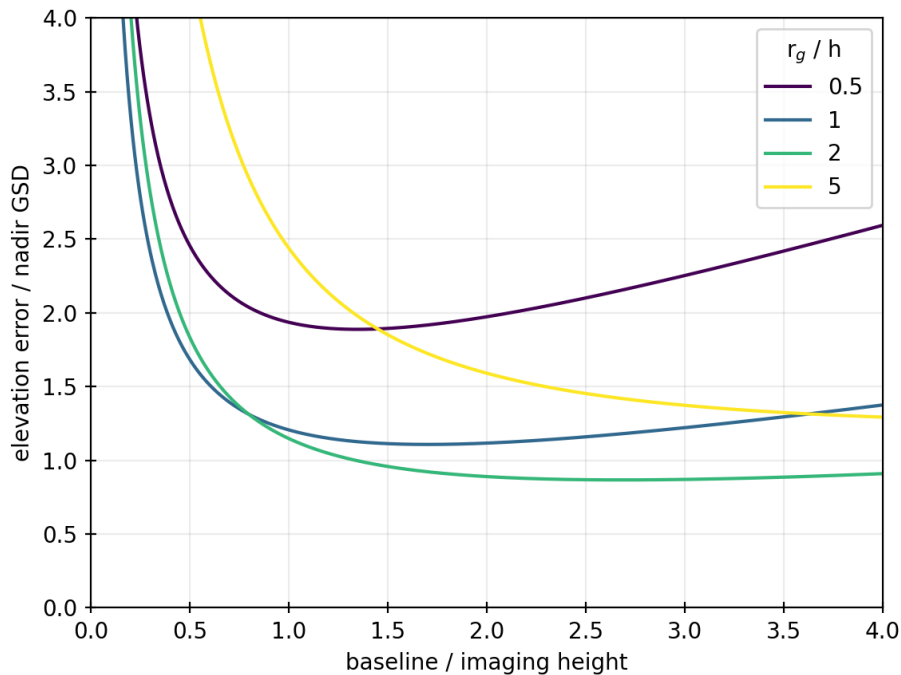
Figure 9: Modelled elevation RMSE as a function of stereo baseline. Surface RMS slope is 0.19 and angular error is 0.2x the camera's angular resolution.

Modelled curves of the types plotted in Figure 7 and Figure 8 are plotted alongside the measured results presented in section 3 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 flattened 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 2.4 (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 $\sigma(a_2)$ (on the assumption that $\sigma(a_1) = \sigma(a_2)$), 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. By this method, 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 10.
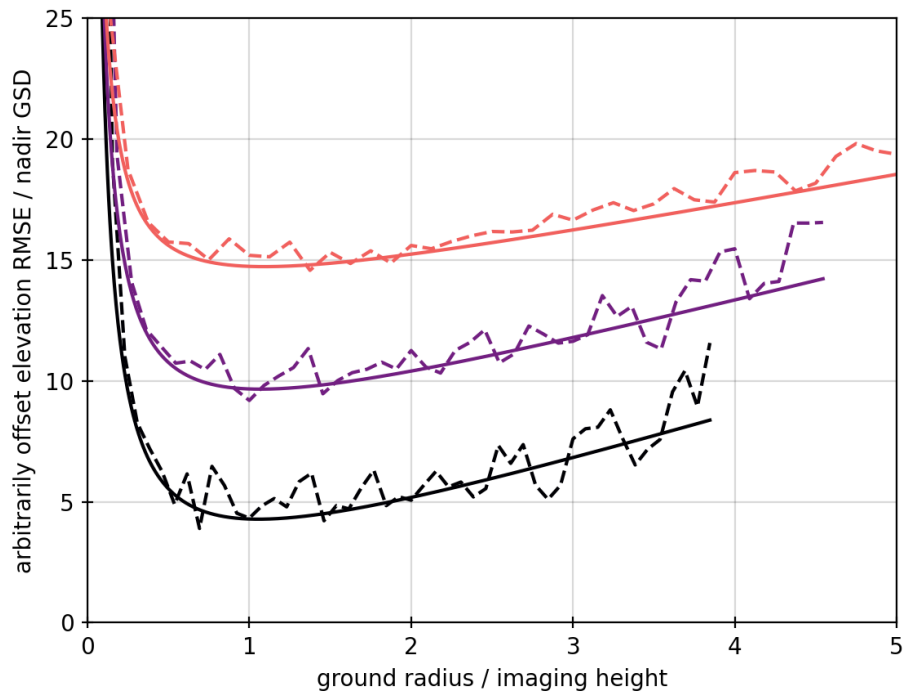
Figure 10: Example modelled elevation RMSEs (solid) and measured elevation RMSEs (dashed) for a flat surface, from three different stereo pairs of simulated images (offset vertically for clarity). Modelled angular error is 0.11x pixel angular resolution for all three curves.

Because the modelled RMSE curves assume constant values for the angular errors $\sigma(a_1)$ and $\sigma(a_2)$, differences between modelled and measured curves provide insight into the effect and occurrence of stereo matching errors (for example those potentially caused by scale changes between images, as discussed in section 2.4).

3 Results

Figure 11 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. Surface detail can be discerned in the diagrams, particularly where the points are so dense that they form contiguous surfaces. 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.

There are clear visible differences between the four example point clouds of Figure 11. 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 are, at least in part, a feature of 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 failed matching around nadir is significant, and the low altitude results in widespread blind spots.
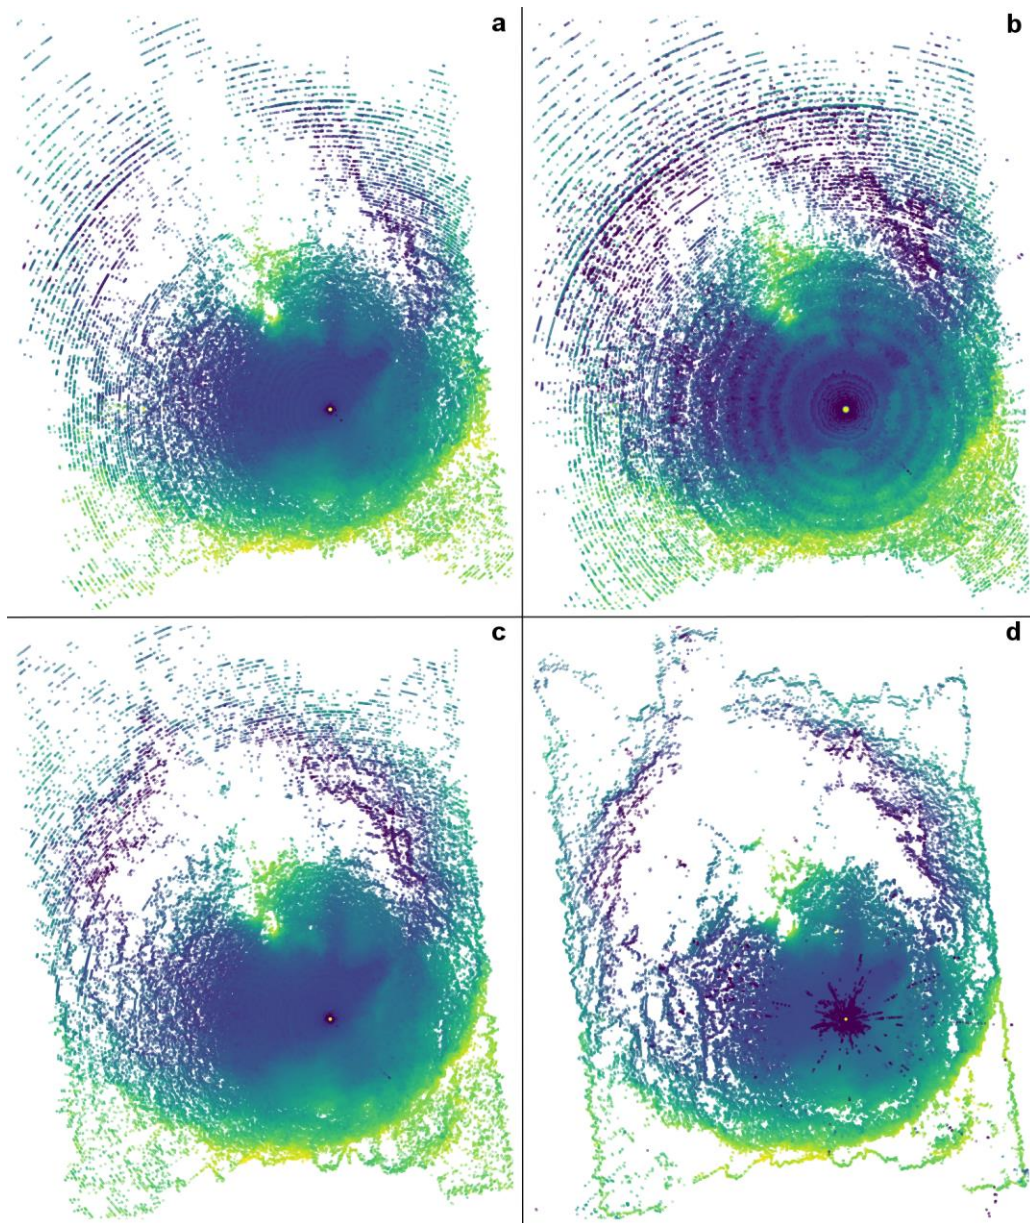
Figure 11: Four example elevation point clouds derived in this study. They were obtained from images with baseline/altitude of (in the ground truth's physical units) 1/5, 1/12, 2/7, 6/5 km for a, b, c and d respectively. All are displayed with the same elevation scale with dark (light) shading indicating low (high) terrain. White indicates no data.

3.1   Error Versus Ground Radius

Figure 12 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 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.
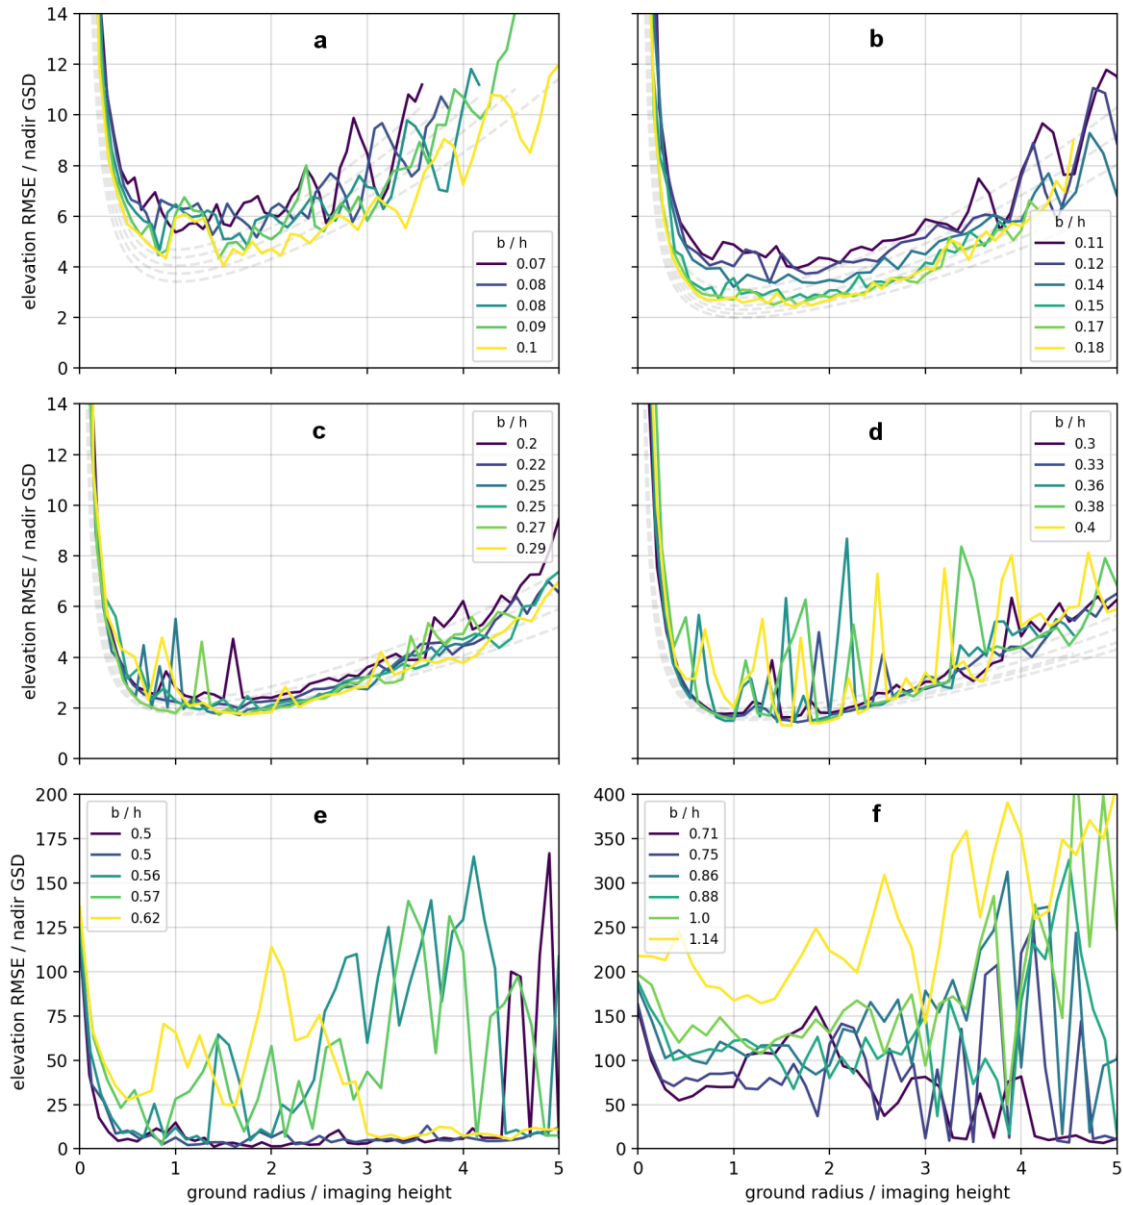

Figure 12: Measured elevation RMS error as a function of radius from nadir for 33 stereo pairs (solid coloured lines), segmented into six separate plots based on their different b/h values. Dashed grey lines are the modelled error for each stereo geometry, assuming a surface RMS slope of 0.19 and angular 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 12 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 to 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.

3.2   Error Versus Imaging Height

Figure 13 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 12. 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 when imaging from 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 12).
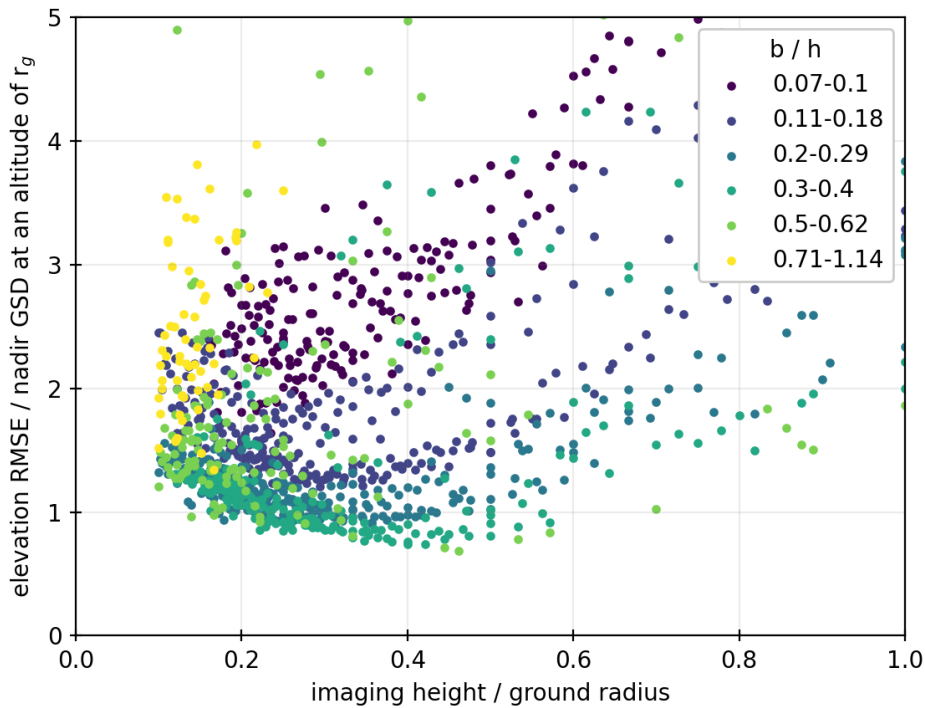
Figure 13: Measured elevation RMS error as a function of imaging height for a fixed ground radius, for 33 stereo pairs (same as are plotted in Figure 12).

There is a large spread in the full dataset plotted in Figure 13, and the shape of the modelled curves presented in section 2.6 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 0.2-0.4.

3.3   Discussion

The geometry of a stereo pair has a significant impact on the accuracy of elevation measurement, with two key contributing factors: the length of the stereo baseline, and the angles at which the surface is viewed in the two images (the difference between these two angles being the mechanism by which points' positions are triangulated). In our results, and the modelled data of section 2.6, the latter factor

is evident from two optimal relationships: the optimum ground radius to image using a stereo pair of a given height, and the optimum height of a stereo pair with which to image a given ground radius. Recall that the 'height' of a stereo pair here means the height of its lower image. A ground-radius-to-height or height-to-ground-radius ratio therefore constrains the viewing angle of a stereo pair's bottom image, whilst providing an approximation of the upper image's viewing angle (assuming a limited range of possible baseline lengths). The latter relationship (optimum height from which to image a given ground radius) has practical value, as it can translate to mission observation planning.

Defining an ideal height-to-ground-radius ratio range (characterised by its minimum $(h/r_g)_{min}$, and maximum $(h/r_g)_{max}$) can constrain the altitudes at which stereo pairs should be captured. For a stereo pair with height h, the ground radii of points on the surface whose elevations are measured with minimum error conform to the relationship

$$\frac{1}{(h/r_g)_{max}} \leq \frac{r_g}{h} \leq \frac{1}{(h/r_g)_{min}} \tag{17}$$

Let $h_1, h_2, \dots, h_n$ denote the decreasing heights of a set of stereo pairs captured during a descent. For contiguity of images' regions of minimum elevation error:

$$\frac{1}{(h/r_g)_{max}} h_n \leq \frac{1}{(h/r_g)_{min}} h_{n+1} \tag{18}$$

and therefore

$$h_{n+1} \geq \frac{(h/r_g)_{min}}{(h/r_g)_{max}} h_n \tag{19}$$

describes the required heights of the lower images of stereo pairs. The ideal heights of upper images are also informed by equation 19, in conjunction with any constraints on the stereo baseline size.

4 Conclusions

It is of no surprise that imaging geometry has a strong impact on the outcomes of descent stereophotogrammetry presented in this paper, given that the connection between baseline and depth measurement error is well established for the more conventional cases of aerial and satellite imagery (e.g. Hallert (1960); Johnsson (1960)). The accuracy of descent stereophotogrammetry is a strong function of viewing angle, thus varying significantly over the surface of the observed body and resulting in optimal stereo geometries. From the measured elevation RMSEs, 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 2.6). Additionally, the elevation RMSEs of our simulated topography measurements exhibit an optimum baseline-to-height ratio of 0.2-0.4.

Notably, the model curves of Figure 7, Figure 8 and Figure 9 indicate significantly larger optimum baseline-to-height ratios of at least 1 for near surface, and as much as 3-4 for distant surface. For the larger baselines in our measured data, elevation measurement error grows rapidly (by two orders of magnitude once b/h reaches 1), and ceases to follow the model behaviour. The reason for this difference 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.

An optimum viewing geometry, whether the same as that found in this study or different due to an adapted imaging/stereo-matching setup, should be used to inform the acquisition of stereo descent images. Equation 19 can be used to determine the heights at which the bottom images of stereo pairs should be obtained, and the heights of the upper images can subsequently be determined based on baseline restrictions. Equation 19 requires that images be captured at increasing frequency as the camera altitude decreases if optimum geometry is to be maintained. It is important to note also that imaging geometry does not only dictate the accuracy of stereoscopic elevation measurements. It also

influences the horizontal resolution of the images (in particular, their radial GSD), and by extension the horizontal resolution of the topographic measurement.

The findings of this study apply to the general concept of wide-angle descent stereophotogrammetry, and represent the first comprehensive quantitative analysis of the technique. Descent stereophotogrammetry could be carried out with a wide variety of planetary exploration missions, including penetrators, powered descent stages, balloon probes and aerial vehicles, meaning the findings of this study have wide reaching value. Suitable images could be captured with a range of imager designs, suited to both 3-axis and spin stabilised platforms, which could be compatible with compact, low mass missions. For certain planetary missions (e.g. a surface penetrator mission) descent stereophotogrammetry may be the only means by which topography measurement can be obtained, and its inclusion can therefore provide a valuable science boost.

Aside from these conclusions, this study also demonstrates the applicability of high fidelity DTMs and surface images to simulated studies of imaging instrumentation. The methods presented in this paper could be readily utilised and adapted for the study of specific descent stereophotogrammetry setups.

5 References

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.

Bradski, G., 2000. The OpenCV library. *Dr. Dobb's Journal of Software Tools.*

Brydon, G. & Jones, G. H., 2018. Rotational push-broom imaging from a planetary penetrator. *EPSC meeting, EPSC2018-994,* Volume 12.

Crisp, J. A. et al., 2003. Mars Exploration Rover mission. *J. Geophys. Res.,* Volume 108.

Daudon, C. et al., 2020. A new digital terrain model of the Huygens landing site on Staurn's largest moon, Titan. *Earth and Space Science,* Volume 7.

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.

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.

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.*

Gowen, R. A. et al., 2011. Penetrators for in situ subsurface investigations of Europa. *Adv. Space Res.,* Volume 48, pp. 725-742.

Grotzinger, J. P. et al., 2012. Mars Science Laboratory mission and science investigation. *Space Science Reviews,* Volume 170, pp. 5-56.

Hallert, B., 1960. *Photogrammetry basic principles and general survey.* 1 ed. New York: McGraw-Hill.

Heipke, C. et al., 2007. Evaluating planetary digital terrain models—The HRSC DTM test. *Plan. and Space Sci,* Volume 55, pp. 2173-2191.

Hirschmuller, H., 2008. Stereo processing by semiglobal matching and mutual information. *Pattern Analysis and Machine Learning, IEEE Transactions on,* Volume 30, pp. 328-341.

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.

Jia, Y. et al., 2018. The scientific objectives and payloads of Chang'E-4 mission. *Plan. and Space Sci.,* Volume 162, pp. 207-215.

Johnsson, K., 1960. On the accuracy of stereo-plotting of convergent aerial photographs. *Photogrammetria,* Volume 17, pp. 83-98.

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.

Kim, J. R. & Muller, J. -P., 2009. Multi-resolution topographic data extraction from Martian stereo imagery. *Plan. and Space Sci,* Volume 57, pp. 2095-2112.

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.

Li, C. et al., 2015. The Chang'e 3 mission overview. *Space Sci. Rev.,* Volume 190, pp. 85-101.

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., 2011. Planetary penetrators: their origins, history and future. *Advances in Space Research,* Volume 48, pp. 403-431.

Maki, J. N. et al., 2003. Mars Exploration Rover Engineering Cameras. *J. Geophys. Res,* Volume 108, pp. 1-24.

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. 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.

Moratto, Z. M. et al., 2010. Ames Stereo Pipeline, NASA's open source automated stereogrammetry software. *Lunar and Planetary Science Conference, 41st,* p. 2364.

Mottola, S. et al., 2015. The structure of the regolith on 67P/Churyumov-Gerasimenko from ROLIS descent imaging. *Science,* Volume 349.

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.

Schneider, D. & Schwalbe, E., 2005. Design and testing of mathematical models for a full-spherical camera on the basis of a rotating linear array sensor and fisheye lens. *Proc. 7th Conference on Optical 3-D Measurement Techniques,* pp. 245-254.

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.

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.

Tao, Y. et al., 2018. Massive stereo-based DTM production for Mars on cloud computers. *Plan. Space Sci,* Volume 154, pp. 30-58.

Tomasko, M. G. et al., 2002. The Descent Imager/Spectral Radiometer (DISR) experiment on the Huygens entry probe of Titan. *Space Sci. Rev.,* Volume 104, pp. 469-551.

Walker, S., 2007. New features in SOCET SET. *Photogrammetric Week,* Volume 7.

Wray, J., 2013. Gale crater: The Mars Science Laboratory/Curiosity rover landing site. *International Journal of Astrobiology,* Volume 12, pp. 25-38.