

# Hardware-in-the-Loop Testing of Stereo Vision-Based Hazard-Detection Method for Planetary Landing

Svenja Woicke\*

*Delft University of Technology, Faculty of Aerospace Engineering, Delft, The Netherlands*

Hans Krüger †

*German Aerospace Center, Institute of Space Systems, Bremen, Germany*

Erwin Mooij‡

*Delft University of Technology, Faculty of Aerospace Engineering, Delft, The Netherlands*

**Hazard-detection and avoidance systems will become an important asset for next-generation landing and exploration missions. To date, multiple studies into these systems were conducted to develop the methods needed and to demonstrate their performance in mainly software-based tests. Few studies were able to demonstrate the performance of the algorithm in hardware-in-the-loop tests, as these are usually difficult to set up and expensive to execute. In this paper, the hardware-in-the-loop testing of a stereo-vision based hazard-detection algorithm is presented. It was performed with the Testbed for Robotic Optical Navigation (TRON) at the German Aerospace Center (DLR) in Bremen, Germany. Since this testbed only allows for testing in a scaled environment, one of the challenging tasks was to design a scaled-down test set-up to represent a real-life lunar descent. The hardware-in-the-loop testing confirmed the results obtained during the earlier software-in-the-loop testing, that stereo vision can successfully be used for hazard detection during planetary descent.**

## Nomenclature

$b$	= baseline, [m]
$C$	= cost, -
$d$	= disparity, [pixel]
$f$	= focal length, [pixel]
$I$	= image intensity, -
$u$	= image coordinates in right-direction, [pixel]
$v$	= image coordinates in downwards-direction, [pixel]
$Z$	= boresight distance from the camera to the surface, [m]
$\mu$	= mean
$\sigma$	= standard deviation

## I. Introduction

In the context of performing precise landings in poorly known or challenging planetary environments, hazard-detection and avoidance (HDA) systems will become increasingly important. Examples for such scenarios are landings on bodies with only limited surface knowledge, such as Jupiter's moon Europa. Also landing in more challenging regions on Mars may require HDA systems. Here, also landing close to, while not crashing into, pre-existing surface assets, for example, sample containers, will be a future application of HDA.

---

\*PhD candidate, Section of Astrodynamics and Space Missions (s.woicke@tudelft.nl).

†Researcher, Department of GNC Systems (hans.krueger@dlr.de).

‡Assistant professor, Section of Astrodynamics and Space Missions. Associate Fellow, AIAA (e.mooij@tudelft.nl).

Within the HDA system, reconstruction of the surface in the landing region and analysing it for safety is one of the main tasks, the so-called hazard detection (HD). The other main task is the aspect of reaching and touching down at a safe landing site, called hazard avoidance (HA). This paper focuses on the task of hazard detection.

Hazard detection can be approached in two ways: using passive sensors, i.e., cameras,<sup>1,2</sup> or active sensors, i.e., lidar.<sup>2-6</sup> In this research, a camera-based method is used, more precisely a method using the principle of stereo vision for reconstruction of the digital elevation model (DEM) of the landing region. It should be noted that both camera and lidar-based approaches have their own advantages and disadvantages, with cameras being superior from the point of view of financial, power and weight budgeting, and lidars being a more direct method of measuring the surface elevation. To date, no consensus towards a clear preference for either one of the two approaches have been reached.

This work is the extension of the work presented in Refs. 1 and 7, where detailed software-in-the-loop testing (SILT) of the stereo vision-based hazard detection algorithm is presented. In this paper, the hardware-in-the-loop testing (HILT) of this algorithm is presented, demonstrating that the algorithm is indeed a method that can be used for hazard detection using real camera images as input.

This paper is organised as follows. In Sec. II, a short introduction to stereo vision as a hazard detection method is given. Moreover, some exemplary hazard-detection results of this method obtained from software-in-the-loop tests are presented. Next, the set-up of the hardware-in-the-loop test and an introduction of the facility used are given in Sec. III. This is followed by a presentation and discussion of the results in Sec. IV. A more detailed analysis of the results is performed in Sec. V. The paper closes with conclusions drawn from the testing in Sec. VI

## II. Stereo-Vision Hazard-detection Algorithm

The stereo vision algorithm presented in this section is discussed in more detail in Ref. 1. This section will introduce the reader to the topic of stereo vision and briefly outline the algorithm and design choices made. To perform stereo vision, two cameras need to be attached to the lander with a known baseline. Preferably, this baseline is as rigid as possible, as any unknown relative movement of the cameras is equivalent to errors to uncertainties in the extrinsic calibration of the set-up. An on-board recalibration of a stereo set-up may be possible to some extent, but should be avoided, as it is time consuming and adds complexity to the system.




Figure 1: Stereo-vision principle.

The two cameras will generate two input images for the stereo vision algorithm. Figure 1 shows an illustration of a stereo vision set-up. It is important to note that the cameras are oriented in such a way that their optical axes are parallel. This results in two images, which are not fully overlapping; towards very low altitudes the overlapping portion shrinks and finally vanishes. However, this is not considered to be a problem, since hazard detection is performed at higher altitudes.




Figure 2: Matching of stereo-vision images.

Stereo vision makes use of the fact that when taking two pictures of the same scene from different viewpoints, the projections of the same ground points will appear at different image positions within the two images. Figure 2c shows this principle, as well as the two input images (Figs. 2a and b) used during the software-in-the-loop testing. The change in image position for a given ground point, can be used to compute the distance towards this point from triangulation. Finding a pixel in the left image corresponding to the pixel representing the same ground point in the right image is called image matching. Comparing the two different locations results in a value for the shift of that pixel. This shift is the so-called disparity,  $d$ . Corresponding pixels will be located along the same epipolar line. The epipolar line is the intersection of the epipolar plane and the image plane, where the epipolar plane is formed by the two optical centres and an image point. For the SILT testing presented here, perfect cameras and perfect parallel alignment of the cameras is assumed. In such a case the epipolar lines of are inherently aligned and no further actions have to be taken.

The maximum depth resolution (*or*: the minimum detectable difference in  $z$ ),  $\delta Z$ , the paramount factor for creating useful HD surface maps, is given by

$$\delta Z = \frac{Z^2}{bf} \delta d \quad (1)$$

where  $Z$  is the boresight-distance from the camera to the surface,  $b$  is the baseline,  $f$  is the camera's focal length and  $\delta d$  is the minimum detectable disparity. From this equation it can be concluded that the depth resolution scales linearly with increasing baseline and that the baseline should be maximised. However, the maximum achievable baseline is limited by the lander dimension and structural rigidity and is assumed to be limited to 2 m to 3 m, based on current lander designs. To work with the limiting case rather than the more relaxed constraint, a maximum baseline of 2 m will be assumed in this research.

For image matching, a cost function is used to determine which two pixels in their respective images represent the same ground points. The algorithm used employs the sum of squared difference (SSD) as cost function, since this was found to deliver the best results after an extensive trade-off.<sup>1</sup>

$$\text{SSD} = \sum_{u,v} (I_1(u,v) - I_2(u+d,v))^2 \quad (2)$$

Since Eq. (1) shows that the depth resolution also scales with the minimum detectable disparity, it is important to achieve non-integer, so-called sub-pixel, disparities. This is done using a parabolic fit through the point of minimum disparity:

$$\hat{d}_{\min} = d_{\min} - \frac{C(d_{\min} - 1) - C(d_{\min} + 1)}{4C(d_{\min}) - 2C(d_{\min} - 1) - 2C(d_{\min} + 1)} \quad (3)$$

Lastly, knowing the disparity, the distance from the cameras to the surface can be computed using

$$z(d) = \frac{fb}{d}, \quad (4)$$

where both the focal length and the baseline are fixed by the design, and the disparity is found from image matching. Finding the match for each pixel in the left image and computing the distances from the camera to the surface for each of them will lead to a dense DEM.

From Eq. (1) it is apparent that the distance to the surface is the driving factor for the depth resolution, as the depth resolution scales inversely with the square of this distance. From this observation it is obvious that stereo vision is not a suitable method to detect surface hazards at very high altitudes without considering unrealistic baseline values. However, boulders, the major landing surface hazard, can only be detected if the image resolution is high enough to resolve these features. Therefore, boulders can only be successfully detected at low altitudes, which are also the altitudes at which stereo vision performs well. Using this algorithm it was established that for a 2 m stereo baseline, stereo vision can be applied for altitudes of 200 m and below, which is a very suitable altitude for detection of boulders. For a field of view (FOV) of 30° and a image size of 512 x 512 pixel, this leads to a minimum resolution of  $\approx 3$  m at 200 m altitude and 1 m at 150 m altitude.

Using these dense DEMs, it is possible to compute a hazard map of the landing region by computing surface characteristics which can cause a landing hazard. In the algorithm used these hazards are:

- Slope
- Roughness
- Insufficient illumination/shadow

Slope and roughness can both be computed from the DEMs and are quantitative assessments of the surface. Illumination is determined from the input images. In addition to the roughness computed from the DEM, it is also possible to perform a qualitative assessment of surface roughness based on image texture.

Slope is computed as the mean plane over a patch with the size of the lander footprint, while roughness is calculated as the local deviation from this mean plane, in the direction of the mean-plan's normal. The texture-based roughness is computed from a histogram-based variance. A detailed discussion of the analysis of different methods that lead to these choices can be found in Ref. 1.

Figure 3 shows the result of a software-in-the-loop test, for a landing site imaged from 200 m altitude. Figure 3a presents the resulting DEM, Fig. 3b the slope map and Fig. 3c the roughness map, as computed from the DEM. The roughness map computed from the image texture is shown in Figure 3d. The final hazard map obtained by merging all the inputs presented in Figs. 3a-d is presented in Fig. 4a. These maps are rescaled on a scale from 0 to 1, where all values above the of the selected hazard thresholds (slope  $\geq 15^\circ$  and roughness  $\geq 0.5$  m) are assigned the value 1. The merging is performed by adding all individual, rescaled, maps with equal weights.

To evaluate the performance of the hazard detection function, it is necessary to compare the outcome of the function to the ground truth hazard map, computed from the ground truth DEM; the resulting map is given in Fig. 4b. From these maps it is apparent that the algorithm tends to slightly overestimate hazards, as some safe sites are labelled hazardous. However, less than 1% of all hazards are overseen. Therefore, the algorithm presented here delivers good performance in a conservative sense.

### III. Hardware-in-the-loop Testing

To extend the performance evaluation of the algorithm, HILT is performed. These tests are executed at the Testbed for Robotic Optical Navigation (TRON) at the German Aerospace Center (DLR) in Bremen, Germany<sup>8,9</sup>. In the following, first some general aspects of the facility will be introduced, after which the test set-up will be described.

#### A. Facility

TRON is a HILT facility with the purpose of supporting the development of optical navigation technology. TRON provides an environment, which allows qualifying breadboards to Technology Readiness Level (TRL) 4, and qualifying flight models to TRL 5-6. Typical sensor hardware, which can be tested in TRON, are






Figure 3: Intermediate results of stereo vision hazard mapping SILT.<sup>1</sup>

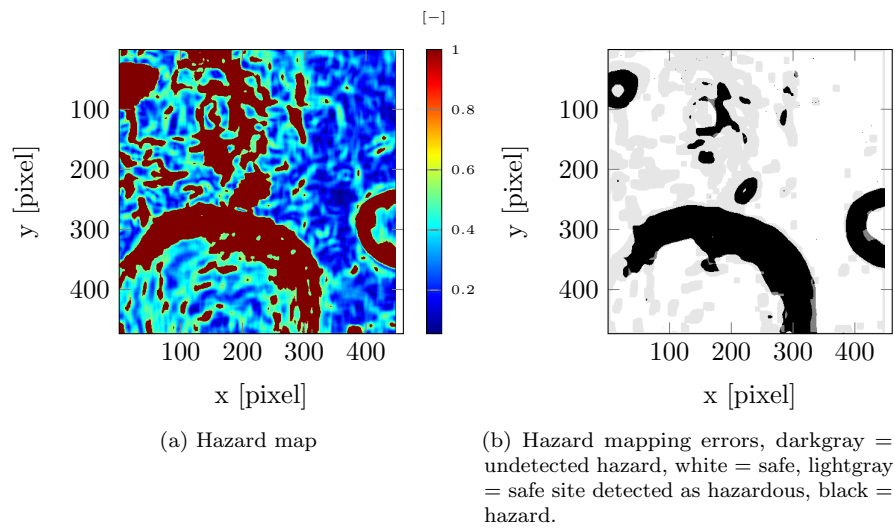


Figure 4: Hazard mapping results of SILT.<sup>1</sup>

active and passive optical sensors like lidars and cameras. The major components of the laboratory are a robot on a rail for dynamic positioning of the sensor under testing, a dynamic lighting system for illumination of the targets, laser metrology equipment for high precision ground truth, and a dSPACE real-time system for test observation and control and synchronization of ground truth and sensor data. The laboratory can be customised with user-defined hardware, such as models of a Lunar or Martian surface. Due to this flexibility, as well as its extensive dimensions, TRON is well suited for simulating missions representative of the ones encountered by optical sensors during exploration missions.




Figure 5: Layout of the TRON facility.

As shown in Fig. 5, TRON features three different terrain models located at different walls inside the laboratory. Each of these models is designed for a specific phase of approaching and landing on a body. Figure 6 shows a picture of the facility. All three terrain models are visible, as well as the robot, the laser tracker and the lighting system. Figure 5 also shows the T-Mac mounted on the payload platform together with the stereo cameras. The T-Mac serves as a target for the laser tracker to measure the distance from the target to the tracker as well as its orientation with respect to the laser tracker.

Terrain 3 is designed for simulating the last phase of a landing mission. In this way, the terrain can be used to not only perform hardware-in-the-loop tests of hazard detection and avoidance methods, as presented in this paper, but also methods for terrain relative navigation with respect to the landing site. Therefore, terrain model 3 is used for this research. An image of this model is given in Fig. 7.




Figure 6: Picture of the TRON facility.

Terrain model 3 is installed at the front wall of TRON. Its size is about 4.2 m x 2.2 m, the terrain dynamics




Figure 7: Picture of surface model 3.

is  $\approx 0.26$  m. The model reference data were obtained entirely by DLR via a process beginning with hand-modeling, and ending with 3D scanning and post-processing, as described in Ref. 10 . The model was then manufactured in two steps. At first, a coarse milling step obtained the rough terrain structure. Afterwards, a finishing surface layer was applied manually. Due to the hand-made finishing, no manufacture marks such as milling lines are visible, leaving the model with a practically infinite resolution. The self-similarity with respect to scale of the model and the Moon can be exploited by applying a different scale to this model. This landing-site model is not only representative for the Lunar surface, but also for many asteroid surfaces. Combining this model with low scale factors makes it a useful sensor target for 3D imaging sensors.

Moreover, a full ground-truth model of the surface exists, and its location is precisely known within the laboratory reference frame. This enables a very accurate comparison of simulation results to reality. The ground truth of Terrain 3 is presented in Fig. 8.

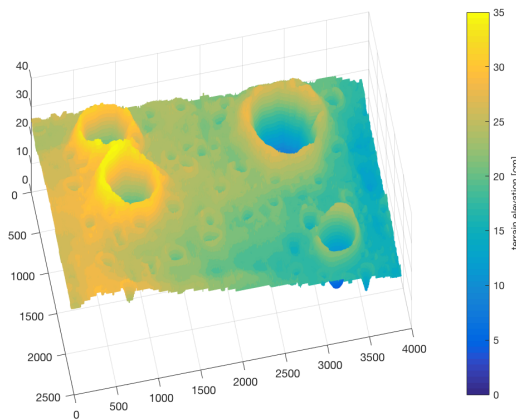


Figure 8: Laser scan of Terrain 3, the terrain dynamics in z-direction are exaggerated with respect to the other axes for better visibility.

The robot, carrying the (optical) sensors, can be moved along a rail from a distance of 11 m to a distance of 1 m to Terrain model 3, as also indicated in Fig 5. At the same time, it can be moved with a radius of  $\approx 1$  m perpendicular to the rail.

More detailed information on the set-up of TRON and the surface model can be found in Ref. 8 and 10.

## B. Test set-up

For testing the stereo vision algorithm, two cameras (the stereo pair) are attached to a plate on the robot's arm, see Fig. 9. These cameras and the laser tracker are hardware-triggered by the dSPACE system, which enables synchronised image capturing of both cameras, as well as recording the robot/camera position during the simulation. The cameras' full pose with respect to the laboratory reference system is determined by the laser tracker, which measures the T-Mac pose at 100 Hz. The T-Mac is fixed to the camera platform on the robot. This information is necessary for comparison of the simulation results to the ground-truth data, to perform a quantitative assessment of the algorithms performance.




Figure 9: Camera set-up on payload platform.

The cameras used are two Allied Vision Proscila GC1380H cameras, both with a Schneider Kreuznach Cinegon 1.9/10 10.4mm-F1.9 lens. The cameras are set to deliver an image of  $1024 \times 1024$  pixels, which is down-sampled to  $512 \times 512$  pixels to limit execution time. The field of view is  $\approx 37^\circ$ . The cameras' intrinsics and extrinsics are calibrated using checker-board patterns and DLR's camera calibration toolbox, DLR CalDe and DLR CalLab.<sup>11</sup> The same tool is used for hand-eye calibration of the camera set-up.

As discussed in Sec. II, the working envelope of the stereo vision HD algorithm starts at 200 m above the surface and current lander designs allow for baselines of 2 m to 3 m. Since testing in TRON is done in a scaled environment, the baseline has to be scaled accordingly. TRON's 10 m rail is used to simulate the 200 m of the final descent towards the ground. This represents a scale factor of 1:20. To obtain a set-up of a comparable depth resolution as the proposed landing set-up, a 30 cm baseline will be used. Here, it should be noted that the baseline is not scaled at the same ratio as the remaining set-up. This is due to the fact that the depth resolution was supposed to be kept comparable to the full set-up, i.e., scaled at 1:20. However, as discussed in Sec. II and supported by equation Eq. (1) the depth resolution does not scale linearly with the rest of the set-up. Therefore, 30 cm was established as a value representative of the full set-up by an iterative process.

In a real landing scenario, slopes of more than  $15^\circ$  over the size of the lander footprint (commonly on the order of 2 m to 5 m) and roughness of more than 0.3 m to 0.5 m, are considered to be hazardous.<sup>1</sup> As with the baseline, these values have to be adapted to the scaled environment. For the 1:20 scale used in this set-up, this leads to the following hazard thresholds:

- Slope:  $\geq 15^\circ$  over a patch of 10 cm to 25 cm
- Roughness:  $\geq 1.5$  cm

## IV. Results

In this section the DEM and the hazard map created from pictures taken at  $\approx 3$  m from the terrain model are presented. At this distance, the camera images show only the terrain model. Since a 1:20 scale environment is used, this is equivalent to a 60 m distance (at 2 m baseline).

After one stereo pair of two images is acquired, these images first need to be rectified, using the results from the intrinsic calibration. Rectification is necessary in the hardware test, opposed to the SILT, as epipolar



lines need to be aligned for stereo reconstruction. For purely translated cameras, as used during SILT, the epipolar lines are already aligned by default.

Figure 10 shows a composition of the rectified right and left stereo image. Careful visual inspection will lead to the conclusion that tpixels pertaining to the same physical points are indeed on the same row. This leads to the conclusions that a) rectification worked successfully and b) camera calibration was successful.




Figure 10: Rectified stereo pair(red - left image, cyan - right image).

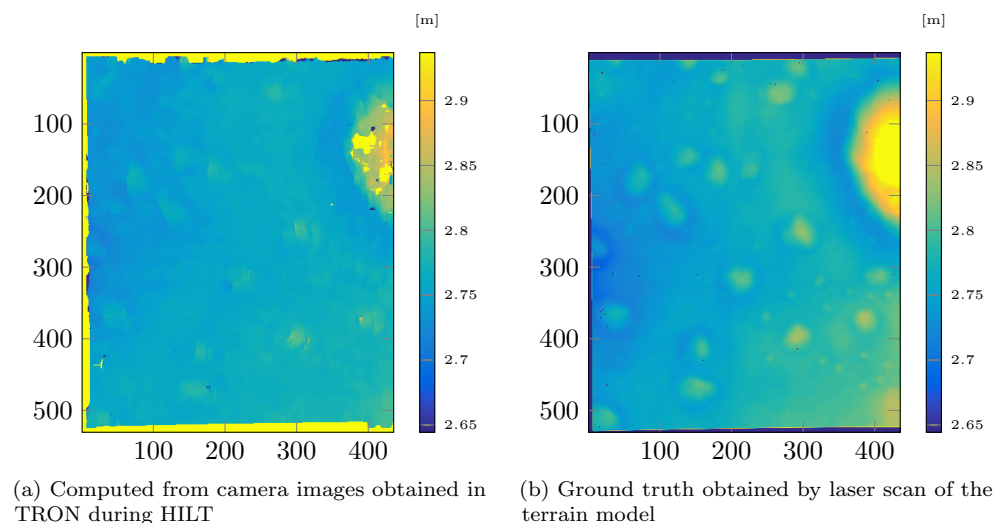


Figure 11: DEM from HILT obtained from TRON measurements.

Figure 11a presents the DEM as computed from the input pair shown in Fig. 10. The corresponding ground truth is given in Fig. 11b.

It can be concluded that over all the scene is reconstructed well, all craters seem to be in the DEM. It is clear that some errors are present in the computed DEM close to the image borders, which is due to the stereo reconstruction having more difficulties finding the appropriate matches here. However, it should also be noted that these regions are by design more prone to problems with calibration and optics. The extensive testing of the algorithm on a larger number of HILT stereo images also showed that insufficient and/or very uneven illumination can also result in reconstruction errors. It is thus very important to create realistic illumination conditions during HILT. In theory one would try to achieve a parallel lighting condition, however this is very difficult to reach in a laboratory where the distance between the light source and the terrain is

small and limited by its size. The DEM mean and median error in the z-direction of a rectified Cartesian coordinate frame are 0.037 m and 0.066 m, respectively, with a standard deviation of 0.034 m and a median absolute deviation (MAD) of 0.016 m. The error histogram in Fig. 12 shows that the error is not completely Gaussian, which is why it was chosen to give both the mean and the median. These error values permit successful hazard detection as envisioned in this paper.




Figure 12: Error histogram of the DEM error in z-direction in rectified Cartesian coordinates (HILT).

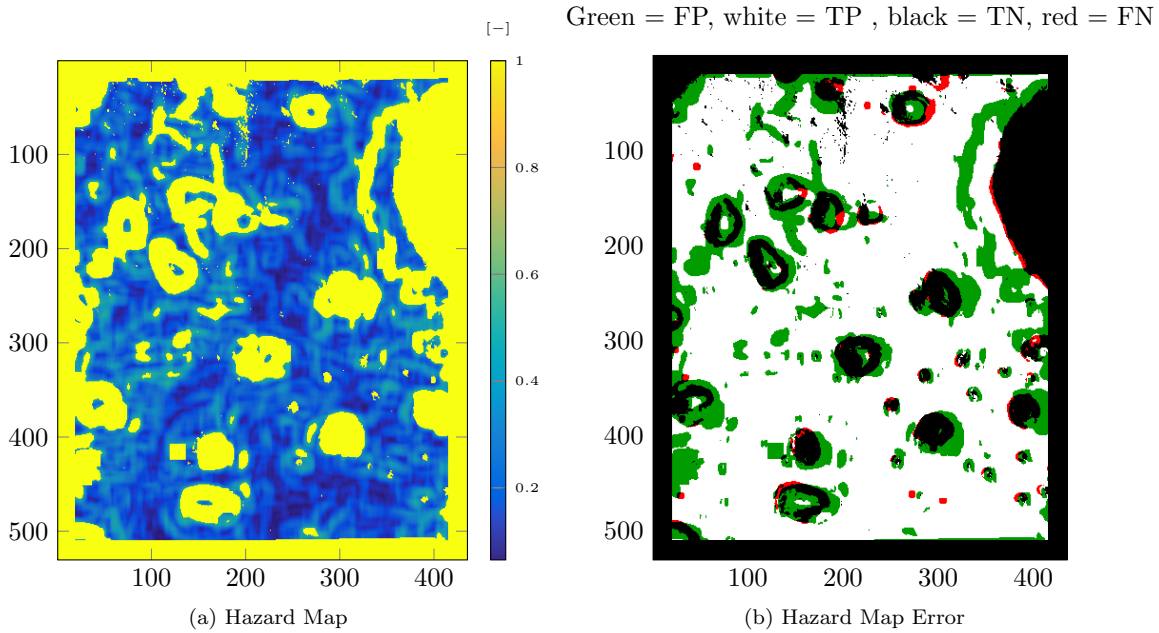


Figure 13: Hazard map, and hazard map compared to real hazard map (HILT).

Figure 13a presents the computed hazard map based on the output of the algorithm. This hazard map is based on slope, roughness and illumination extracted from both the DEMs and the input image. The hazard thresholds are set as discussed in Sec. III B. A value of 1 represents a location that would be hazardous. It can be seen that the craters are visible in the hazard map. Also, the higher slope on the left of the DEM is identified and labelled as hazardous. In Fig. 13b a comparison of the computed hazard map and the true hazard map is shown. This figure is colour-coded as follows: a red region represents a hazard, undetected by the algorithm (false negative, FN). A green region is a safe site which is erroneously labelled as an unsafe site

by the algorithm (false positive, FP). A white pixel represents a safe location that is correctly identified by the algorithm (true positive, TP). A black pixel is hazardous and correctly detected as such by the algorithm (true negative, TN). False negatives are the most dangerous errors, as these can lead to a mission loss when selecting a landing site that is actually unsafe.

Overall, there are 1.7% FN, 16.5% FP, 68.5% TP, and 13.3% TN in this scene. Overall, 88.6% of all hazards in the scene are correctly identified, while 80.6% of all safes sites are correctly identified. In total, 81.8% of all detections are correct.

Closely investigating Fig. 13b, one can see that both FN and FP errors are not randomly occurring, but undetected hazards occur in two different situations:

- in locations right next to correctly detected hazards. This could be accounted for by applying a safety factor to not select landing sites right next to detected hazards.
- The algorithm sometimes fails to detect small boulders. This is due to the smoothing effect of the window size used during stereo matching and the comparably small size of the boulder in the scene (just 2 pixels to 3 pixels each), which makes them very difficult to distinguish from image noise, while in the high quality ground truth they can be identified very accurately. This problem might be overcome by implementing a more advanced boulder detection algorithm in addition to the roughness detection already included.

Figure 14 shows the individual contributions that lead to the hazard map as presented in Fig. 13a, namely the slope in Fig. 14a, the roughness in Fig. 14c, with the ground truth given in Fig. 14b and Fig. 14d, respectively. The texture-based roughness is shown in Fig. 14e and the shadow map in Fig. 14f. It can be seen that the illumination conditions during the HILT were chosen in such a way that quite some shadows were present on the model. Moreover, it can be seen that DEM-based roughness (Fig. 14d) and texture-based roughness (Fig. 14e) show comparable, but not identical, results. Therefore, it is clearly advisable to use both methods simultaneously. This conclusion was also already reached during the earlier SILT evaluation.

Comparing the computed slope in Fig. 14a to its ground truth in Fig. 14b one can see that the algorithm performs quite well in computing the terrain's slope. As the computed DEMs are slightly smoother than the ground truth DEM, the computed slope is smoother than the true slope as well.

Figure 14c presents the true roughness. Comparing this to the computed roughness in Fig. 14d, one can see that the overall performance of roughness detection is good for larger-scale features, while the algorithm fails to detect small roughness features, e.g., very small rocks. This behaviour was already addressed above.

In summary, the results presented in this section demonstrate that stereo vision is a suitable solution for hazard detection, not only on simulated images, but also when using real images.

## V. Analysis

To draw more general conclusions on the performance of the algorithm, eight datasets were recorded and used to reconstruct stereo DEMs, as well as the resulting hazard maps. Both DEM errors and hazard-mapping errors were computed for all of these datasets. The datasets were recorded at terrain distances of 2.5 m, 3 m, 4 m, 5 m, 6 m, 7 m, and 8 m. At a 1:20 scale, as discussed earlier, this corresponds to 50 m, 60 m, 80 m, 100 m, 120 m, 140 m, and 160 m above the planet's surface during a real scenario (at a 2 m baseline).

Figure 15 reports the DEM errors of the stereo reconstruction. Both mean, as well as median errors, are given. This is due to the fact that very small regions of strong noise may cause strong effects on the mean while the median is affected less by these regions. Therefore, the combination of the mean and the median together allow a more precise assessment of the actual performance of the reconstruction. Moreover, as discussed in Sec. IV, the error distribution is not fully Gaussian. Based on Fig. 15 it can be concluded that the DEMs are reconstructed accurate enough to perform HD. The large standard deviation for 5 m distance is caused by some very large errors in a very small region of the DEM ( $\leq 10$  pixel), most likely due to problems with the lighting conditions of the image (the standard deviation is caused by some very large reconstruction errors in a very small region of the DEM). Overall, it can be seen that the reconstruction improves if the cameras are moved closer to the surface, which is to be expected as depth resolution is a function of the square-root of the distance to the surface (see Eq. (1))

Figure 16 shows the percentages of correctly detected hazards, correctly detected safe sites, undetected hazards, and undetected safe sites. Moreover, the sum of all correct detections is given. It is important to



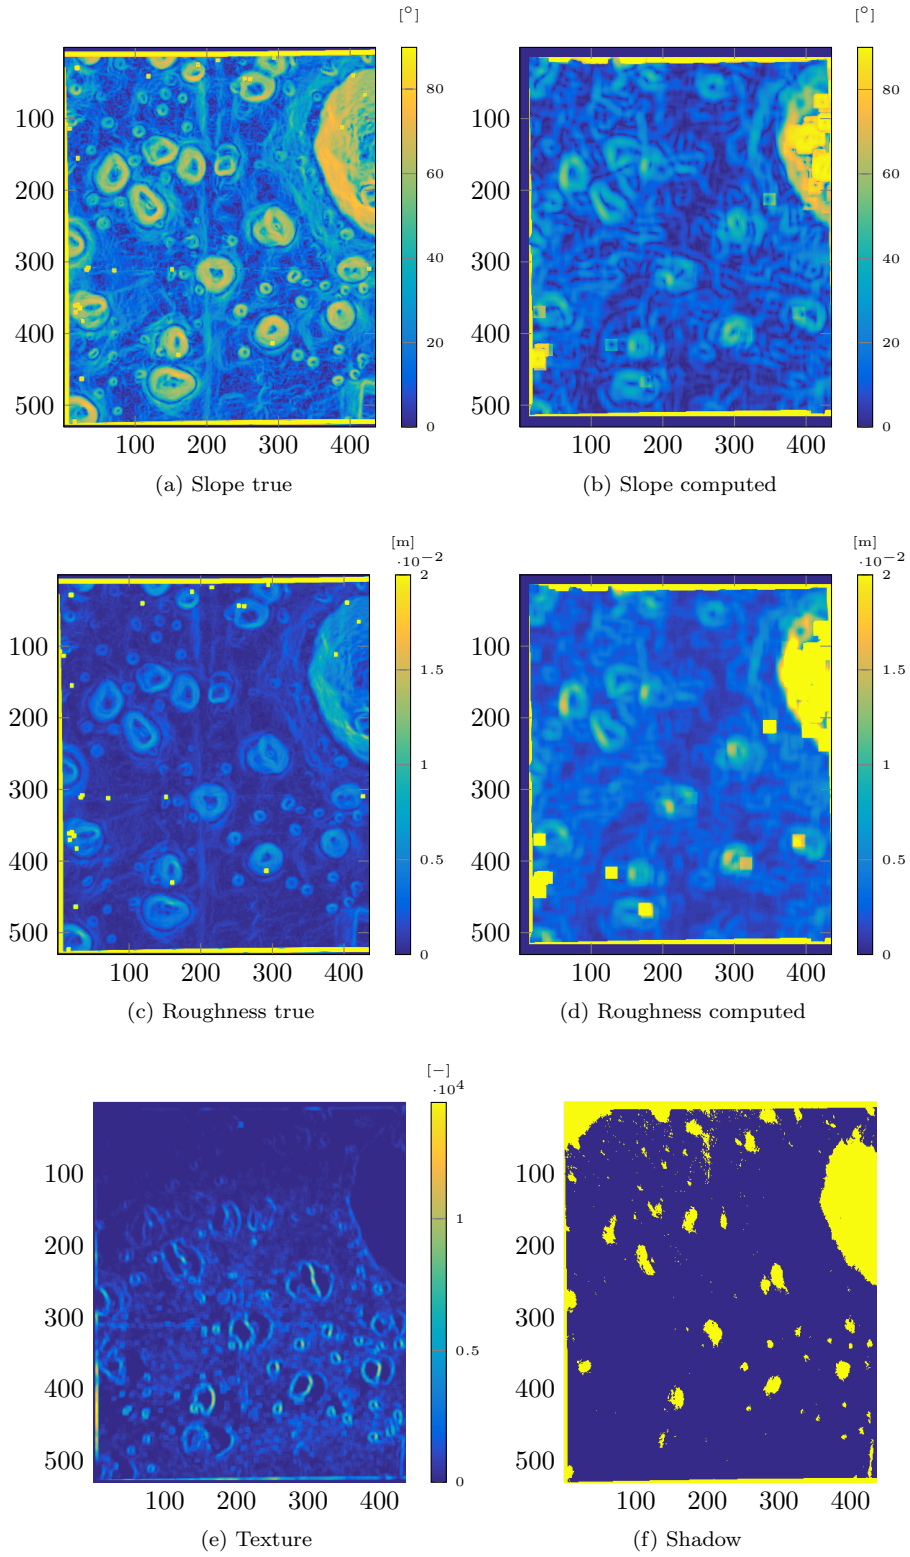


Figure 14: Hazard map contributions (HILT).




Figure 15: DEM mean error and standard deviation (HILT).




Figure 16: DEM mean error and standard deviation (HILT).

note that the there are never more than 2.5% undetected hazards in the scene. By implication this means that the probability of a safe landing based on these hazard maps is always greater than 95.6%, assuming that the function selecting the landing site will only select FN sites proportionally to the size of the DEM/hazard map. However, it is also clearly noticeable that the algorithm does miss a substantial number of safe sites in the scene. This can be seen as a safety factor, as these areas largely correspond to regions that are close to unsafe(e.g., containing large slopes slightly below the threshold) and are close to regions which are actually unsafe. Again, performance does improve when moving closer to the surface, as the DEMs are more accurate at smaller distances.

Lastly, Fig. 17 shows the percentage of all hazards detected. It can be seen that this number is always above 75%, but there is no clear trend in this performance.




Figure 17: DEM mean error and standard deviation (HILT).

Overall, this analysis clearly demonstrates that HDA is possible with the result of a stereo vision surface reconstruction. Therefore, the HILT supports the conclusion from the SILT performed previously, that stereo vision is a feasible candidate for on-board hazard detection.

It is a known fact, that stereo reconstruction becomes challenging - if not impossible - for scenes with little to no texture. During the laboratory experiments at the TRON facility, we were able to reproduce this behaviour on the black wall behind the terrain, as well as on the very smoothly painted ceiling of the laboratory (see Fig. 18), whereas reconstruction on the strongly textured carpet works well. Which demonstrates that the lunar analogue terrain has sufficient texture for conduction stereo based hazard-detection. This supports the assumption that stereo based hazard detection is possible for lunar landings.




Figure 18: Sufficient texture is needed for stereo reconstruction.

It comes as no surprise that performing hazard detection on real images is more challenging than when using artificial imagery, such as those presented in Sec. IV. It is also far more challenging to trade the robustness of noise against the ability to resolve very small features in the scene.

## VI. Conclusions

Based on the hardware-in-the-loop testing (HILT) performed at the TRON facility, it is possible to confirm the conclusions obtained during software-in-the-loop testing (SILT): stereo vision is a capable method for performing on-board hazard detection during planetary descents. A scaled set-up was developed representing a Lunar landing scenario at 1:20 scale. It was shown that the transition from software simulations to simulations including real hardware is a challenging, but feasible path.

We were able to demonstrate that for all test cases considered, successful hazard detection was performed. As expected, the algorithm performs better, the closer the cameras are to the terrain. Both mean and median DEM errors always remain below 10 cm. The algorithm always detects more than 75% of all hazards. Overall, the algorithm tends to overestimate hazards, which serves as a safety factor. These findings support the feasibility of stereo-based hazard detection for planetary landing.

Due to these promising results, this approach will be investigated through further hardware-in-the-loop testing with the long-term vision of offering a system for reliable hazard detection for space exploration missions.

## References

- <sup>1</sup>Woicke, S. and Mooij, E., “A stereo-vision hazard-detection algorithm to increase planetary lander autonomy,” *Acta Astronautica*, 2016, pp. 42–62.
- <sup>2</sup>Neveu, D., Mercier, G., Hamel, J.-F., Simard Bilodeau, V., Woicke, S., Alger, M., and Beaudette, D., “Passive versus active hazard detection and avoidance systems,” *CEAS Space Journal*, Vol. 7, No. 2, 2015, pp. 159–185.
- <sup>3</sup>de Lafontaine, J., Ulitsky, A., Tripp, J. W., Richards, R., Daly, M., and Sallaberger, C., “LAPS: the development of a scanning lidar system with GNC for autonomous hazard avoidance and precision landing,” *Defense and Security*, International Society for Optics and Photonics, 2004, pp. 81–93.
- <sup>4</sup>Johnson, A. E., Klumpp, A. R., Collier, J. B., and Wolf, A. A., “LIDAR-based Hazard Avoidance for Safe Landing on Mars,” *Journal of guidance, control, and dynamics*, Vol. 25, No. 6, 2002, pp. 1091–1099.
- <sup>5</sup>Langley, C., Mukherji, R., Yang, G., Allen, A., and Tripp, J., “Full-scale Static Testing of the Lidar-based Autonomous Planetary Landing System (LAPS),” *Space Exploration Technologies*, Vol. 6960, 2007.

<sup>6</sup>Epp, C. D. and Smith, T. B., “Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT),” *2007 IEEE Aerospace Conference*, March 2007, pp. 1–7.

<sup>7</sup>Woicke, S. and Mooij, E., “Stereo-Vision Algorithm for Hazard Detection during Planetary Landings,” *AIAA Guidance, Navigation, and Control Conference, AIAA-2014-0272*, 2014.

<sup>8</sup>Krüger, H. and Theil, S., “TRON - Hardware-in-the-Loop Test Facility for Lunar Descent and Landing Optical Navigation,” *IFAC-ACA 2010 Automatic Control in Aerospace*, September 2010.

<sup>9</sup>Krüger, H., Theil, S., Sagliano, M., and Hartkopf, S., “On-Ground Testing Optical Navigation System for Exploration Missions,” *9th International ESA Conference on Guidance, Navigation & Control Systems*, June 2014.

<sup>10</sup>Lingenauber, M., Bodenmüller, T., Bartelsen, J., Maass, B., Krüger, H., Paproth, C., Kuß, S., and Suppa, M., “Rapid modeling of high resolution moon-like terrain models for testing of optical localization methods,” *ASTRA 2013 - 12th Symposium on advanced space technologies in robotics and automation*, European Space Agency, June 2013.

<sup>11</sup>Strobl, K. H., Sepp, W., Fuchs, S., Paredes, C., Smisek, M., and Arbter, K., “DLR CalDe and DLR CalLab,” <http://www.robotic.dlr.de/callab/>.