

Robust approximation of image illumination direction in a segmentation-based crater detection algorithm for spacecraft navigation

Bolko Maass¹

Received: 21 March 2016 / Revised: 25 May 2016 / Accepted: 25 May 2016
© CEAS 2016

Abstract This paper describes an efficient and easily implemented algorithmic approach to extracting an approximation to an image's dominant projected illumination direction, based on intermediary results from a segmentation-based crater detection algorithm (CDA), at a computational cost that is negligible in comparison to that of the prior stages of the CDA. Most contemporary CDAs built for spacecraft navigation use this illumination direction as a means of improving performance or even require it to function at all. Deducing the illumination vector from the image alone reduces the reliance on external information such as the accurate knowledge of the spacecraft inertial state, accurate time base and solar system ephemerides. Therefore, a method such as the one described in this paper is a prerequisite for true “Lost in Space” operation of a purely segmentation-based crater detecting and matching method for spacecraft navigation. The proposed method is verified using ray-traced lunar elevation model data, asteroid image data, and in a laboratory setting with a camera in the loop.

Keywords Crater detection · Illumination estimation · Spacecraft navigation

1 Introduction

Impact craters are a surface feature that is shared by large asteroids and moons which lack a thick atmosphere that could cause their quick erosion. For that reason, algorithmic

optical detection of such craters is not only attractive for geological and other scientific purposes, but also in autonomous spacecraft navigation. Here, impact craters are used as readily detectable and highly pose- and illumination-invariant landmarks that are not only persistent within geological time frames, but are also very memory-efficient in terms of map data to be carried.

CDAs used to match detected craters to a catalog of known craters (CDMAs, Crater Detection and Matching Algorithms) can generally be used for navigation purposes in a straightforward manner, as a predictor–corrector step to cancel a state estimator's accumulated error by introducing measurements that are absolute with regard to the global coordinate frame and not relative to the spacecraft's prior state, as is the case with input from inertial measurement units or feature trackers, for example. These applications require a valid prior state estimate to make accurate predictions of the image locations in which the known craters should be seen. In this way, these methods move the task of matching measured craters to the catalog of known craters into the image space, and are, therefore, limited to correcting small affine transformation errors. Such a system is extensively documented by Simard Bilodeau [13].

Moving to the more general case, the so-called “Lost in Space” scenario, there can be no reliance on prior state information. Such a situation may occur in the absence of an external measurement of the spacecraft state, in the case of attitude sensor or state estimator failure, after takeoff of a landed spacecraft, in asteroid navigation scenarios, where the distance to earth is too large to measure position to sufficient accuracy, etc.

The state acquisition problem in scenarios of this type has only been partially solved, for example, by Hanak, where the global lunar position can be fully unknown, but the spacecraft's attitude must be properly controlled to

✉ Bolko Maass
bolko.maass@dlr.de

¹ GNC Systems Department, Institute of Space Systems, German Aerospace Center (DLR), Bremen, Germany

guarantee nadir pointing [3], or by Cheng and Miller, who exploit invariants of coplanar conics to match detected craters to a known catalog from an arbitrary unknown position and attitude, as long as the craters lie on a roughly planar surface [1].

Capabilities of any Lost in Space-capable CDMA are a strict superset of those CDMA's used to correct small estimation errors in the way described above, they perform the same function with less inputs (ideally, just an image and the catalog of known craters). One such input important to many of the CDAs described in Sect. 2 that can no longer be taken for granted is the image's dominant projected illumination direction.

In the above-mentioned Lost in Space scenario and/or in cases, where knowledge of the spacecraft state and the target body's shape, orientation, and position is not sufficient to deduce the illumination direction, e.g., because of unevenly illuminated highly irregular target terrain, a method to compute the illumination direction from the image alone is required.

The main innovation presented in this paper is a method of extracting this direction as simply one more step in the process of crater detection by image segmentation.

This paper is organized as follows: Sect. 2 will introduce a number crater detection (and matching) method categories that have been proposed for spacecraft navigation with a main focus on their reliance on knowledge about the image illumination direction, and in what way, this reliance may prevent them from being used in a Lost in Space scenario.

Section 3 will briefly define the scope of the proposed method for illumination estimation in conjunction with the CDA developed by the author at the German Aerospace Center, in which scenarios it is expected to work, and its known limitations.

The proposed method itself and its derivation from, and interaction with, the CDA will be introduced in Sects. 4 and 5.

Section 6 will elaborate the verification strategy pursued in this paper and which performance measures and test parameters are considered the most relevant to robustly assess the proposed method's performance, and how different application scenarios call for different performance criteria. The actual test data sets and their parameters are summarized in Sect. 7.

The results of the methods application to the test data sets are presented and discussed in Sect. 8. Section 9 closes with a summary.

2 Reliance on known illumination in various categories of state-of-the-art CDAs

For spacecraft navigation specifically, there exist today mainly three more or less mature conceptual approaches

in CDAs, all in part or fully relying on external knowledge about the illumination direction. Methods following the first approach operate under the assumption that visible crater rims produce the most salient structures in the image and use as a base stage an edge extraction scheme, such as the well-known Canny Edge Detector. They assemble the pixels of the resulting image edge response first into segments of consistent curvature representing parts of the crater rim and then group these segments into elliptical approximations of crater candidates to be used as landmark measurement points for navigation. The well-known method by [2] is in this category, employing a number of heuristics, such as inter-edge gradients and curvature estimation to group edge segments into image point sets that can be used to fit candidate ellipses [2]. As described by its authors, this algorithm uses prior knowledge about the global illumination direction to discard edge segments that are almost parallel to the illumination direction and to establish a directional order between edge segments for more robust grouping. The method by Cheng et al. has been picked up more recently by Mammarella et al. [9].

A second conceptual approach to crater detection is segmentation-based, which means that the lowest level image feature extracted is a descriptor of a connected image area corresponding to the illuminated and shadowed areas within the bowl of a visible crater, instead of segments of its rim edges. This means that the image-processing operation used to build these descriptors is a simple connectivity test instead of a potentially sensitive curvature test. Segmentation-based CDAs are, therefore, naturally scale-agnostic, because the descriptor construction operates on the image support and not on the image intensity function, in contrast to the edge segment curvature heuristics that frequently risk missing craters, whose edges are partially non-convex on the edge detector's inherent scale. This category of CDAs is the method on which the proposed illumination estimation approach is based. It is derived from maximally stable extremal regions (MSER, [10]). To obtain a leaner area descriptor, it is simpler than MSER in the sense that this algorithm does not track the extracted image areas' borders to evaluate their rates of change during intensity transition, but employs a roughly equivalent stability measure based on the movement of their centroids. Importantly, however, grouping the extracted stable areas in the method as published still requires prior knowledge about the illumination direction [8].

A method that is mainly segmentation-based but somewhat hybrid with the edge approach has been developed by Spigai et al. who extract connected light and dark patches using an adaptive intensity clustering, but incorporate the image edge response in the process. Grouping one light and one dark image area to form a crater candidate region again requires knowledge of the illumination direction, to establish a directional order of the extracted areas [15].

There also exist fully hybrid methods combining features of the above-described basic approaches in the sense that their methods are used to mutually support each other but are not tightly integrated. The one by Singh and Lim is of this type, performing a binary thresholding to extract the image areas of interest which are then used to group extracted image edges for ellipse fitting [14], or the one by Yu et al. who run a full MSER to extract the areas of interest and use the result to reinforce the grouping decisions for crater rim edge segments that were also extracted [17]. Both methods as described require the illumination direction to be known to pair one dark and one light image area in the correct directional order.

A third category of crater detection methods employs derivatives of the Hough Transform to perform the crater detection, such as an early work on the asteroid navigation case [7] which actually does not incorporate known illumination direction, but performs a costly tensor voting to associate edge segments of consistent curvature and is still subject to the scale sensitivity of edge detection, or the work by Hanak on the actual lunar Lost in Space scenario [3]. The latter method in fact has an option to derive the image illumination direction algorithmically, but uses a method by Varma and Zisserman [16] that employs a statistical approach and makes strong assumptions regarding camera viewpoint, surface texture and albedo, and light source elevation to prevent hard shadows. The method needs to touch all image pixels and is, therefore, quite complex and computationally expensive in comparison to the method that will be suggested in this paper.

Recapitulating, it has been established above that edge-based and/or Hough-based methods may work without external knowledge of the illumination direction, but their performance depends heavily on the scale of crater rim structures and their conformance to the convexity hypothesis that the grouping heuristics are based on. Segmentation-based methods do not impose this specific class of constraints and, therefore, have the potential to detect a whole class of craters that would be rejected by CDAs based upon the modelled ellipticity of the crater rim, which makes them attractive especially in the cases of very low camera elevation with respect to the observed surface, where increasing perspective effects reduces the crater rims' conformance to the mentioned assumptions. The central flaw of segmentation-based CDAs is that because of this same lack of model assumptions, they cannot robustly group their intermediate descriptors, the illuminated and shadowed areas, without external knowledge of the illumination direction.

3 Scope of the proposed method

The method presented in this paper shall be understood in the context of developing a Lost in Space-capable

segmentation-based CDA in a local sense, meaning that in all intended application scenarios, the rough position in large target bodies' reference frames can be expected to be such that for any image taken of the target body's surface, the illumination conditions satisfy the CDA's assumption in principle, but the camera pose knowledge is not good enough to derive an image illumination direction that allows grouping of the detector's intermediate results into proper crater detection candidates.

This of course excludes cases in which the spacecraft is in areas of local night or noon on the target body, where the CDA would not even work with perfect camera pose knowledge, but this holds true for all crater detection methods discussed above, be they edge-based, Hough-based, or hybrid. Fortunately, these cases are less likely specifically in asteroid navigation scenarios for which any given image containing a large portion of the target body is likely to have areas of local morning or evening, where the CDA can provide correct detections.

The definition of correctness of detections and the related but separate problem of false alarms is addressed in Sect. 6, the range of external illumination conditions and the camera pose for which the method performs well, as well as the effects in cases, where these conditions are not met, will be discussed in the results in Sect. 8. For the following explanation of the method, it shall be assumed that imaging conditions are broadly suitable.

4 Building crater candidates from image segmentation

The method described in this paper has been developed in the context of our segmentation-based CDA [8]. Motivating its basic functional principle is easiest done in context with Fig. 1: the task of detection of a crater (left image) is to be accomplished by detecting the two proximal connected areas *A* and *B*, as illustrated in the image on the right of the figure, that cover image points of below and above average intensity, respectively, and are arranged in such a way that *B* is the closest area of similar size and opposite intensity as *A* with respect to the positive illumination direction *s*, indicated by the arrow. In our CDA, the extraction of *A* and *B* from the image is performed by tracking stable connected areas in binary threshold images while gradually raising the intensity threshold, with the stability being measured in terms of the movement of the areas centroids. For formal details of this process, refer to [8].

The result of this operation when applied to an example image of the lunar crater Clavius can be seen in Fig. 2, with the "B" areas shown within boxes in Fig. 2b and the "A" areas shown within boxes in Fig. 2c, along with their respective centroid traces from tracking the areas

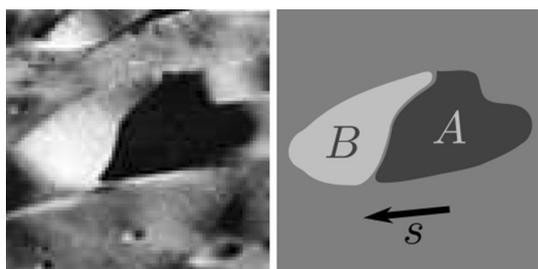


Fig. 1 Shadowed area A and illuminated area B inside a crater after extraction from the image according to [8] in correct arrangement with respect to the directional order imposed by the projected illumination direction s

throughout the increasing (or decreasing, in the case of the illuminated areas) threshold.

At this point in processing, the CDA only has information about the sets of areas as shown in Fig. 2b, c: (1) their size in pixels, (2) their centroid's position in the image, (3) a covariance matrix of their constituent points, and (4) their category (illuminated or shadowed). It specifically has none of the contextual cues that are available to a human observer from Fig. 2a, such as a spatial understanding of the scene (i.e., attitude of the camera with respect to the surface) or the physical configuration of the observed surface itself. Regrettably, areas cannot simply be matched by proximity, as this would risk assigning crater status to hills, rocks or the shadow or reflection behind or in front of the outer ejecta ridges that larger craters show. The missing piece of context that allows to robustly reject such false matches is the illumination direction. The following section describes a method to estimate the illumination direction from the above-described information available to the CDA at this point.

5 Estimating the illumination direction

The central insight of the method proposed in this paper is that an approximation to the required illumination direction information is encoded in the spatial distribution of the extracted illuminated and shadowed image areas. It can be obtained by building a directed graph of adjacent extracted areas and investigating the distribution of the graph's edges' directions in the image.

This process is best motivated by considering a minimal “image” of two craters, one large with a prominent ejecta ridge with four resulting image areas extracted, and one small crater with two areas (cf. Fig. 3). The extracted shadowed areas A_1, \dots, A_3 and illuminated areas B_1, \dots, B_3 are related to each other primarily by size and relative location in the image. If we identify each area with its designator A_i or B_j , which we assume to be located at the respective area's centroid in the image, we can draw edges from each dark area to its closest two neighboring light areas, as illustrated.

The choice of considering exactly two neighboring areas is a model assumption and based on the observation that, in the absence of spurious segmented areas, each light or dark area inside a large crater of the type shown on top in Fig. 3 should have exactly two competing matches of similar scale: the one on the outside of the crater rim, and the one on the inside of the crater bowl. Assuming this, taking two matches into account should guarantee that the correct match is covered. Only considering one match might systematically prefer the incorrect match inside-to-outside, and considering more than two matches would only introduce noise, as the correct match was already captured and any additional matches would be between different craters or other image structures.

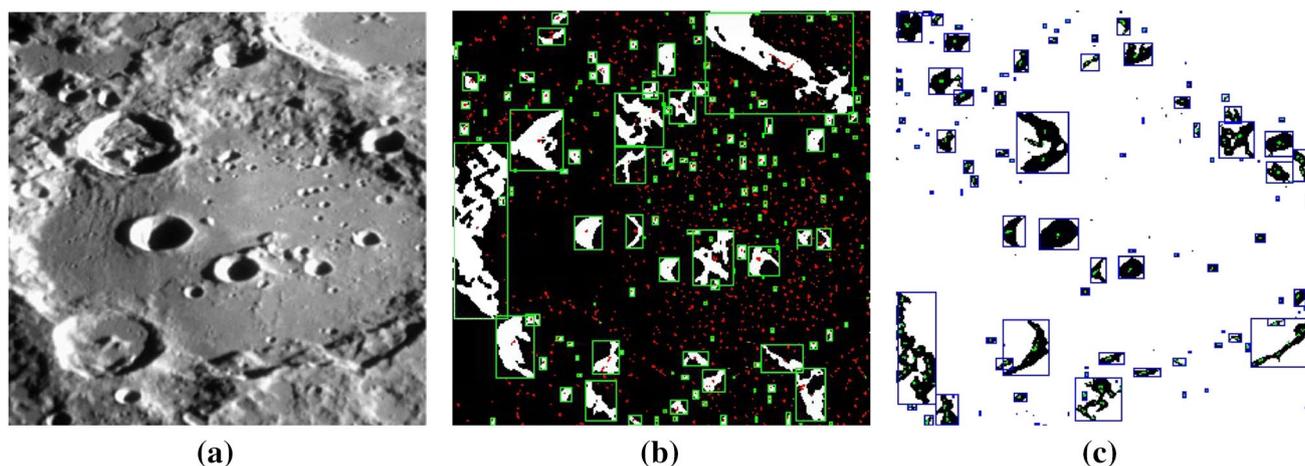


Fig. 2 Extracting illuminated and shadowed areas from an image of the lunar surface. **a** Example image of crater Clavius image: Conrad Jung, Chabot Space and Science Center. **b** Illuminated areas extracted

with centroid traces in red. **c** Shadowed areas extracted with centroid traces in green

large crater, four areas

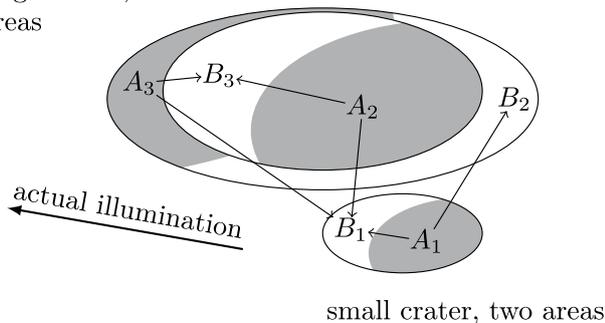


Fig. 3 Different numbers of extracted *contrast areas* depending on crater size and their complete oriented neighborhood graph

In practice, as an additional step, matches that are too dissimilar in size are excluded to make the process more robust against image and terrain noise (a good practical threshold is obtained by disallowing matches, where the larger area is more than ten times the size of the smaller area). Since the areas' centroids are localized precisely in image coordinates, this produces a well-defined directed neighborhood graph in the image space.

This process can be formalized: let n_A be the number of extracted dark areas and n_B be the number of extracted light areas (in the image, both are 3). Let every dark area A_i 's position be identified by its centroid c_i^A for $i = 1, \dots, n_A$ and every light area B_j 's position by its centroid c_j^B for $j = 1, \dots, n_B$. With this, we define the graph's adjacency matrix \mathbf{G} row-wise by its rows (g_{i1}, g_{i2}) , where

$$\begin{aligned} g_{i1} &:= \arg \min_{j \in \{1, \dots, n_B\}} \left\{ \|c_i^A - c_j^B\|_2 \right\} \quad \text{and} \\ g_{i2} &:= \arg \min_{j \in \{1, \dots, n_B\} \setminus g_{i1}} \left\{ \|c_i^A - c_j^B\|_2 \right\}, \end{aligned} \quad (1)$$

giving

$$\mathbf{G} := \begin{bmatrix} g_{11} & g_{12} \\ \vdots & \vdots \\ g_{n_A 1} & g_{n_A 2} \end{bmatrix}. \quad (2)$$

By this construction, the resulting adjacency matrix, for example, in Fig. 3 would be

$$\mathbf{G}_{\text{Fig.3}} = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 3 & 1 \end{bmatrix}. \quad (3)$$

Having defined the neighborhood graph, it is straightforward to compute the angles Θ of all its edges with respect to the image coordinate system's first axis (usually, the image matrix's "down" direction):

$$\Theta := \bigcup_{i=1}^{n_A} \left\{ \text{atan2}(c_{g_{i1}}^B - c_i^A), \text{atan2}(c_{g_{i2}}^B - c_i^A) \right\} \quad (4)$$

Overlaying the oriented histogram of the set of angles Θ resulting from processing the example image 2a in this manner gives Fig. 4a. There are two obvious peaks in the direction set described by the neighborhood graph: one between 250° and 270° , which a human observer would associate with the actual illumination direction, and one between 70° and 90° . This second peak results from the spurious matches of areas at isolated hills or ridges, and the matches of extracted image areas within the craters to their exterior ejecta ridges' reflections and shadows.

The important observation is that because of the above-mentioned various sources of angles around it, this secondary peak in the histogram is robustly wider and less sharply defined than the peak around angles of edges corresponding to the correct matches between illuminated and shadowed areas. To measure and exploit this behaviour, an estimator ϕ operating on the angle set Θ (cf Eq. 4) can be defined using the wrapping angular difference between two angles θ and $\bar{\theta}$

$$\theta \circlearrowleft \bar{\theta} := \begin{cases} \theta - \bar{\theta}, & \text{where } |\theta - \bar{\theta}| \leq \pi \\ -2\pi + \theta - \bar{\theta}, & \text{where } \theta - \bar{\theta} > \pi \\ 2\pi + \theta - \bar{\theta}, & \text{where } \theta - \bar{\theta} < -\pi \end{cases} \quad (5)$$

and, based on this, the wrapping window function of width π

$$w_{\bar{\theta}}(\Theta) := \left\{ \theta \in \Theta \mid -\frac{\pi}{2} \leq (\theta \circlearrowleft \bar{\theta}) \leq \frac{\pi}{2} \right\}, \quad (6)$$

giving an approximate of the wrapping variance $\sigma(\bar{\theta})^2$ as

$$\sigma(\bar{\theta}) := \sqrt{\frac{1}{|w_{\bar{\theta}}(\Theta)|} \sum_{\theta \in w_{\bar{\theta}}(\Theta)} (\theta \circlearrowleft \bar{\theta})^2}. \quad (7)$$

The expression $|w_{\bar{\theta}}(\Theta)|$ denotes the number of angles within the window. The choice of a half circle window around the currently considered angle $\bar{\theta}$ is the model assumption derived from the observation that the histogram peaks occur nearly diametrically opposed and symbolize two "competing" illumination directions. The window divides the histogram bins into sets supporting either one of the directions.

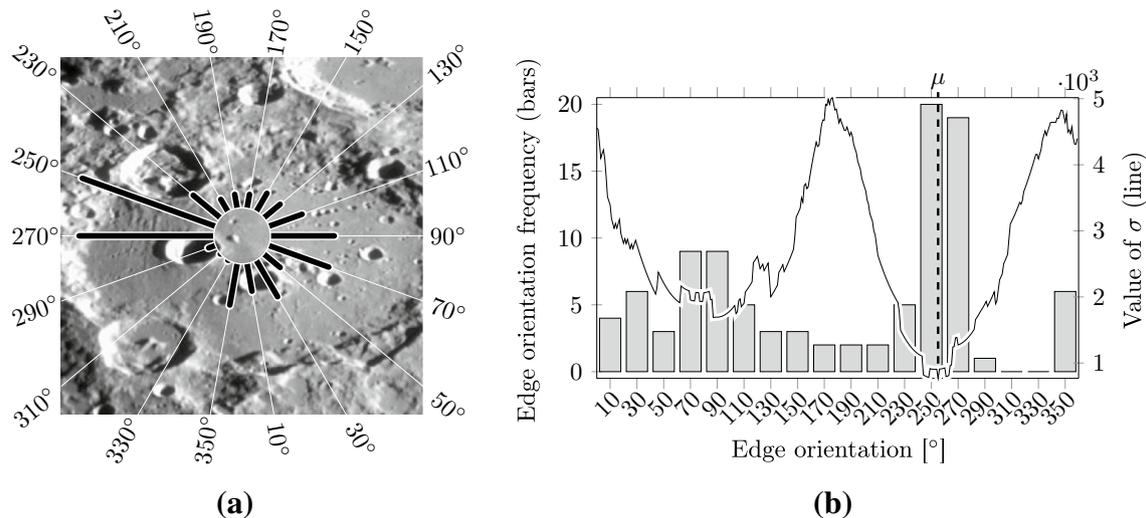


Fig. 4 Neighborhood graph’s edge orientation histogram for the example image and the resulting illumination direction estimate. **a** Image of crater Clavius overlaid with the edge orientation histo-

gram of the neighborhood graph. **b** Unwrapped orientation histogram and values of σ with minimum at the estimated illumination direction

Equation (7) now allows constructing an estimator that is based on the idea of computing the graph edge direction variances for all subsets of range π or less in Θ and choosing the smallest such variance. This estimator shall be called ϕ and be defined as

$$\phi(\Theta) := \arg \min_{\bar{\theta} \in [0, 2\pi)} \{\sigma(\bar{\theta})^2\}. \tag{8}$$

Defined this way, the function ϕ returns from a set of angles Θ the specific angle μ for which the variance within a wrapping window of width π around it is minimal. Processing the example Fig. 2a’s histogram using this function results in a plausible value for μ , as shown in Fig. 4b.

The specific angle μ returned by ϕ is the estimate for the dominant illumination direction in the image. It can be used to robustly perform the grouping of the previously extracted illuminated and shadowed areas A_i and B_j into crater candidates by pruning the neighborhood graph of all edges that deviate from this estimated illumination direction by more than $\sigma(\mu = \phi(\Theta))$.

It is important to emphasize that the illumination direction estimation method presented in this section uses no data that was not previously available from the area descriptor extraction stage of the CDA and simply exploits the observed distribution of edge directions in the neighborhood graph between the area segments of opposite intensity. As for any estimator, there will occur cases that break its underlying assumptions and, therefore, have the algorithm fail to return a useful result. Testing how often this is the case is subject of the following section.

6 Verification strategy

The central object of the following investigation is how well the CDA, supported by the proposed method, is able to recover a sufficient number of craters for matching to a catalog.

6.1 Performance measures

Since the illumination estimation proposed in this paper is meant to support the segmentation-based CDA, its performance must be judged by how well it allows the CDA to group extracted light and dark image areas into crater candidates. Depending on the perspective properties of a scene and the distribution of actual craters in the image, the value for the direction that best allows this grouping may in fact be different from the “true” image-global mean illumination direction, or the irregularity of the target body makes the concept of a global illumination less well-defined (as is often the case with asteroid image data). How well a true direction is recovered is, therefore, secondary and only included for completeness and to give rough overview of the estimator’s performance.

As has been elaborated in Sect. 3, the CDA that is being investigated is meant to work in a Lost in Space situation, and consequently, at least in scenarios with very large target bodies, the available catalog of known craters cannot reasonably be expected to be complete. This, however, means that the crater detection performance can only be measured in true positives: known catalog craters

that also have been detected in the image. In these cases, the notion of false alarm detections becomes largely meaningless, because without visual inspection by a human observer, a false detection (the algorithm considering something that is not a crater to be a crater) cannot be distinguished from the detection of an actual crater that is not in the catalog.

The central success criterion in these cases is, therefore, the question whether enough known craters have been detected to recover a camera pose. To perform pose updates in the context of a navigation system, even a single good detection is helpful, but a number of five matched detections have to be considered a lower threshold for the initial pose acquisition in the Lost in Space scenario, as has also been noted by Hanak [3, p. 186].

For smaller target bodies with nearly complete crater catalogs, detection-to-catalog matching methods may be used that exploit the denseness of the catalog, but at the same time are more sensitive to the number of spurious detections. Therefore, a small sample of asteroid images will be inspected visually to establish an estimate of the false alarm rate of our CDA when used in conjunction with the illumination estimation method proposed in this paper. The inspection will be performed on an image set of the asteroid Eros, where illumination conditions are largely non-uniform. For this purpose, the definition of a false alarm shall cover those cases, where the CDA assigns a crater candidate to an image area that does not show an actual crater or erroneously merges adjacent craters of similar size. Correctly assigned but poorly fitting crater candidates are not counted as false alarms, as they contain useful information and their parameters' errors can be best treated (or rejected) during fusion in a navigation filter.

6.2 Test parameters

The primary verification strategy has the goal of feeding the estimator constructed above with data that covers in a meaningful way the parameter space of the two main external influences on any CDA: (1) the illumination elevation with respect to the observed surface and (2) the elevation of the camera's optical axis with respect to the observed surface. In aggregate over the below described test data sets featuring nearly planar target surfaces, the results of applying the CDA with illumination estimation can then be stratified by these external parameters of illumination and view. In the Eros data set, similar variability in illumination and view are produced naturally because of the irregularity of the target, but cannot be quantified in a global sense for the same reason. Stratification of the results by these parameters is, therefore, not possible.

7 Test data sets

7.1 Lunar artificial data set

An image set that comprehensively covers the parameters mentioned above and also provides accurate ground truth for the estimated quantity, the illumination direction, has to be produced artificially. This has been accomplished using lunar surface elevation data produced from stereo image data captured during the Kaguya mission, cf. Fig. 5a (for an overview on the Kaguya mission, refer to [4]), rendered by a camera simulation software developed within DLR [12] for a virtual camera of 66° horizontal and vertical fields of view at a resolution of 1024×1024 pixels.

From the rectangular elevation model, a circular patch is extracted to guarantee rotational symmetry. The virtual camera is sequentially moved around the terrain in 5° steps (see parameter φ_c in Fig. 5b), and increased in elevation in 5° steps from 5° to 90° elevation with respect to the terrain base plane (parameter ν_c), while always pointing at the terrain model center. For every one of these poses, the virtual sun is moved around the terrain model in steps of 45° (parameter φ_s) and raised in elevation in steps of 5° from 5° to 65° (parameter ν_s). The light is modelled as parallel to simulate an illumination source at infinity.

In this way, 134784 distinct images are created that can be used to test the illumination direction estimator. Figure 5c shows a small sample. For all but the lowest camera elevations, the terrain can be considered flat, and the actual illumination vector is known for every image. Therefore, its projection onto the elevation model base plane and then into the image plane is a good value to use as ground truth for the illumination direction to be estimated.

7.2 Lunar laboratory data set

To confirm the results obtained from processing the artificial image set in a more realistic trajectory scenario, a second test campaign with an actual camera in the loop is performed in the laboratory Testbed for Robotic Optical Navigation (TRON) at the German aerospace center's Institute of Space Systems in Bremen. Here, a physical 1:125000 scale model of the elevation data shown in Fig. 5a is available, together with a KUKA industrial robot that can move a camera on realistic scaled lunar trajectories while illuminated by a movable stage lamp mechanism, cf. [5, 6] and Fig. 6a.

The test trajectory moves the camera seven times over the terrain, modelling a realistic lunar descent orbit trajectory. During each overflight 206 images of size 1024×1024 pixels are taken for a total of 1442 images, while care is taken that the camera's 66° field of view touches the

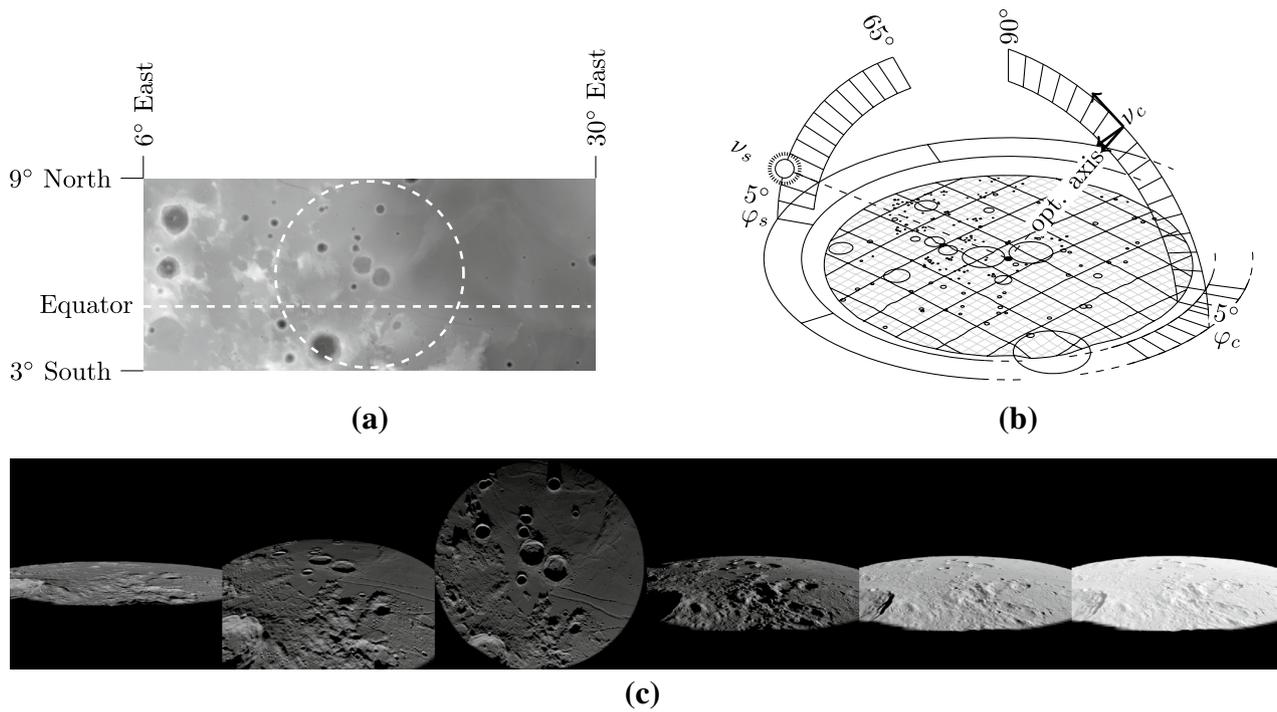


Fig. 5 Creating simulated images of lunar terrain by equidistant sampling of azimuth and elevation of a virtual camera and light source. **a** Source terrain elevation model with extracted terrain circular patch. **b**

Covering the pose space of the camera and the light source. **c** Sample of resulting image set: 3 images with increasing camera elevation and 3 images with increasing sun elevation

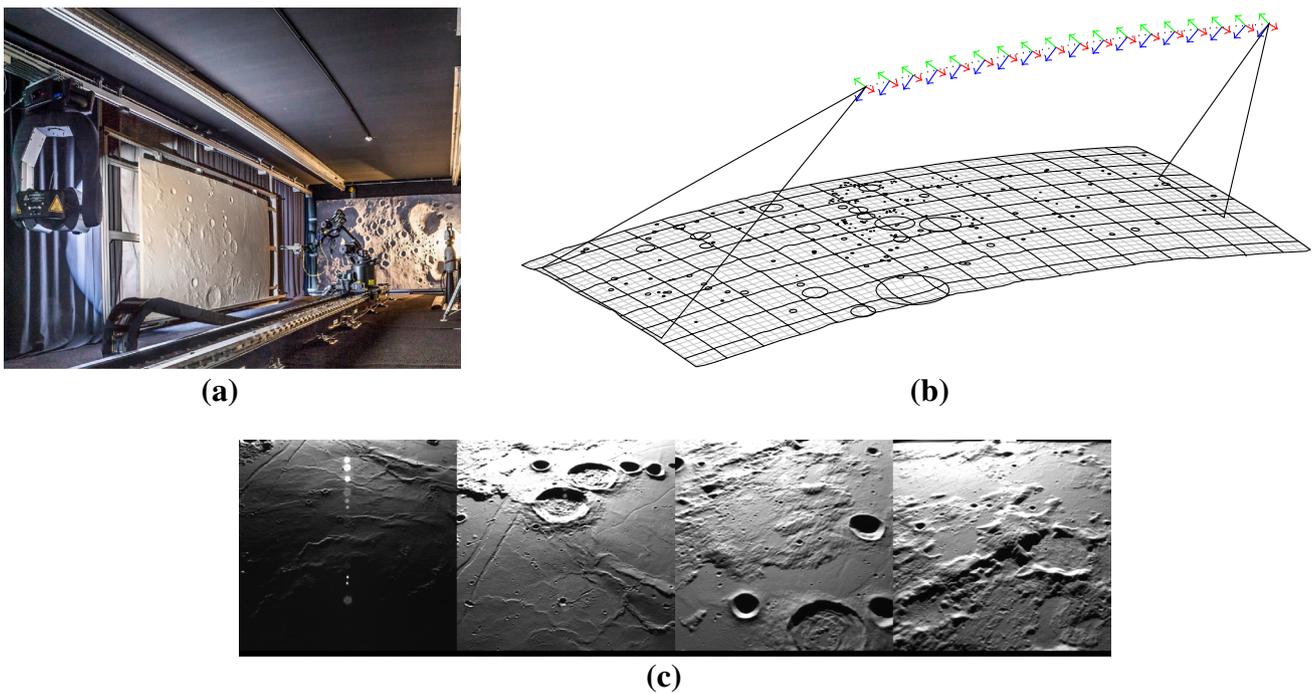


Fig. 6 Taking realistic descent orbit images in TRON. **a** Testbed for Robotic Optical Navigation, TRON. **b** Test trajectory over the terrain model with first and last horizontal field of view boundaries. **c** Sample of resulting image set

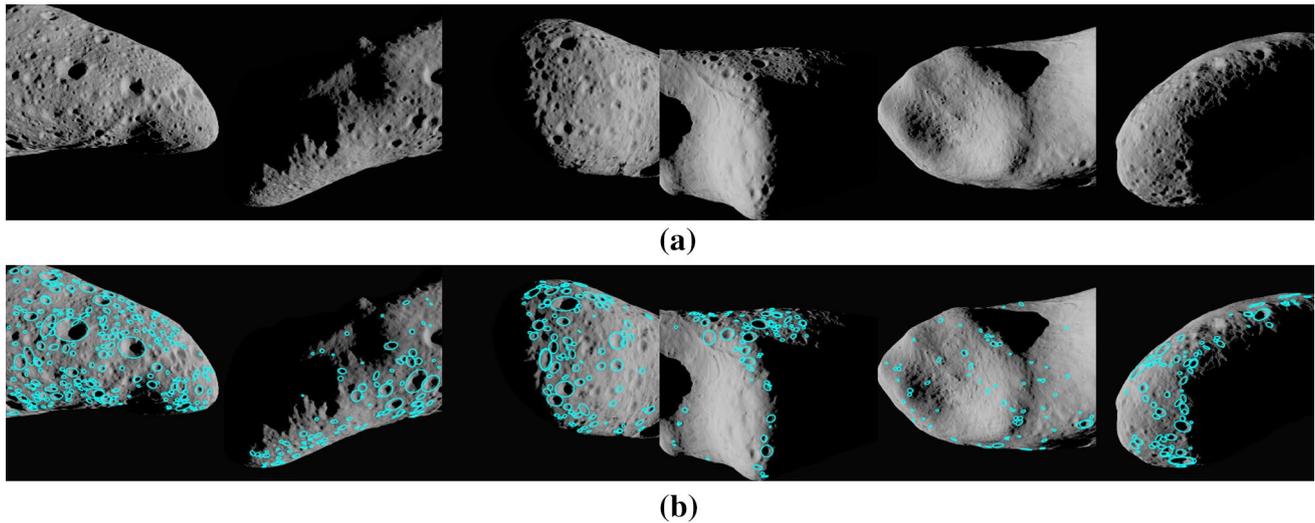


Fig. 7 Crater detection performance on Eros image data set. **a** Eros simulated flyby image set random sample. **b** Crater detection results for image sample

terrain model's boundaries as little as possible. The camera's elevation with respect to the local tangent plane is held constant at 45° . An illustration of the configuration of terrain model, camera coordinate system, and the first and last boundaries of the camera's field of view are shown in Fig. 6b.

Between subsequent overflights, the illumination elevation with respect to the surface is increased, providing a wide range of illumination elevation from approximately 2.5° to 60° . The illumination always comes from straight ahead of the camera, making the ground truth illumination direction in all images 0° . A sample of the resulting images can be seen in Fig. 6c, highlighting some of the additional visual effects introduced by the laboratory setup, such as lens artifacts and a different saturation behaviour. Because there is significant downrange coverage of the camera field of view's footprint on the terrain, and the light source is not at practical infinity, the illumination elevation as understood in the artificial image tests is no longer well-defined for this image set. We, therefore, only consider the overall performance and no longer stratify by illumination elevation.

To obtain a performance indicator with respect to the crater candidates formed using the estimated illumination direction, a ground truth crater catalog of 1000 elements is used for the model. This means that an estimated 20 % of all surface craters visible in the image data set are known and could actually be matched to the detections by a later stage algorithm. This again means that spurious crater detections cannot be distinguished from good detections of unknown craters, which is why only the true positive rate will be evaluated.

7.3 Asteroid artificial data set

A final test is performed on an artificial image set of the near-earth asteroid 433 Eros created by rendering an STL model obtained from NASA [11]. In total, 3400 images are taken that simulate medium-distance flybys over the length of Eros with constant external illumination in the image horizontal to optical axis plane, resulting in varying illumination conditions on Eros' surface. A ground truth illumination direction with a planar reference would be 270° for images 1–1700 and 90° for images 1701–3400.

For an example of the resulting diversity in images, see Fig. 7a. This data set will be used to draw a sample for visual inspection, to establish an estimate of the true false alarm rate of the crater detector.

8 Results

8.1 Lunar artificial image set

Evaluation of the artificial image set gives the results shown in Fig. 8. The evaluation of the strict correctness of the estimate, defined as having the ground truth illumination direction within $\mu \pm \sigma$, is shown in Fig. 8a. As already mentioned above, at the camera elevations used in some of the artificial scenes, the perspective effects become significant. For planar surfaces, these effects may average out over the whole image, but in cases where the segmented image areas are concentrated in one part of the image, the resulting illumination direction estimate will be skewed accordingly. However, since this estimate reflects the actual

predominant edge directions in the image area neighborhood graph, the resulting pairings of areas will generate correct crater candidates.

The results concerning this quality measure, reflecting the success rate in correct crater candidate creation, is shown in Fig. 8b. The choice of the minimum number of correct detections was explained in Sect. 6

It can be seen in Fig. 8a that even the strict ground truth comparison breaks down only for very low camera

elevations and very high illumination elevations. There is an extensive area of these parameters where the estimator has a high chance of capturing the ground truth illumination direction within its natural uncertainty measure σ . The application-specific result, as shown in Fig. 8b, is even stronger. Correct detection of the necessary number of craters for position estimation is almost certain for all but the lowest camera elevation of 5° and illumination elevations above 55° .

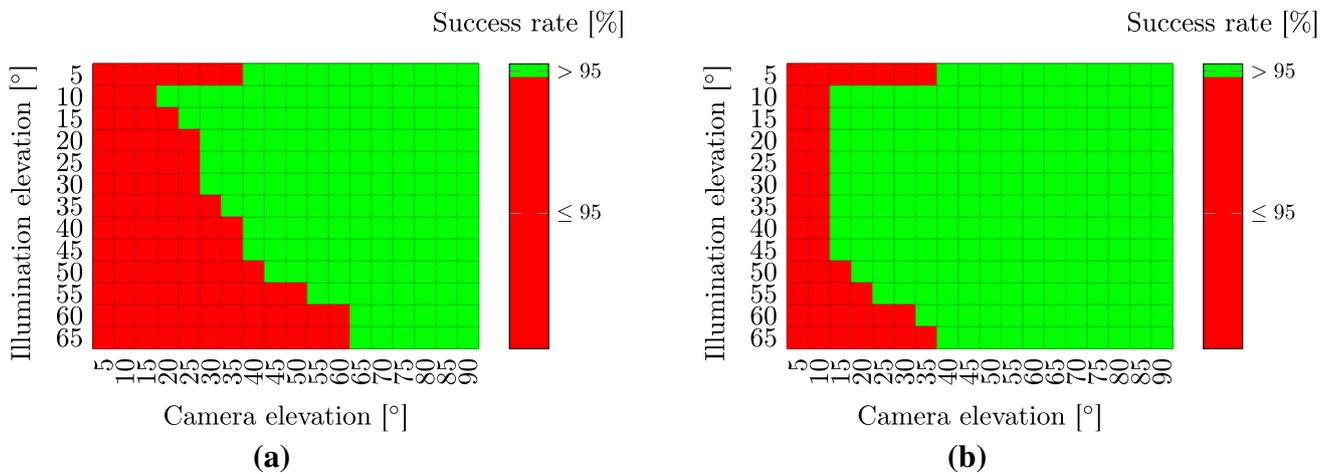


Fig. 8 Method performance with respect to ground truth illumination estimation quality and crater detection success for the artificial data set, stratified by input illumination elevation and input camera elevation. **a**

Fractions of test cases with ground truth illumination direction within estimated illumination direction $\mu + \sigma$. **b** Fractions of test cases with correct detection of at least 5 known craters

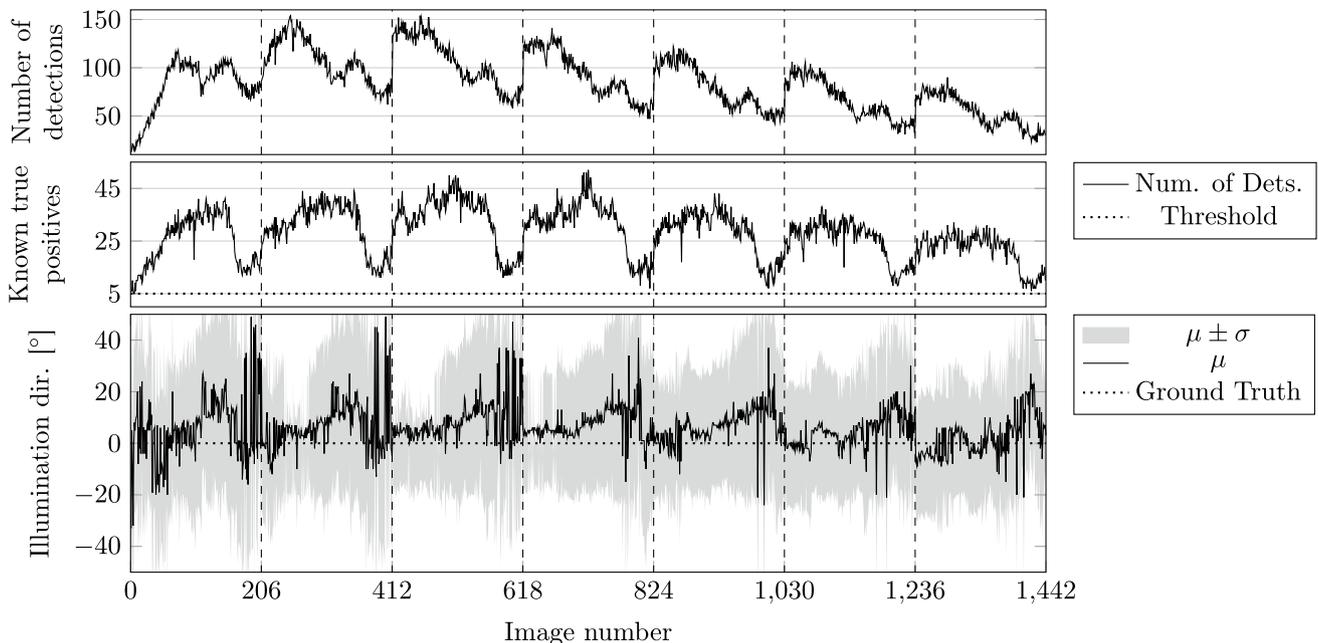


Fig. 9 Test on the TRON data set with crater detection results and number of true positives (known catalog craters detected) for the estimated illumination direction μ and σ . Ground truth illumination direction at 0° , separate terrain overflights delimited at *dashed lines*

8.2 Lunar laboratory image set

The bottom graph of Fig. 9 shows that there is robust estimation of the ground truth illumination direction of constantly 0° . Only in 11 out of 1442 attempts, this true illumination direction lies outside the $\mu \pm \sigma$ area. Those cases do not, however, indicate a breakdown of the estimator, but were symptom of the systematic behaviour originating in the perspective effects of the scene: it can be seen in the top graph of Fig. 9 that the number of crater detections never drops to zero. The detection number also does not strongly correlate with the illumination direction estimate, as long as the $\mu \pm \sigma$ envelope is wide enough to include the image area graph edges that represent the true matches. The true positive rate in the middle graph reflects the local density variation of the known catalog and is included to give an estimate of the number of “true” matches that a matching algorithm can expect to work with. The true positive rates realized here are sufficient to compute an initial camera pose on average within two images, using a Lost in Space matching algorithm based upon [3] that has been extended at DLR for non-nadir camera-scene configurations.

It can be observed that towards the end of each trajectory over the terrain model (separated by vertical dashed lines in Fig. 9), the illumination estimate is significantly less stable than in the beginning. This is an artifact of the decreasing distance between camera and light source. Because of decreasing distance to the lamp when moving towards the end of the terrain model, the non-parallelity of the illumination and in consequence the actual shadow direction on the terrain has a much larger variance than at the near end of the terrain model. If the extracted crater interior areas are not uniformly distributed in the image and some orientations in the areas’ neighborhood graph become overrepresented, the large observed jumps between the estimate of the illumination direction occur. This is not a concern for realistic lunar scenarios, as the illumination in space is nearly parallel. The relatively low (but still sufficient) number of detections of craters in these parts of the trajectory is also a result in the very high illumination elevation because of the proximity to the laboratory lamp.

The perspective effect mentioned above also explains the local bias toward the positive direction estimates that can be seen in the graph: for large parts of the terrain, there is a preference within the CDA for areas extracted predominantly in the right part of the image. This naturally shifts the mean direction of their connecting lines to the positive values observed. The later overflights with an overall higher illumination elevation see a shift of the mean estimate towards back the ground truth and a lower uncertainty overall. In any case, the wide swath of admissible image area graph edge angles indicated by $\mu \pm \sigma$ makes sure that crater candidates on the far side of the image are still considered admissible and are correctly formed.

8.3 Asteroid image set

Figure 7b shows typical crater detection results for the image set shown in Fig. 7a. The whole data set of 3400 images shows that the ground truth illumination direction of 270° for images 1–1700 and 90° for images 1701–3400 is captured by the illumination estimator’s output of $\mu \pm \sigma$ in 94.35 % of cases. Statistics for the results of the visual inspection for false alarm detections in a random sample of images are given in Table 1.

The sample of 1 % of the 3400 images has been chosen uniformly at random. Neglecting possible correlation of the

Table 1 Crater detection quality by visual inspection for a random sample of Eros images

Image #	Number of detections	Number of false alarms	False alarm percentage
73	60	4	6.67
344	128	2	1.56
411	67	2	2.99
504	32	3	9.38
630	150	5	3.33
704	48	3	6.25
848	266	10	3.76
1047	233	9	4.04
1142	239	10	4.18
1288	82	2	2.44
1412	113	5	4.42
1484	71	2	2.82
1565	115	3	2.61
1608	80	2	2.50
1622	175	7	4.00
1654	191	13	6.81
1686	25	2	8.00
1892	41	2	4.88
1981	60	3	5.00
2160	60	6	9.09
2173	57	9	15.78
2290	23	1	4.35
2422	27	1	3.70
2465	106	8	7.55
2581	56	1	1.79
2593	33	1	3.03
2778	70	3	4.29
2793	56	3	5.36
2878	67	2	2.99
2986	74	2	2.70
3124	83	2	2.41
3233	123	7	5.69
3346	104	8	7.69
3390	45	0	0.00

crater detection results for overlapping images, bootstrapping the observed imagewise false alarm rates in Table 1 suggests a 95 % confidence interval for the false alarm rate as (3.86, 5.91 %).

9 Conclusion

In this paper, an estimator for the effective illumination direction in an image of a cratered surface has been introduced that uses only data already available from the previous processing stages of a segmentation-based crater detection algorithm. Its purpose is to allow the CDA to robustly group its intermediate descriptors, the segmented image areas belonging to the shaded and lit interior parts of a crater, into proper crater detection candidates. The performance overhead of applying this method is very well tolerable. In the straightforward C implementation of the CDA, the additional computational cost of building the neighborhood graph in relation to the cost of the segmentation stage before it is less than 5 % for all tested images. The computational cost of building and evaluating the edge direction histogram is negligible. Of course, these ratios might change when performing the segmentation in specialized hardware.

In the context of the CDA from which it is derived, the estimator can be considered robust and accurate enough to fully replace external knowledge about the illumination direction the grouping stage when making sure a certain envelope of surface illumination parameters is not left. This is a necessary step towards usability of purely segmentation-based CDAs in Lost in Space scenarios, where such external knowledge should not be presumed available.

References

- Cheng, Y., Miller, J.: Autonomous landmark based spacecraft navigation system. *Adv. Astron. Sci.* **114**(3), 1769–1783 (2003)
- Cheng, Y., Johnson, A.E., Matthies, L.H., Olson, C.F.: Optical landmark detection for spacecraft navigation. *Adv. Astron. Sci.* **114**(3), 1785–1803 (2003)
- Hanak, F.C.: Lost in low lunar orbit crater pattern detection and identification (Doctoral dissertation). University of Texas at Austin, Texas, United States of America (2009)
- Kato, M., Sasaki, S., Takizawa, Y., et al.: The Kaguya mission overview. *Space Sci. Rev.* **154**(1–4), 3–19 (2010)
- Krüger, H., Theil, S.: Tron-hardware-in-the-loop test facility for lunar descent and landing optical navigation. In: 18th IFAC symposium on automatic control in aerospace (2010)
- Krüger, H., Theil, S., Sagliano, M., Hartkopf, S.: On-ground testing optical navigation systems for exploration missions. In: 9th international ESA conference on guidance, navigation and control systems (2014)
- Leroy, B., Medioni, G., Johnson, E., Matthies, L.: Crater detection for autonomous landing on asteroids. *Image Vis. Comput.* **19**(11), 787–792 (2001)
- Maass, B., Krüger, H., Theil, S.: An edge-free, scale-, pose- and illumination-invariant approach to crater detection for spacecraft navigation. In: Image and signal processing and analysis (ISPA), 2011 7th international symposium on, IEEE, pp. 603–608 (2011)
- Mammarella, M., Rodríguez, M.A., Pizzichini, A., Montero, A.M.S.: Advanced optical terrain absolute navigation for pinpoint lunar landing. *Advances in Aerospace Guidance, Navigation and Control*, pp. 419–430. Springer, Berlin (2011)
- Matas, J., Chum, O., Urban, M., Pajdla, T.: Robust wide-baseline stereo from maximally stable extremal regions. *Image Vis. Comput.* **22**(10), 761–767 (2004)
- NASA (2014) Eros printable STL model. <http://nasa3d.arc.nasa.gov/detail/eros>. Accessed 4 Dec 2016
- Paproth, C., Schlüssler, E., Scherbaum, P., Börner, A.: Sensor++: simulation of remote sensing systems from visible to thermal infrared. In: International archives of the photogrammetry, remote sensing and spatial information sciences, XXII ISPRS Congress, pp. 257–260 (2012)
- Simard Bilodeau, V.: Autonomous vision-based terrain-relative navigation for planetary exploration (Doctoral Dissertation). Université de Sherbrooke, Québec, Canada (2015)
- Singh, L., Lim, S.: On lunar on-orbit vision-based navigation: terrain mapping, feature tracking driven ekf. In: Proceedings of the AIAA guidance, navigation and control conference (2008)
- Spigai, M., Clerc, S., Simard Bilodeau, V.: An image segmentation-based crater detection and identification algorithm for planetary navigation. *Intelligent Autonomous Vehicles*, Lecce, Italy (2010)
- Varma, M., Zisserman, A.: Estimating illumination direction from textured images. In: Computer vision and pattern recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE computer society conference on, IEEE, vol. 1, pp. 1–179 (2004)
- Yu, M., Cui, H., Tian, Y.: A new approach based on crater detection and matching for visual navigation in planetary landing. *Adv. Space Res.* **53**(12), 1810–1821 (2014)