



Performance of widely used feature detection techniques for co-registration of planetary images

Antonia Schriever and Klaus Gwinner

German Aerospace Center (DLR), Institute of Planetary Research, Planetary Geodesy, Rutherfordstr. 2, 12489 Berlin, Germany

Introduction: Image co-registration, which is an important preprocessing step in planetary remote sensing for various scientific use-cases, is the alignment of images of the same scene acquired at a different time, from different viewpoints or using different sensors. Automated techniques for image alignment can be based on correlating points in the images. This creates the need for accurate and precise feature detection and matching.

In feature detection, first distinctive keypoints within an image, such as corners, lines, edges or blobs are identified. Many methods are generating a feature descriptor for each keypoint, which forms a numerical representation of the keypoint that describes the local image characteristics at the point. Finally, those described keypoints are matched to create a correspondence between homologous local features in a second image.

While recent feature detection and matching methods have been widely tested on terrestrial images, we aim to test their performance for planetary images. In the context of this work, the image data comes from the High Resolution Stereo Camera (HRSC) on ESA's MarsExpress mission [1]. The images are orthorectified with the global DEM from Mars Orbiter Laser Altimeter (MOLA, [2]) data. This ensures a fairly good alignment between the images of the HRSC Nadir, Stereo and Photometry images while still typically including pixel offsets of a few pixels.

Methods: The scope of our tests includes the FAST, SUSAN, Harris, ORB and BRISK Corner detectors, the Canny Edge detector, KAZE, A-KAZE and SIFT (Blob detectors), and the LSD Line Segment Detector. An overview of these methods including original references can be found in [3]. After the feature detection step, the detectors FAST, SUSAN, Harris, Canny and LSD are combined with a SIFT descriptor for feature matching. The detectors KAZE, A-KAZE, SIFT and BRISK are evaluated individually alongside their respective descriptors. Additionally, ORB is created through the combination of FAST with BRIEF, which is solely a descriptor that needs identified keypoints as input.

Finally, descriptor matching uses the nearest neighbor-based FLANN (Fast Library for Approximate Nearest Neighbors) matcher. To ensure robust matching, a Nearest-Neighbor Distance-Ratio is set, where a match is accepted only if the distance ratio between nearest and second nearest neighbor is below the threshold ratio [4]. Furthermore, matches with a pixel distance exceeding a threshold determined by the standard deviation of distances among all matches are eliminated. For comparison purposes, a well-established least-squares matching (LSM) technique which is also applied for HRSC stereo reconstruction [5] is applied, as well. It uses approximate keypoint

coordinates as input which can either be generated internally using cross-correlation (dense matching method) or can come from external feature detectors.

Results: First, we tested invariance in rotation, scale, brightness and noise by determining the proportion of matched keypoints for modified and original images. Different detectors performed best in different invariance tests: SIFT, KAZE and AKAZE (LSM small angles) for rotation, SIFT for scale, FAST for brightness (LSM not tested) and LSM, KAZE and AKAZE for noise.

Second, we evaluated the runtime of the different feature detectors, which is a relevant property in time-critical applications such as navigation. The computation time needed for a single match (detection, description and matching) is lowest for FAST and Harris, while KAZE and AKAZE are more than 300% slower than all other detectors.

Lastly, to evaluate point precision and accuracy and the number of matches, we created two datasets with 15 images each, that contain either only feature-rich or feature-poor images (Fig. 1). The feature-rich images show well developed features and texture, while the feature-poor images show weak textures and little to no distinctive features. Image entropy and contrast serve as quantitative measure for the texture quality. As expected, they show lower values for the feature-poor set. The evaluation for both image sets showed higher number of matched keypoints and lower number of completely unmatched image pairs in the feature-rich set. ORB and LSM are the only methods that allow to match all image pairs of the feature-rich set. KAZE, only topped by LSM, finds the most matches in both sets and also has the lowest number of unmatched pairs in the feature-poor set.

Precision is measured as the standard deviation of the co-ordinate differences between the matched points in the two images, as the images are already closely aligned through ortho-rectification. Among the feature detectors, Harris showed the most precise results, but also has the highest number of unmatched image pairs, while SIFT, and FAST also show better values than the remaining methods.. However, the results of LSM were by a factor of higher than two more precise than those of each of the feature detectors. Moreover, notably, none of the feature detectors did achieve sub-pixel precision for these images. In the feature-poor group, similar results were achieved, but low number of matches prevented the derivation of reliable precision values for some of the feature matchers.

The results described above have been complemented by runs performing testing of matched points by using the detected keypoints also as input to LSM. In most cases, the precision of the matches could be further refined by the LSM step. However, for many of the detected keypoints no LSM matches were produced. This demonstrates the complementarity between feature-based and dense matching approaches, and limits the use of LSM for checking individual feature-based matches.

Since we found that some of tested methods can provide accurate results on planetary images but were not able to detect any features in some of the images, usage of different methods seems to be advisable, as a more general conclusion. We will also discuss our results obtained for HRSC data in comparison to those obtained for other planetary image data sources (e.g. HiRISE and CTX).

References:

- [1] Jaumann, R., et al. PSS 55 (7-8), 928-952 (2007) DOI: 10.1016/j.pss.2006.12.003.
- [2] Smith, D.E., et al. JGR 106 (E10), 23689-23722 (2001), DOI: 10.1029/2000JE001364.
- [3] Misra, I. et al. IJRS 43:12, 4477-4516 (2022), DOI:10.1080/01431161.2022.2114112
- [4] Lowe, D.G. et al. IJCV 60, 91-110 (2004), DOI: 10.1023/B:VISI.0000029664.99615.94
- [5] Gwinner, K., et al. PE&RS 75(9), 1127-1142 (2009), DOI 10.14358/PERS.75.9.1127

Figure 1: Example feature-rich (left) and feature-poor (right) image.

