

# ASSESSMENT OF FEATURE DETECTORS AND DESCRIPTORS IN REMOTE IMAGES OF PLANETARY BODIES

E. J. Speyerer<sup>1,\*</sup>

<sup>1</sup> School of Earth and Space Exploration, Arizona State University, Tempe, AZ - [espeyer@asu.edu](mailto:espeyer@asu.edu)

Commission III, ICWG III/H

**KEY WORDS:** Feature, Detectors, Descriptors, Matching, Registration

## ABSTRACT:

Algorithms to detect, describe, and match common features in image sets are expanding to new worlds with the integration of the OpenCV feature matching framework in a popular mapping tool for planetary images (ISIS3). These algorithms provide a new approach to register images and build image-based control networks. However, the natural landscape of the Moon and other planetary bodies pose new challenges such as numerous features that appear similar (i.e., impact craters and boulders). In addition, planetary image sets vary in scale, orientation, and noise properties, especially when conducting cross instrument and cross mission comparisons. This study assesses a collection of common feature detector and descriptor algorithms to examine how they adapt to these challenges. With our analysis, we did not identify an ideal detector and descriptor combination that exists for our diverse lunar dataset. However, we did identify where particular algorithms succeed and identify their shortcomings. By knowing these capabilities, users can identify the proper set of algorithms to apply to an image set given the presence of noise and variations to scale and orientation.

## 1. INTRODUCTION

### 1.1 Background

The terrestrial remote sensing community adopted feature detection, description, and matching techniques from the machine vision community to identify and compare similar landforms and other distinctive attributes in vast image sets. Some applications of these tools include image registration, object localization, and 3D terrain reconstruction. With the recent addition of feature detection and matching routines in Integrated Software for Imagers and Spectrometers 3 (ISIS3; Anderson et al. 2004), a popular image processing tool for remotely sensed observations from planetary missions, feature-based matching is expanding to new worlds. However, these new planetary bodies pose challenges for some feature detection and matching routines that were originally derived for terrestrial applications. Countless non-unique and repetitive surface features (e.g. impact craters with similar appearance, boulder fields, etc.) cover many planetary bodies such as the Earth's Moon. This study provides insight into the effectiveness of various feature detectors and descriptors on images acquired of planetary surfaces under various lighting and viewing geometries.

### 1.2 Objectives

The objective of this study is to analyze how feature detection, description, and matching techniques handle planetary images that lack hard corners seen in terrestrial images (i.e., roof lines and/or road intersections) and contain many features that look similar (i.e. crater rims and boulder fields). By comparing the performance of each algorithm and identifying optimal routines for a given type of image, one can increase the success of registering image pairs, locating objects/features in a scene, and/or reconstructing the scene using multiple images.

Since the inclusion of feature matching into ISIS3, Becker and co-workers (2016) used feature matching routines to create a global control network of the planet Mercury. In doing so, they controlled 100,432 images with sub pixel accuracy. Using a

FASTX feature detector and SIFT descriptor, and a brute force matcher they derived 12,596,336 control points and 94,745,475 tie point measurements. From the global control point cloud, Becker *et al.* (2016) created a 64 pixel/degree digital elevation model of the planet.

In addition, Speyerer et al. (2016) used feature matching to calibrate the Clementine UVVIS camera system. In this study, they identified common features in UVVIS and LROC WAC observations acquired under similar lighting and viewing geometries. Using these control points, they calculated the actual focal length for each UVVIS band, the optical distortion, and the orientation of the camera during each observation. By registering the two datasets, comparative studies can be carried out using UVVIS images and datasets tied to the LRO coordinate system, which is the generally accepted system for reporting and describing lunar features.

This study investigates the performance and robustness of common feature detectors and descriptors on optical images as well as examines how illumination of the terrain effects the functionality. By understanding the capabilities and limits of various feature detectors and descriptors, users can improve their selection for studies that use images of planetary bodies. Eventually, feature matching routines could be used to create a global control mosaic of all the images acquired of a planetary body simplifying cross instrument and cross mission studies.

## 2. MATERIALS AND METHODS

### 2.1 Image Set

For this analysis, we used images acquired by the Lunar Reconnaissance Orbiter. The mission launched in June of 2009 and carries three science cameras that encompass the Lunar Reconnaissance Orbiter Camera payload (a single Wide Angle Camera and twin Narrow Angle Cameras; WAC and NACs). Each NAC incorporates a Ritchey-Chretien telescope with an

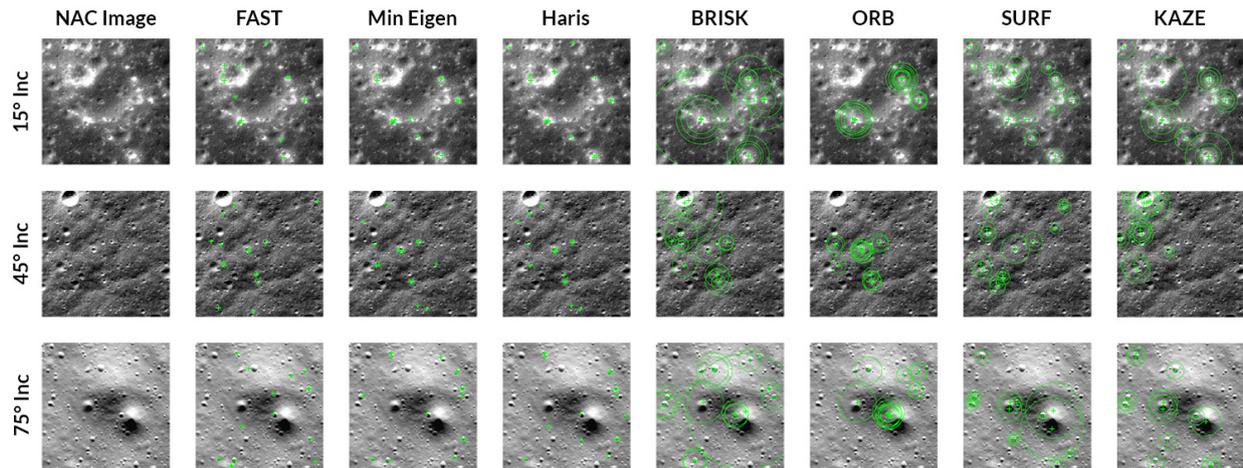


Figure 1. Application of seven feature detectors on three 256x256 pixel sub-images of LROC NAC observations with incidence angles of 15° (M109753063L), 45° (M139396321R), and 75° (M1285183949L).

effective focal length of 700 mm and a line array CCD with an instantaneous field of view of 10 micro-radians (Robinson et al., 2010). With this design, the NACs have a pixel scale of 0.5 m from an altitude of 50 km, which was the altitude during the primary mission. Since launch, the NACs have acquired over 1.6 million images of the lunar surface over a broad range of lighting and viewing geometries. Prior to use in this study, we calibrated each LROC NAC image using the standard processing procedures outlined in Humm et al. (2016), which decompresses the image to the original 12-bit space and translates the raw digital number (DN) values into reflectance (I/F).

## 2.2 Feature Detectors and Descriptors

Feature registration and matching can be segmented into three parts: feature detection, feature description, and matching. In feature detection, a set of algorithms are applied to an image in order to identify “interesting” points or regions of an image. Among feature detectors, there are two general types we investigated: corner and blob. As the names suggest, corner detectors identify the corners or the intersections of two edges in an image and tags them as interest points. In terrestrial images, corner detectors are useful for tagging the corner of man-made structures, such as buildings and road intersections. blob detectors aim to detect unique regions in an image that have different properties than the remaining portion of the image, such as brightness. In this study, we evaluated several common feature detectors (see Table 1), including: Features from Accelerated Segment Test (FAST) (Rosten and Drummond, 2005; 2006), Minimum Eigenvalue Algorithm (Shi and Tomasi, 1994), Harris (Harris, 1988), Speeded-Up Robust Features (SURF) (Bay et al., 2008), KAZE (Alcantarilla et al., 2012), Binary Robust Invariant Scalable Keypoints (BRISK) (Leutengger, 2011) and Oriented FAST and Rotated BREIF (ORB) (Rublee et al., 2011).

Feature Detector	Feature Type	Scale Invariance
FAST	Corner	No
Minimum Eigenvalue Algorithm	Corner	No
Harris	Corner	No
BRISK	Corner	Yes
ORB	Corner	No
SURF	Blob	Yes
KAZE	Blob	Yes

Table 1. Feature detectors used in this study.

After detection, a separate set of algorithms are used to described the unique feature. As with the detectors, there are several types of feature descriptors (Table 2). In this analysis, we investigate the effectiveness of four feature descriptors and paired them with their detector counterparts: BRISK, ORB, SURF, and KAZE. In addition, we analyzed the Fast Retina Keypoint (FREAK; Alahi et al., 2012) descriptor using the input from the Harris detector.

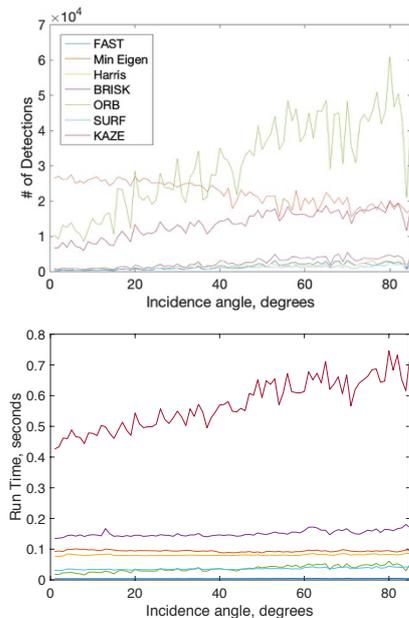
Feature Descriptors	Binary	Scale Invariance	Rotation Invariance
FREAK (input: Harris)	Yes	Yes	Yes
BRISK	No	Yes	Yes
ORB	Yes	No	Yes
SURF	Yes	Yes	Yes
KAZE	No	Yes	Yes

Table 2. Feature descriptors used in this study.

## 2.3 Influence of Illumination on Feature Detectors

As stated, the goal of a feature detector is to identify interesting portions of an image such as the corners of objects or regions of the image that appear similar. Planetary missions capture images of the terrain under a variety of lighting and viewing conditions. Therefore, an effective feature detector must work on images acquired under a variety of solar incidence angles. As seen in **Figure 1**, we selected a 256 x 256 pixel region from three LROC NAC observations with a solar incidence angles of 15°, 45°, and 75°. The seven feature detectors were then applied to the cropped images and the 15 strongest detections from each algorithm were subsequently plotted. At small incidence angles (15°), we see the corner detectors focus on the boundaries of albedo variations around small fresh craters. As the incidence angle increases to 75° the corner detectors start identifying the sharp shadow boundaries around small craters. When applying the Oriented FAST and Rotated BREIF (ORB) detector (Rublee et al., 2011) to the three cropped images, we see concentric circles around a small number of surface features. This indicates that the 15 strongest detection signatures are occurring over the same region and at slightly different scales. For the two blob detectors we investigated, we see that the center of the detections generally lie in the middle of regions with the same intensity values. In the case for the small incidence angle images (15°) the KAZE detector identifies the center of regions where the albedo is higher than the surrounding regolith while at larger incidence angles (75°), the detector is sensitive to the regions of the image in shadow.

To further assess the effect of illumination, we selected a 1024 x 1024 pixel region out of the center of 850 NAC images acquired in incidence angle bins of 1° to 85° (10 images in each bin) and applied the seven feature detectors to the image set. **Figure 2** shows the number of detected features using each algorithm. Apart from the Minimum Eigenvalue Algorithm, the number of detections increase as the incidence angle increases from 1 to 85° with SURF having the smallest increase from about 500 to 1500 detections and ORB having the largest increase from about 1000 to 5000 detections. In terms of processing time, in general, the time did not change with incidence angle except for ORB that had an increase from 0.02 to 0.05 seconds to process the sub-region and KAZE that had an increase of around 150% to process the larger incidence angle images.



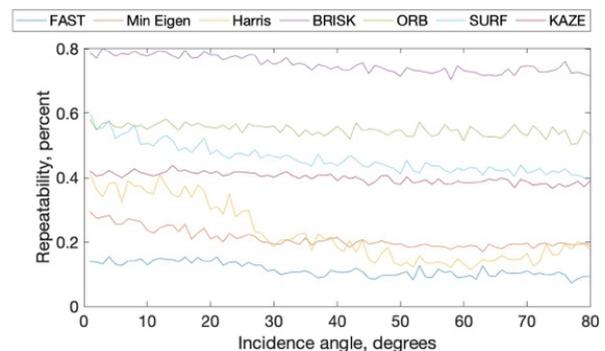
**Figure 2.** Number of detections (top) and run time (bottom) vs. solar incidence angle.

#### 2.4 Assessing the Robustness of Feature Detectors

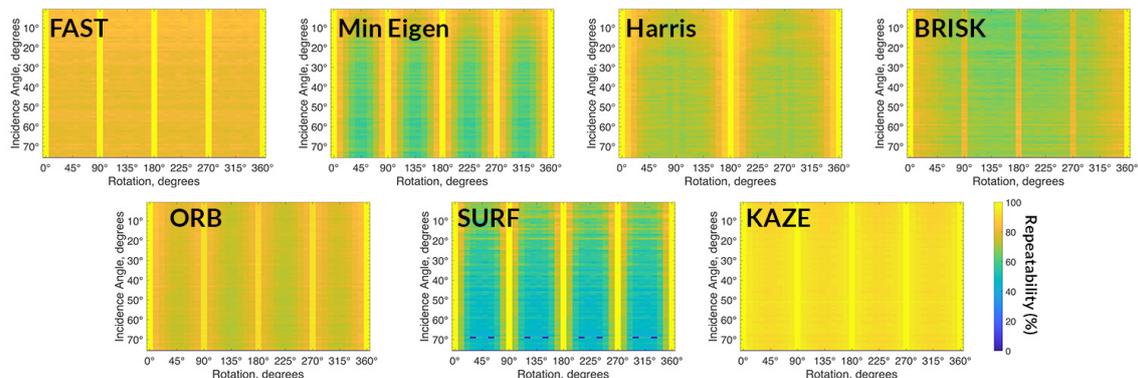
A robust feature detector will identify the same point or region of interest regardless of rotation, scale, and noise. To examine the invariance due to rotation as a function of incidence angle, we selected a 1024 x 1024 region out of each NAC image. We then rotated the sub-image between 0 and 360° in 10° increments. To remove the impact of resampling and interpolation on the rotated images, we then reduced the image size of both the unmodified and the modified image by a factor of two making a 512 x 512

image. We applied the feature detector to the image before and after rotation and computed the ratio of common features detected in both images and the total number of features detected (**Figure 3**). Since the field of view is different in each image due to the rotation, we only counted detected features within 256 pixels of the image center. While the SURF descriptor is rotation invariant, the SURF detector struggled to detect the same surface feature when a slight rotation was applied to the image. Likewise, Minimum Eigenvalue Algorithm, BRISK, and Harris detectors did not perform as well when the image was rotated. On the other hand, the FAST, ORB, and KAZE detectors continued to match the same features when comparing to a rotated image.

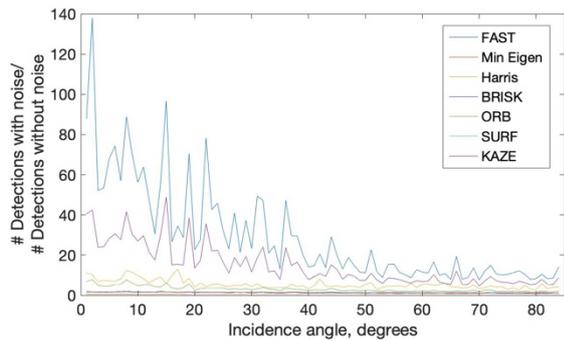
Next, we assessed the feature detectors to changes in scale. In this case we resampled the LROC NAC sub-image from 1024 x 1024 to 512 x 512 and applied each feature detector to both images then compared them to identify how many features were detectable in the lower resolution image (**Figure 4**). As expected, the corner detectors that are scale independent fared better than the corner detectors that were not designed for detecting features over a range of scales. Ranked from best to worst in terms of repeatability at small incidence angles: BRISK, ORB, SURF, KAZE, Harris, Min Eigen and FAST. Overall, the repeatability was not affected by incidence angle with the exception of the Harris detector that was repeatable 40% of the time when the incidence angle was small (<20°), but the repeatability decreased to 20% when the incidence angle was >30°. Furthermore, as seen in **Figure 1**, the corner detectors that are not designed for changes in scale identified small contrast boundaries located around the smallest detectible impact craters. When the scale of the image was changed, these small craters were no longer visible. Therefore, the detector was not able to locate the same surface features in the reduced image.



**Figure 4.** Repeatability of detection when the images is resized by a factor of two.



**Figure 3.** Repeatability of detection when the image is rotated.



**Figure 5.** Increase in the number of false feature detections when salt and pepper noise is applied to 2.5% of the frame.

Finally, we assessed the performance of each detector to noise. Depending on the planetary mission and instrument, different types of noise patterns effect the image. For example, some instruments suffer from salt and pepper noise due to bit errors while the quality of other images are reduced by Gaussian noise. **Figure 5** shows the impact of salt and pepper noise modifying 2.5% of the pixels on the detector results. At small incidence angles, the performance of the FAST and BRISK detectors suffered greatly with the inclusion of noise resulting in 40 to 140 times as many detections as the original images without noise. This is likely due to corner detectors matching the smallest detectable features, which in this case is the salt and pepper noise pattern. However, at high incidence angles, the effect of the noise was reduced on both detectors. Meanwhile, the two blob detectors suffered the least from the additional image noise.

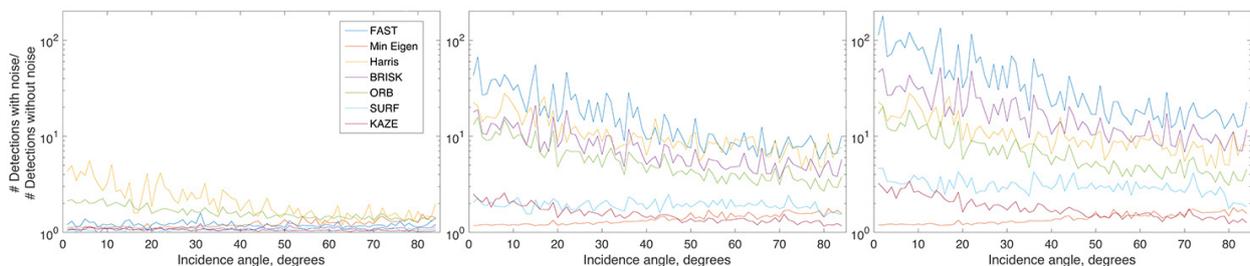
When examining the robustness to Gaussian noise with a mean of zero and a variance of 0.001, 0.011 and 0.021, we found that at small variance levels, the Harris corner detector suffered the

worst and at higher levels the FAST detector was more susceptible to the Gaussian noise. Again, we see that the effect of the noise decreases as the incidence angle increases and the two blob detectors (SURF and KAZE) as well as the Minimum Eigenvalue Algorithm suffered the least from the inclusion of additional Gaussian noise.

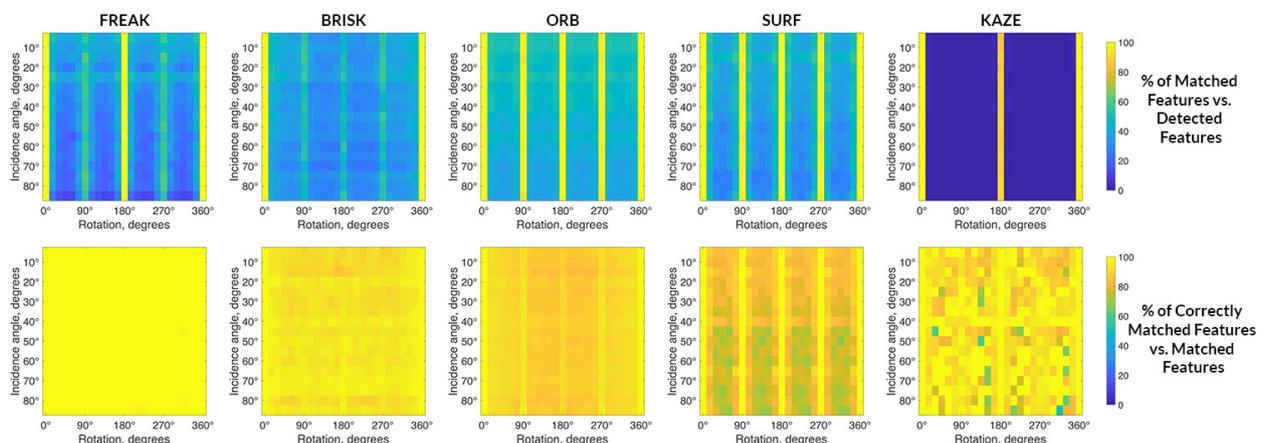
### 2.5 Examining the Robustness of Feature Descriptors

As with the feature detectors, we assessed the robustness of feature descriptors to rotation, scale, and noise. Likewise, we selected a 1024 x 1024 region out of each NAC image. We then rotated the sub-image between 0 and 360° in 15° increments. To remove the impact of resampling and interpolation on the rotated images, we then reduced the image size of both the unmodified and the modified image by a factor of two making a 512 x 512 image. We used the result of the appropriate feature detectors from section 2.4 as input into the descriptor algorithm (FREAK → FREAK; BRISK → BRISK; ORB → ORB; SURF → SURF; KAZE → KAZE). We then used a brute force feature matcher to compare the descriptions and produce a set of matched points. By knowing the actual displacement caused by the image rotation, we calculated the number of correct and incorrect matches. **Figure 7** shows the results to the number of matched features vs. the number of detected features as well as the percentage of correctly matched features. Most descriptors performed well, except for the KAZE descriptor, which was not able to correlate features when the image was rotated. The bottom row of Figure 7 shows that when matched, the accuracy was high (>85%) for all the descriptors tested.

Next, we examined the robustness of the five feature descriptors to a change in image scale by resizing the image to half its length on each side. We then compared the number of true matches to the total number of matches identified (Figure 8). With the input provided by the Harris detector, the FREAK descriptor did not



**Figure 6.** Increase in the number of false feature detections when Gaussian noise is applied to image with a mean of zero and a variance of 0.001 (Left), 0.011 (middle) and 0.021 (right).



**Figure 7.** Performance of feature descriptors when the image is rotated.

identify any correct matches. Of the remaining descriptors, the SURF and BRISK descriptors worked the best with our dataset.

Finally, we examined the robustness of the feature descriptors to noise (Figure 9) following the same techniques discussed in the previous section and found that the SURF and KAZE performed well (>60% correctly matched features vs the number of detected features) when a small amount of noise was included, but all descriptors had less success when additional noise was added to the image.

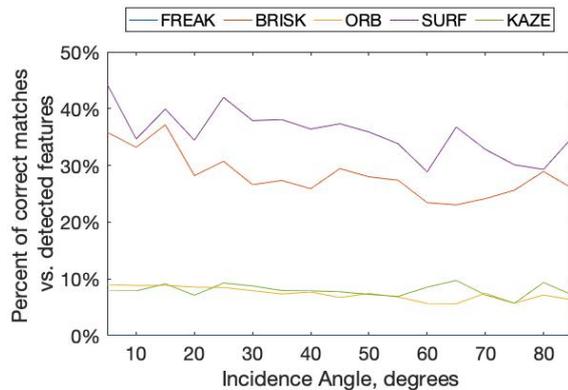


Figure 8. Performance of feature descriptors when the scale of the image was reduced by a factor of two.

### 3. CONCLUSIONS AND FUTURE WORK

Our analysis of feature detectors and descriptors show there is not an ideal set of algorithms that can successfully tie every image of our diverse dataset and maintain robustness to rotation variation, scale, and image noise. In general, the corner detectors focused on small features in the image and slight variations in surface reflectance, while the blob detectors identified regional trends in the dataset. In addition, at high incidence angles the blob detectors focused on uniform shadowed regions. As a result, the performance of corner detectors suffered from changes in image scale and noise than either blob detector we analysed.

In terms of the feature descriptors, the SURF and BRISK descriptors each paired with their detector counterparts generally outperformed the other three detectors in terms of robustness when applying a rotation, scaling, and adding noise to the image. However, when matching the SURF descriptors, we did find that about 25% of the descriptors were incorrectly matched when rotation was applied to the image. Thus, additional tests may be required to remove false matches.

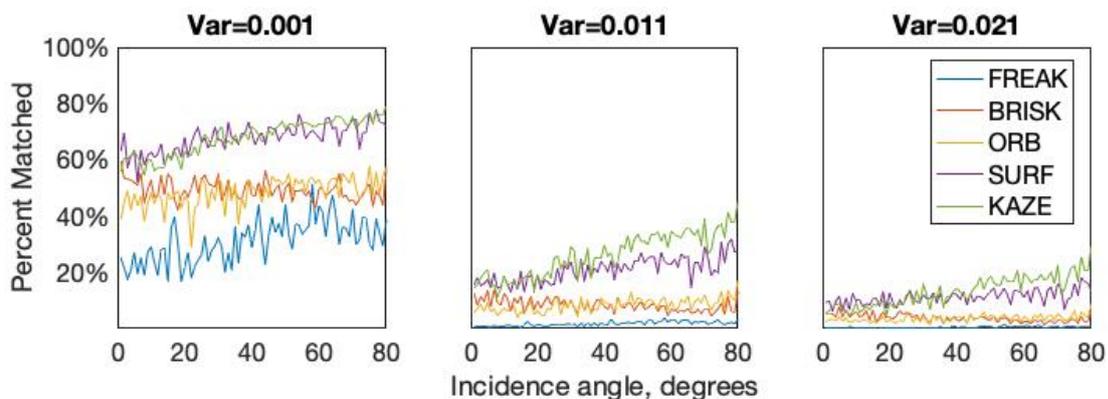


Figure 9. Performance of feature descriptors when Gaussian noise is applied to image with a mean of zero and a variance of 0.001, 0.011, and 0.021.

It should also be noted that some algorithms (i.e., KAZE detector) required much more time to execute than others (Figure 2). Future work will help quantify the resource usage of each algorithm in terms of CPU usage and memory to better understand which set of algorithms could be applied to real-time applications, such as image based navigations and other routines needed for pinpoint landings.

While this study investigated the impact of the solar incidence angle, further work will address how feature detectors and descriptors handle images with different emission angles. This analysis will help identify proper routines needed to register features in geometric stereo observations and can provide a sparse set of tie points for image triangulations and terrain reconstruction.

### 4. ACKNOWLEDGEMENTS

The authors would like to acknowledge support for this work from a NASA Planetary Data Archiving, Restoration, and Tools (PDART) program under NASA grant NNX15AJ59G (PI: E. Speyerer) and the dedication of the LRO Science Teams that produced an invaluable foundational dataset for science and exploration analysis.

### 5. REFERENCES

Alahi, A., Ortiz, R., and Vandergheynst, P., 2012. FREAK: Fast Retina Keypoint. In: *IEEE Conference on Computer Vision and Pattern Recognition*. doi.org/10.1109/CVPR.2012.6247715.

Alcantarilla, P.F., Bartoli, A., Davison, A.J., 2012. KAZE Features. In: *Computer Vision – European Conference on Computer Vision 2012*. 214-227.

Alcantarilla, P.F., Bartoli, A., Davison, A.J., 2012. KAZE Features. In: *Computer Vision – European Conference on Computer Vision 2012*. 214-227.

Anderson, J.A., Sides, S.C., Soltesz, D.L., Sucharski, T.L., and Becker, K.J., 2004. Modernization of Integrated Software for Images and Spectrometers. In: *Lunar and Planetary Science Conference XXXV*. Abstract #2039.

Bay, H., Tuytelaars, T., Gool, L., 2008. SURF: Speeded Up Robust Features. In: *European Conference on Computer Vision 2006*. pp 404-417.

Becker, K. J., Robinson, M. S., Becker, T. L., Weller, L. A., Edmundson, K. L., Neumann, G. A., et al., 2016, First Global Digital Elevation Model of Mercury. In: *47th Lunar and Planetary Science Conference*. Abstract #2959.

Harris, C. and Stephens, M., 1988. A Combined Corner and Edge Detector. *Alvey Vision Conference*. 15.

Humm, D.C., Tschimmel, M., Brylow, S. M., Mahanti, P., Tran, T. N., Braden, S. E., Wiseman, S. Danton, J., Eliason, E. M., Robinson, M. S., 2016. Flight Calibration of the LROC Narrow Angle Camera. *Space Science Reviews*, 200, pp 431-473.

Leutenegger, S., Chli, M. and Siegwart, R., 2011. BRISK: Binary Robust Invariant Scalable Keypoints. In: *2011 International Conference on Computer Vision*. pp. 2548-2555. doi.org/10.1109/ICCV.2011.6126542.

Robinson, M. S., Brylow, S. M., Tschimmel, M., Humm, D., Lawrence, S. J., Thomas, P. C. et al., 2010. Lunar Reconnaissance Orbiter Camera (LROC) instrument overview. *Space Science Reviews*. 150(1–4), 81–124. doi.org/10.1007/s11214-010-9634-2.

Rosten, E.; Drummond, T., 2005. Fusing points and lines for high performance tracking. In: *IEEE International Conference on Computer Vision*. 2: 1508–1511. doi.org/10.1109/ICCV.2005.104.

Rosten, E.; Drummond, T., 2006. Machine learning for high-speed corner detection. In: *European Conference on Computer Vision*. 1: 430–443. doi.org:10.1007/11744023\_34.

Rublee, E., V. Rabaud, K. Konolige and G. Bradski, 2011, ORB: An efficient alternative to SIFT or SURF. In: *Proceedings of the 2011 International Conference on Computer Vision*, 2564–2571.

Shi, J., and Tomasi, C., 1994. Good Features to Track. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. June 1994, pp. 593–600.

Speyerer, E. J., Wagner, R. V., & Robinson, M. S., 2016. Geometric calibration of the Clementine UVVIS camera using images acquired by the Lunar Reconnaissance Orbiter. In: *ISPRS - International Archives of the Photogrammetry, Remote Sensing*