(Preprint) AAS XX-XXX

USING ONE-DIMENSIONAL SIGNAL FEATURES FOR IMAGE-BASED LUNAR SKYLINE RECOGNITION

Christopher Land*, Paul D. McKee†, Siddarth Kaki‡, Jacob Sullivan§

As lunar surface exploration becomes a higher priority there is increased need for methods of lunar surface navigation. Autonomous navigation (i.e., without reliance on human operators or other assets) will become essential for safe and robust surface operations without overburdening Surface Communications and Navigation (SCAN) resources. A recent development in this field is the utilization of the visible skyline for localization, either by comparison to pre-rendered skylines or to skylines rendered in real-time. This requires a fast and accurate way to compare two skylines (rendered and observed) and to rank skyline similarity.

This work develops a novel one-dimensional (1D) signal feature. This is the lower-dimensional analog of familiar 2D image features, often used in image processing and computer vision. This signal feature is then used in the context of finding common points between disparate sets of skyline points, which are expressed as sets of corresponding azimuth and elevation angle pairs. This feature can be used to rank the similarity between skylines rendered at different locations on the lunar surface. It is shown that this technique is less computationally expensive and has higher accuracy when compared to other skyline similarity metrics. Finally, these 1D features are demonstrated in the context of a full surface navigation simulation using Apollo 17 images and data.

INTRODUCTION

More and more spacecraft are traveling to the lunar surface for science and exploration. This number will only increase as the Artemis program advances. Any mission in which a landed vehicle moves appreciably from its landing site (i.e., a rover) will need to navigate on the lunar surface. This could take the form of determining relative distance and direction to a nearby waypoint, referred to here as "local navigation," or finding one's global position on the lunar surface (e.g., latitude and longitude), referred to here as "global navigation." The ability to do so autonomously—that is, without reliance on Earth-based ground stations or other lunar surface or orbital assets—will provide increased redundancy and can help alleviate strain on already overtaxed systems. In the case of human lunar surface exploration, this capability could be a matter of crew safety. This is especially so under loss-of-communication scenarios.

The method of lunar surface navigation pioneered during the Apollo program, which was comparatively primitive due to cost and schedule constraints,[1] can be greatly improved upon using modern

*OSTEM Intern, Aeroscience and Flight Mechanics Division, NASA-JSC, 2101 NASA Pkwy, Houston, TX 77058
†GN&C Engineer, Aeroscience and Flight Mechanics Division, NASA-JSC, 2101 NASA Pkwy, Houston, TX 77058
‡Navigation Discipline Engineer, Aeroscience and Flight Mechanics Division, NASA-JSC, 2101 NASA Pkwy, Houston, TX 77058
§LTV GN&C System Manager, Aeroscience and Flight Mechanics Division, NASA-JSC, 2101 NASA Pkwy, Houston, TX 77058

technology and techniques. The Apollo Lunar Roving Vehicle (LRV, a.k.a. "Moon Buggy") had no global navigation capability, and relied on integrating data from a gyroscope and wheel encoders to estimate distance and direction to the lander.[2] Engineers at NASA's Jet Propulsion Laboratory (JPL) have been working on surface navigation techniques for the Mars rovers for decades. These largely rely on stereo image processing and computer vision for hazard detection and avoidance, as well as visual odometry.[3–5] Several methods of lunar surface navigation have been developed since the Apollo days, and several more are still under development today.[6] If current trends continue, lunar surface navigation by a rover will almost certainly leverage one or more optical cameras.

One method of surface navigation which has seen recent development is the utilization of the observable skyline for localization.[7–10] This technique relies heavily on the uniqueness of the visible skyline at a particular point. The relationship between the shape of the observed skyline and an observer's position is highly nonlinear—one cannot perform localization directly from skyline data. Rather, one must leverage pre-rendered skylines at known locations or some form of on-board capability to render the expected skyline at the current location. This dichotomy leads to two skyline-based navigation paradigms: coarse and fine skyline navigation. In the coarse case, skylines can be pre-rendered from Digital Elevation Map (DEM) data at pre-determined waypoints near a region of a planned traverse. A rover can then take an image or series of images (to construct a panorama) and compare the observed skyline to the pre-rendered skylines for closest match, thus finding the nearest waypoint. In the fine sense, a skyline can be rendered from a DEM on-board using an initial guess of the vehicle pose. Points on the expected and observed skylines can be correlated, and the rover pose can be updated by solving the perspective-n-point (PnP) problem. These steps are repeated until the pose solution converges to satisfaction. Both the coarse and fine skyline-based navigation techniques rely on the ability to correlate points between skylines, which are readily expressed as corresponding azimuth and elevation angle pairs. This essentially amounts to a one-dimensional signal correspondence problem.

This work develops novel one-dimensional (1D) signal features which can be employed for skyline comparison. These are analogous to familiar two-dimensional (2D) image features such as SIFT,[11] SURF,[12] ORB,[13] KAZE,[14] and BRISK.[15] The features described in this work utilize a SIFT-like difference-of-Gaussians approach to identify features, and perform a local polynomial fit of skyline points around each feature to develop a corresponding feature descriptor. It is observed that these 1D features will have applications outside of lunar surface navigation whenever one wants to compare two signals. It is further observed that these are not the first or only 1D signal features. Substantial work has been done towards developing 1D convolutional neural networks for processing biomedical signals (e.g., ECG) and other similar applications.[16,17] There has also been prior work on developing 1D features for signal processing,[18,19] even in the field of vision-based navigation.[20] Researchers have also considered the application of 2D features to images of signal plots.[21] Of these, the closest work to that described in this manuscript is that of Xie and Beigi,[18] who developed SIFT-like signal features and angle-based descriptors. In contrast, horizon-based navigation does not require scale-invariant features—it is reasonable to assume that skylines can be expressed in terms of absolute azimuth and elevation angles, meaning that rendered and observed skylines will have the same scaling (but perhaps different sampling density).

This manuscript will detail the coarse skyline navigation process, which consists of (1) image processing, (2) skyline stitching to produce a panorama, (3) skyline rendering using DEM data, and (4) skyline comparison. Focus will be given to the development of a new 1D signal feature and its application to solving the challenges in skyline comparison. It is demonstrated that the 1D feature technique can find common points in overlapping images taken at the same location (the sorts of images that can be stitched into a panorama), and it is demonstrated that the 1D feature can find common points between observed and rendered skylines at the same location. It will be shown that this navigation technique can find the nearest waypoint for a given set of Apollo 17 surface images and associated Taurus Littrow Valley DEM.

COARSE SKYLINE NAVIGATION

This section details the method of coarse skyline navigation, which stands in contrast to the fine skyline navigation described in Reference 8. The premise is to compare observed skylines (e.g., from a single image or an ensemble of images combined into a panorama) to a set of skylines which are pre-rendered using DEM data at a set of known waypoints. The observable skyline functions like a fingerprint—one cannot determine a person's identity from their fingerprint alone, but their identity can be found by a fingerprint database search. Similarly, a lost rover on the Moon can determine the waypoint whose skyline is closest to its own, provided that the local terrain is sufficiently unique and feature-rich and there is a robust method of scoring skyline similarity. This technique is particularly well-suited for a lost-on-the-Moon filter re-initialization problem, and is less useful during nominal surface operations (wherein a fine skyline navigation technique finds utility). Before proceeding, it is observed that the 1D features detailed in this manuscript are useful not only for coarse skyline navigation, but can help solve the skyline-DEM point correspondence problem inherent in fine skyline navigation.[8]

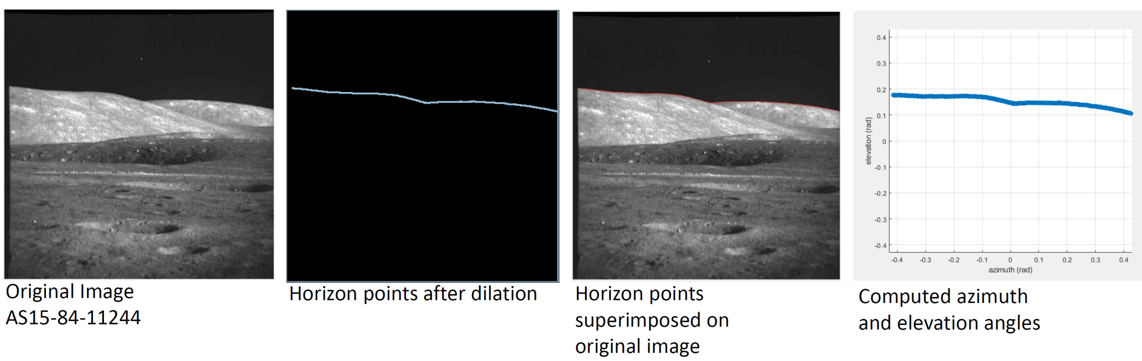

Figure 1. Image processing is performed to extract horizon points, which are converted from pixels to angle pairs.

Coarse skyline navigation consists of four steps: (1) image processing to find skyline points, (2) image stitching to produce a skyline panorama, (3) *a priori* offline skyline rendering from DEM data, and (4) skyline comparison. The first three can be done in any order, but the fourth step must happen last. Image processing begins by scanning each column in the image vertically from top to bottom, looking for a sudden change from dark to bright. If desired, sub-pixel edge finding techniques[22,23] can be used for improved performance. Depending on lighting conditions and other image artifacts, some of the horizon pixels may, in fact, be spurious returns. It is recommended to use image processing techniques such as connected components, dilation, and erosion to find a single continuous set of horizon pixels. When finished, there should be only one skyline pixel per each column in the image. These can then be converted from pixel coordinates to unit vectors in the camera frame by solving the reverse mapping problem,[24] and these unit vectors can be used to find corresponding azimuth and elevation angles for all the points. In this work, azimuth is defined as a horizontal angle which is zero along the camera boresight and positive to the right, and elevation is a vertical angle which is zero along the camera boresight and positive up. This image processing process is shown on an Apollo 15 lunar surface image in Fig. 1.

There are three methods to arrive at a panoramic skyline. The first is to combine images to form a panorama, then to process the panoramic image as a whole to find skyline points. Compared to the other methods, this is slower and more computationally expensive, and can introduce warping or other image stitching artifacts. The second is to find the skyline points in image A, find the skyline points in image B, and use skyline comparison techniques (detailed below) to find how best these fit together. This is the fastest technique, but is the most sensitive to noisy data, and can fail altogether if the skyline is sufficiently feature-poor. The third approach is a hybrid of the first two. Find matching image features (e.g., SIFT) in image A and image B, and then solve Wahba's problem to find the rotation between images. Then, transform the image B skyline angles into the image A frame (or vice versa). This method leverages a feature-rich foreground, and ultimately arrives at a 360-degree panorama of skyline azimuth and elevation angles.

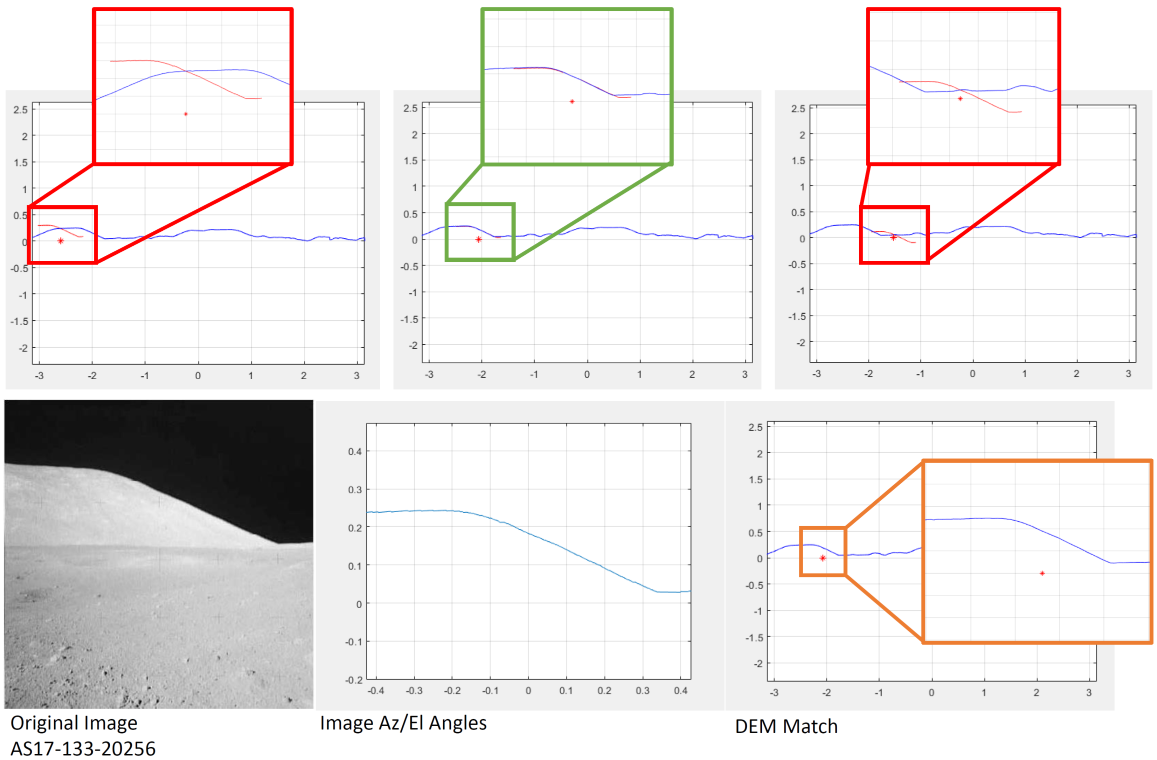

Figure 2. Normalized cross-correlation can be used to detect matching regions of similar skylines, but is less well-suited to skyline comparison when the observer position is unknown.

Skyline comparison is performed in angle space rather than pixel space in this work. This does away with any idiosyncrasies introduced by a particular camera's intrinsic and distortion parameters. As such, skyline rendering does not involve generating a synthetic image of the skyline, but rather computing azimuth and elevation angles that would be seen by an observer at a particular location and orientation. This can be done by placing an observer point on a DEM, then finding and returning the highest terrain height (or the terrain height that produces the highest elevation angle) along each azimuth scan line. Such rendering can be done by brute force (i.e., compute bearing angles to every point in the DEM, then skim along the top and keep only the angles corresponding to the skyline), or by a more efficient binary search for the skyline point (and corresponding elevation) along each azimuth, or through ray tracing.

The overall effectiveness of this navigation technique is very sensitive to the method of skyline comparison that is employed. Perhaps the most obvious comparison method is normalized cross-correlation—sweeping observed over expected skylines at various azimuth offset angles to find an orientation at which the two fit together like puzzle pieces. Indeed, this is well-suited to finding an observer's orientation when comparing a single image to a rendered skyline at that same location, as is illustrated in Fig. 2. In such a situation, there is a correct answer—an azimuth offset at which the two skylines will "agree." It was demonstrated that for Image AS17-133-20256 taken on Apollo 17 on EVA 2 at Station 4, the observed skyline almost perfectly matches a patch of the rendered skyline.

In other cases, normalized cross-correlation does not work well. When comparing two panoramic skylines rendered at different locations, we need a comparison method that scores a match more poorly as the actual distance between the two observer locations becomes greater. There should be a region of convexity around each waypoint within which that waypoint skyline is the best match for the observed skyline, but this is not always true with normalized cross-correlation. Normalized cross-correlation often returns a better match to a skyline rendered kilometers away simply by coincidence. Suppose an observer begins moving in a straight line towards a mountain range in the distance. Skyline features such as peaks and valleys will move away from the center of projection at different rates (points near the edge of the field-of-view will move faster due to perspective). After only a few hundred meters of traverse, the new skyline will score very poorly in terms of overall shape similarity to the original skyline, but all the same features will be present. Thus, a method of skyline comparison is needed that tracks features in a skyline.

SIGNAL FEATURES

This manuscript introduces novel one-dimensional signal features which can be used to assess the similarity between two lunar skylines. This comparison method is more robust to nonlinear deformation than normalized cross-correlation is. Signal features, like image features, consist of feature detection and feature description. These features use a SIFT-like Difference-of-Gaussians (DoG) approach to find points of interest in the skyline signal. Once a skyline has been extracted from the original image, it is convolved with a 1D Gaussian kernel to produce a smoothed signal—analogous to a blurred image. The difference between the original and the smoothed signal is taken, and prominent peaks and valleys in this differenced signal are considered for points of interest. This is repeated multiple times at different smoothing factors to find features of different scales in the signal. This feature extraction process is illustrated for an Apollo 17 lunar surface image in Fig. 3.

Next, these points of interest must be described in such a way that the same point can be detected and matched in different signals. A region of interest is drawn around the feature point, and the whole signal is rotated such that the start and end points of the window have the same elevation (or y-value). Then, a polynomial is fit to the resulting signal, and the coefficients of this polynomial are arranged into a feature descriptor vector. Later, when comparing features, the Euclidean distance between descriptor vectors is taken as a similarity score.

The efficacy of these features is first demonstrated by finding the same skyline points in terrain images with visible overlap. Image AS17-145-22171 and image AS17-145-22172 were taken at the

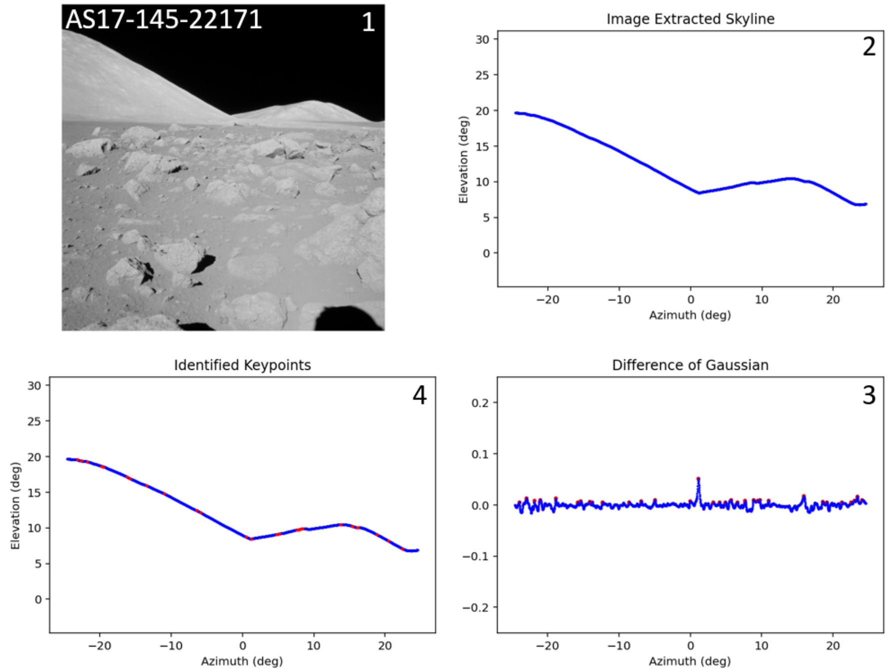

Figure 3. An image is processed to find horizon points, and the "horizon signal" is processed to find keypoints using differences- of-Gaussians.

same location on Apollo 17, and there is some rotation between images. Sufficiently many images of this type can be used to construct a panorama, but we will focus on these two. Both images were processed to find skyline points, and corresponding points were found between images, as shown in Fig. 4. This demonstrates that skylines can be matched between two images using 1D features.

Forward work will include demonstrating the ability to match rendered skylines to image skylines using these same features, and using these features to assign a "similarity score" to two skylines. Then it can be demonstrated that a panoramic image taken anywhere in the Taurus Littrow Valley will be a best match to the nearest rendered waypoint.

CONCLUSION

This work introduces a novel 1D signal feature. This was developed for use in the context of coarse lunar skyline navigation, but is also applicable to fine lunar skyline navigating, and any of a myriad of other signal processing applications (e.g., biomedical signal monitoring, and vibration analysis). The coarse skyline navigation pipeline is detailed, and it is discussed how the signal features out-perform normalized cross-correlation for skyline comparison. Forward work, to be completed between the submission of this extended abstract and the final publication, will consist of fully integrating these signals into an end-to-end coarse skyline navigation simulation. It will be demonstrated that the coarse skyline method using 1D features can find the nearest waypoint in the Taurus Litrow Valley for several images taken by astronauts Cernan and Schmidt. This navigation method can help in future extraterrestrial surface exploration by crewed or autonomous rovers, and would be instrumental in a lost-on-the-surface scenario.

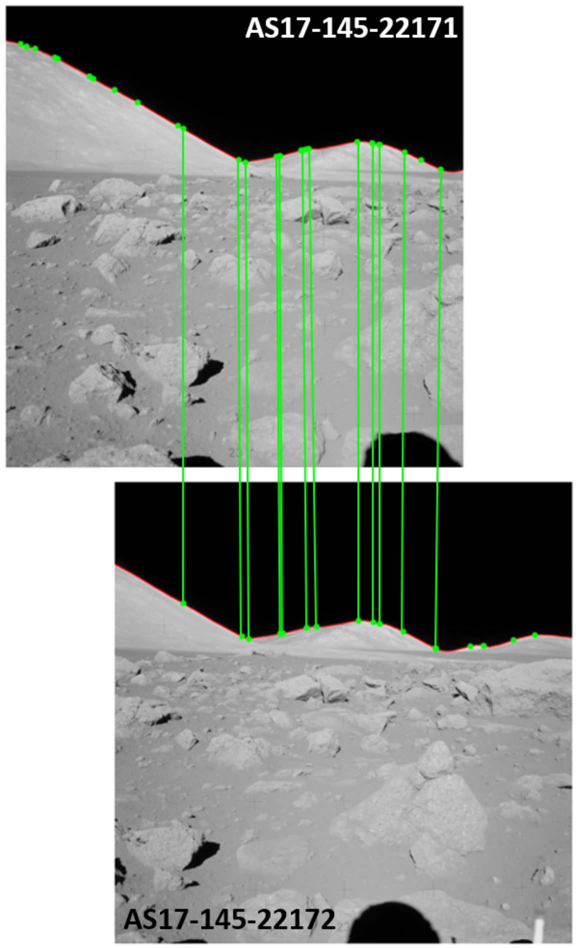

Figure 4. Successive images of the lunar surface with overlapping terrain are processed to find corresponding horizon points.

REFERENCES

[1] H. W. Tindall, "Rover Navigation," Tech. Rep. 69-PA-T-147A, NASA Manned Spacecraft Center, Nov. 1969.

[2] W. G. Heffron and F. LaPiana, "The Navigation System of the Lunar Roving Vehicle," Tech. Rep. TM-70-2014-8, Bellcomm, Inc., Dec. 1970.

[3] M. Bajracharya, M. W. Maimone, and D. Helmick, "Autonomy for mars rovers: Past, present, and future," *Computer*, Vol. 41, No. 12, 2008, pp. 44–50, 10.1109/MC.2008.479.

[4] M. Maimone, Y. Cheng, and L. Matthies, "Two years of visual odometry on the mars exploration rovers," *Journal of Field Robotics*, Vol. 24, No. 3, 2007, pp. 169–186, 10.1002/rob.20184.

[5] V. Verma, M. W. Maimone, D. M. Gaines, R. Francis, T. A. Estlin, S. R. Kuhn, G. R. Rabideau, S. A. Chien, M. M. McHenry, E. J. Graser, *et al.*, "Autonomous robotics is driving Perseverance rover's progress on Mars," *Science Robotics*, Vol. 8, No. 80, 2023, 10.1126/scirobotics.adi3099.

[6] P. D. McKee, "A Survey of Autonomous Navigation Techniques Applicable to Lunar Surface Exploration," *47th Annual American Astronautical Society (AAS) Guidance, Navigation and Control Conference*, No. AAS 25-173, 2025.

[7] A. Liounis, A. Yew, and D. McGann, "Science Geotagging and Localization Using Lunar Horizons," *2023 NASA Exploration Science Forum*, 2023.

[8] W. Driessen, S. Kaki, A. Liounis, D. McGann, P. McKee, A. Tennenbaum, and A. Yew, "Monocular Horizon Navigation," *4th Space Imaging Workshop*, 2024.

[9] E. E. Palmer, J. N. Head, R. W. Gaskell, M. V. Sykes, and B. McComas, "Mercator—Independent rover localization using stereophotoclinometry and panoramic images," *Earth and Space Science*, Vol. 3, No. 12, 2016, pp. 488–509, 10.1002/2016EA000189.

[10] S. Chiodini, M. Pertile, S. Debei, L. Bramante, E. Ferrentino, A. G. Villa, I. Musso, and M. Barrera, "Mars rovers localization by matching local horizon to surface digital elevation models," *2017 IEEE International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, IEEE, 2017, pp. 374–379, 10.1109/MetroAeroSpace.2017.7999600.

[11] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, Vol. 60, 2004, pp. 91–110, 10.1023/B:VISI.0000029664.99615.94.

[12] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," *Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006. Proceedings, Part I 9*, Springer, 2006, pp. 404–417, 10.1007/11744023_32.

[13] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," *2011 International conference on computer vision*, IEEE, 2011, pp. 2564–2571, 10.1109/ICCV.2011.6126544.

[14] P. F. Alcantarilla, A. Bartoli, and A. J. Davison, "KAZE features," *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VI 12*, Springer, 2012, pp. 214–227, $10.1007/978-3-642-33783-3\_16$.

[15] S. Leutenegger, M. Chli, and R. Y. Siegwart, "BRISK: Binary robust invariant scalable keypoints," *2011 International conference on computer vision*, Ieee, 2011, pp. 2548–2555, 10.1109/ICCV.2011.6126542.

[16] S. Kiranyaz, T. Ince, O. Abdeljaber, O. Avci, and M. Gabbouj, "1-D convolutional neural networks for signal processing applications," *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2019, pp. 8360–8364, 10.1109/ICASSP.2019.8682194.

[17] D. Li, J. Zhang, Q. Zhang, and X. Wei, "Classification of ECG signals based on 1D convolution neural network," *2017 IEEE 19th international conference on e-health networking, applications and services (Healthcom)*, IEEE, 2017, pp. 1–6, 10.1109/HealthCom.2017.8210784.

[18] J. Xie and M. S. Beigi, "A scale-invariant local descriptor for event recognition in 1d sensor signals," *2009 IEEE International Conference on Multimedia and Expo*, IEEE, 2009, pp. 1226–1229, 10.1109/ICME.2009.5202722.

[19] Y. Kaya, M. Uyar, R. Tekin, and S. Yıldırım, "1D-local binary pattern based feature extraction for classification of epileptic EEG signals," *Applied Mathematics and Computation*, Vol. 243, 2014, pp. 209–219, 10.1016/j.amc.2014.05.128.

[20] A. Briggs, Y. Li, and D. Scharstein, "Feature matching across 1d panoramas," *6thWorkshop on Omnidirectional Vision*, 2005.

[21] R. Ramele, A. J. Villar, and J. Santos, "BCI classification based on signal plots and SIFT descriptors," *2016 4th International Winter Conference on Brain-Computer Interface (BCI)*, IEEE, 2016, pp. 1–4, 10.1109/IWW-BCI.2016.7457454.

[22] D. T. Renshaw, *Optical navigation using quadrics and indirect indexing & search methods*. Rensselaer Polytechnic Institute, 2022.

[23] J. A. Christian, "Accurate planetary limb localization for image-based spacecraft navigation," *Journal of Spacecraft and Rockets*, Vol. 54, No. 3, 2017, pp. 708–730, 10.2514/1.A33692.

[24] P. D. McKee, R. S. Lovelace, S. A. Lockhart, and J. B. Chong, "Camera Calibration and Alignment Metrology at Johnson Space Center's Electro-Optics Laboratory," *47th Annual AAS Guidance, Navigation and Control (GN&C) Conference*, No. AAS 25-141, 2025.