

# AUTOMATED MAPPING FOR COMPARISON OF SATELLITE IMAGERY

Collin Abidi<sup>1\*</sup> Seth Roffe<sup>1</sup> Evan Gretok<sup>1</sup> and Alan George<sup>1</sup>; <sup>1</sup>NSF Center for Space, High-Performance, and Resilient Computing, University of Pittsburgh, 4420 Bayard Street, Pittsburgh, PA 15213

\*[collin.abidi@nsf-shrec.org]

**Abstract.** *For Earth-observation imagery to become actionable, a rapid means of locating and comparing recent imagery with historically acquired imagery is needed. STP-H7-CASPR has demonstrated the potential of the Satlantis iSIM-90 sensor by successfully capturing and downlinking over 7700 4K images at  $\sim 3.0$  meters per pixel ground-resolved distance. This submission presents a process for geolocating imagery acquired by STP-H7-CASPR using a Python application consisting of API requests, traditional computer-vision techniques, deep-learning image registration, and satellite imagery provided by Google Earth Engine and SentinelHub. The authors demonstrate a method of automatically geolocating images that is accurate to within  $0.04 \text{ deg}^2$  error within the longitude and latitude. The LoFTR method for detector-free feature matching is tested on space-based imagery, with results competitive to classic feature-mapping methods, even without fine-tuning. The pipeline developed in this work enables the comparison of near-real-time Earth-observation imagery to historical imagery for temporal change detection and analysis.*

**Introduction.** The Space Test Program - Houston 7 - Configurable and Autonomous Sensor Processing Research (STP-H7-CASPR) experiment has been in operation on the International Space Station (ISS) since January 12th, 2022. It is an ongoing mission hosted by the Department of Defense (DoD) Space Test Program to develop and test new techniques for automated onboard sensor-data processing [1]. One of the primary sensor payloads aboard CASPR is the Satlantis iSIM-90 multispectral imager, capable of capturing 4K ( $4096 \times 3072$  pixel) images at approximately 3.0 meters per pixel ground-resolved distance (GRD) in four spectral bands [2]. Each image covers roughly 118 square kilometers. CASPR can be observed in Figure 1 as delivered to STP in March 2021.


As of September 2022, STP-H7-CASPR has captured over 7700 4K images. Generating this dataset of high-resolution, low-GRD Earth-observation imagery becomes valuable for downstream temporal-change applications after the imagery is geolocated on the surface of the Earth. An example of such a task is the detection of urban development as addressed in the SpaceNet Challenge [3]. For these downstream applications, consistent meter-level precision feature matching is critical for accurate image alignment. Meeting this requirement with cross-sensor modalities, such as active versus passive imaging, and with variable environmental conditions, such as seasonal change and weather interference, is an active field of research [4, 5, 6, 7].

Traditional approaches to purely visual image geolocation typically use techniques developed in the image-registration domain, such as SIFT-based homography. This method can discover a collection of features within an image and match them to features in mapped satellite imagery. Recently, machine-learning approaches to image registration have emerged, such as LoFTR [8], that use deep-learning feature matching techniques rather than handcrafted features with a detection stage for image registration.


The task of this research is to accurately geolocate CASPR imagery given only a downlink time. By leveraging a variety of open-source satellite-imagery datasets and computer-vision techniques, we demonstrate the benefits and drawbacks of several approaches when considering images collected from our sensor. We leverage Google Earth Engine [9], SentinelHub [10], OpenCV [11], and an ISS locator API, “Where the ISS At?”, [12] to construct a process that automates the georegistration procedure for CASPR imagery.

**Background and Related Work.** The background section introduces the Space Test Program and details on the CASPR mission. It also provides information on Earth-observation platforms and tools as well as feature extraction methods. Several potential applications of this imagery and pipeline are also discussed.

**STP-H7-CASPR.** The Department of Defense (DoD) Space Test Program (STP) was established in 1965 and now provides affordable and efficient spaceflight opportu-



**Figure 1.** STP-H7-CASPR as delivered to STP in March 2021



**Figure 2.** The projective mapping of SIFT features from a CASPR image to a query image (green lines) is shown as it relates to the proper alignment of this CASPR image (red square) onto Google Earth imagery - Imagery Credit Google Earth [13] and STP-H7-CASPR - STP-H7-CASPR was Integrated and Flown by the DoD Space Test Program - *DISTRIBUTION STATEMENT A: Unlimited Distribution (AFRL-2022-1571)*

nities for space experiments on the ISS [14]. STP hosts DoD payloads on the ISS to advance on-orbit research and technology. STP-H7 is a pallet of experiments that was flown to the ISS on the SpaceX Falcon 9 24th Commercial Resupply Service (CRS-24) mission and attached to the outside of the Columbus Module in December of 2021. The CASPR experiment, as one of many on the pallet, is currently operating aboard the ISS performing autonomous-sensing operations and experiments on Earth imagery.

CASPR's primary imager is the iSIM-90, which was designed and developed by the Spanish company Satlantis. The iSIM-90 is a binocular, multispectral telescope which collects data in the visible and near-infrared (VNIR) spectrum at  $\sim 3.0$ -m GRD. The VNIR bands are captured by four filters split between two separate imagers, which allows for multispectral imagery without the need for moving parts such as a filter wheel. Namely, one imager contains a near infrared (NIR) and a red filter split down the center of the image, while the other imager contains a blue and a green filter split down the center of the image. Data from the iSIM-90 is stored in the form of a  $4096 \times 3072$  8- or 12-bit image from each sensor.

While onboard processing is the ultimate goal of CASPR, effective change detection requires large amounts of up-to-date mapping and situational data that is impractical to maintain onboard a space platform. Downlinking imagery and applying this process for mapping combines the best of both approaches. This enables effective use of onboard computing and application of the data to real-world needs.

*Earth-Observation Imagery.* Modern satellite-imaging solutions provide a growing stream of Earth-observation data for scientific and commercial purposes. Governments have long lead the charge in providing this data. Perhaps the most extensive program has been NASA Landsat,

providing consistent global land-cover content since 1972. The Operational Land Imager used on Landsats 8 and 9 can generate 15-m GRD panchromatic and 30-m GRD RGB imagery [15, 16]. Similarly, the European Space Agency (ESA) Copernicus program flies several satellites for Earth-observation. The Sentinel-2 twin-satellite constellation attains 10-m GRD for visible bands [17].

More recently, commercial enterprises have entered the Earth-observation satellite space. Maxar Technologies has produced a series of WorldView satellites since the late 2000s that are capable of panchromatic imaging to as low as 0.3-m GRD [18]. The next-generation WorldView Legion constellation of six satellites will match this spatial resolution while enabling improved revisit times and generation volume [19]. Planet Labs operates multiple constellations totalling over one hundred Earth-observation satellites. The Planet Dove constellation produces 3.0-m GRD imagery of the entire planet's surface each day [20, 21, 22]. Of the platforms considered in this section, this is the most similar in spatial resolution and optical quality to CASPR. Additionally, the Planet SkySat constellation enables tasking of imagery at 0.5-m GRD [23, 24]. The next-generation Planet Pelican offering will generate up to 30 tasked target revisit captures per day from 32 satellites at 0.3-m GRD [25].

As ever larger quantities of Earth-observation data are being generated, many sources of this imagery are developing systems to ease its viewing, comparison, and application. Planet provides a series of tools and interfaces to browse and compare daily image captures and long-term composite "basemaps" generated from them [26, 27]. Less powerful and current but much more widely accessible, Google Earth has provided similar forms of access to and visualization of satellite imagery for the general public since 2005. More recently, Google Earth Engine was introduced to support analysis of this imagery via cloud



resources [28]. Google Earth Engine now provides seamless access to Planet APIs for delivery of imagery [29]. Similar resources are provided by competing platforms, including SecureWatch from Maxar [19], Sentinel Hub for Copernicus missions [30], LandsatLook for Landsat [31], and USGS Earth Explorer for all platforms that feed into the US Geological Survey [32].

Proliferation of Earth-observation imagery is leading to its application across a variety of domains. [33] provides a review of the use of Sentinel-2 imagery for tasks ranging from general land-use classification to forest and agricultural monitoring to disaster response. Higher-resolution Earth-observation imagery, such as the 3-m GRD of CASPR, can enable more complex tasks such as evaluating particulate matter pollutants from orbit [34] or detecting overfishing [35], illegal mining, and other nefarious activities. The SpaceNet Challenges have produced numerous methods of extracting actionable data from Earth-observation imagery, from building footprint and road network extraction to urban development assessment and flood detection [36]. A common thread for all of these tasks: the geolocation offered by our pipeline is critical for validation and directed action based on these methods.

*Feature Extraction.* This research considers four feature-extraction methods in order to perform image matching. Feature-extraction methods are used to convert signals from raw data into numerical formats that contain information from the original data. In the image-processing domain, collections of pixel values can be transformed into useful features using feature-mapping techniques. This section discusses four methods that perform feature mapping in our pipeline: SIFT, AKAZE, ORB, and LoFTR.

Features from the Scale-Invariant Feature Transform (SIFT) algorithm, a classic feature-mapping method, are used to locate and describe keypoints within images [37]. The reader is referred to the original paper for a full discussion of the SIFT algorithm [37]. Liu et al. [38] use SIFT features to increase object-recognition accuracy in satellite imagery. SIFT features, designed to be both scale- and rotation-invariant, are used as a replacement for the object-proposal method stage in R-CNN. Liu et al. evaluate their proposed framework on a database of satellite imagery containing planes and ships. Joshi et al. [39] provide an analysis of several feature-extraction algorithms coupled with the support-vector machine (SVM) algorithm for Landsat 8 image retrieval. The precision and recall results indicate that SIFT fails in both low-texture regions, or areas with low spatial frequency features, and areas of the image with high spatial frequency components, when compared with the Gabor filter used for texture analysis [40]. Xiang et al. [41] explore the failures of SIFT registration in high-resolution optical and synthetic-aperture radar (SAR) imagery. The authors propose optical-to-SAR SIFT (OS-SAR) to perform feature registration in the presence of large variations in il-

lumination and geometry. Results indicate that OS-SIFT features have good performance in image registration and keypoint detection, but could be improved by considering regional data, such as mutual information, rather than just local features.

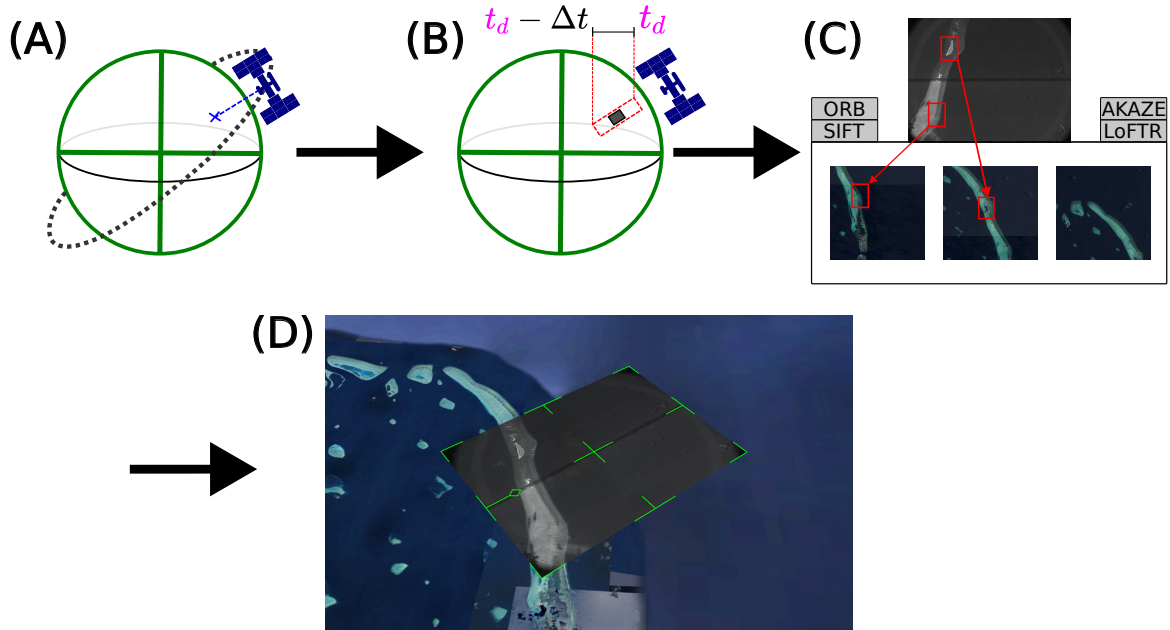
The Accelerated-KAZE (AKAZE) algorithm was developed by Fernández[42] to address the computational challenge of feature detection and description on low-power, camera-enabled devices. The reader is referred to the original AKAZE paper for a full description of the algorithm [42]. Elantcev et al. [43] compare SIFT, SURF, KAZE, ORB, BRISK, and AKAZE features for matching between unmanned aerial vehicle (UAV) and satellite imagery in the presence of significant environmental variation. The authors present a modified statistical differentiation method for increasing image-matching accuracy. The comparison between features demonstrates that AKAZE has the highest percentage of successfully matched images in an urban environment but performs the worst in a rural setting.

The ORB feature descriptor was developed by Rublee et al. [44] as an alternative to SIFT. ORB is compared to SIFT and a deep-learning method for registration of Landsat 7 images by [45]. Their study demonstrates how SIFT and ORB accuracies degrade as the similarity decreases between image pairs. Demchev et al. [46] evaluate ORB, SIFT, and AKAZE on image registration between sequential SAR images. ORB is shown to return far fewer feature vectors when compared to SIFT and AKAZE.

A detector-free method for matching local features using transformers is presented by Sun et al. [8]. This research describes the Local Feature Transformer (LoFTR), which performs pixel-wise dense matching in a coarse-to-fine manner. The transformer architecture provides a global receptive field that enables keypoint detection in low-texture regions, increasing detection repeatability. LoFTR highlights the importance of utilizing the global context of features rather than just those in a local neighborhood and is the only method in this research that does not use descriptors in the local feature-matching stage. Zhao et al. [47] employ LoFTR in a image-matching pipeline similar to the one proposed in this research, though the focus of their study is improving matching with deep-learning style-transfer methods. As of publication, no publications on LoFTR applied to the space-imagery domain aside from this work and [47] have been identified.

**Approach.** This section presents an automated pipeline for geolocating Earth-observation imagery. This pipeline uses key features within the captured imagery to map it to existing Earth-observation data. This geolocation process enables temporal-change detection and other Earth-observation tasks.

First, the ISS location is determined using the timestamp of the downlinked image. This position can be retrieved from a web-API call that returns a latitude and longitude pair. A variable-length region-of-interest is then



**Figure 3.** (A) CASPR image downlink time is associated with the ISS location projected to the surface of the Earth (B) A search space is generated from  $t_d$  to  $t_d - \Delta t$  (C) A query set is downloaded and image registration occurs between  $I_C$  and  $I_Q$  using one of the methods described (D) An example of a successful alignment at Maamutaa Island, Maldives. Image credit Google Earth [48] and STP-H7-CASPR - STP-H7-CASPR was Integrated and Flown by the DoD Space Test Program - DISTRIBUTION STATEMENT A: Unlimited Distribution (AFRL-2022-1571)

generated, with sample density, time window length, and resolution as tunable parameters. This procedure can be seen in Figure 3. Using this region-of-interest, an Earth-observation imagery source is automatically queried using an API. Google Earth Engine and SentinelHub were used to download Sentinel-2 imagery. Images are sorted by amount of cloud cover to minimize error introduced by weather. The SIFT, AKAZE, and ORB methods perform feature extraction on both CASPR images and query images. LoFTR, which is detector-free, extracts feature matches in a coarse-to-fine manner. It should be noted that the LoFTR technique required image downsampling to  $640 \times 480$  to fit architecture parameters. As such, the effective GRD in both CASPR and query images are reduced. Next, the Random Sample Consensus (RANSAC) algorithm is used to produce a projective transformation from features in the image captured from CASPR, denoted  $I_C$ , to query images, denoted  $I_Q$ . Ideal parameters for query image GRD, height, width, and spatial coverage were found primarily by maximizing the number of feature matches that pass Lowe’s Ratio Test [37] during the homography process and by manually examining final alignment quality on a series of test scenes. A subset of 36 images from the full CASPR dataset were selected and manually georeferenced as a source of ground truth.

The homography process begins by extracting feature descriptors from each of the captured and query images

using each of the algorithms described above. Once descriptors are extracted, a k-nearest-neighbor search with  $k = 2$  is performed between all pairs of captured and query image feature descriptors using the Fast Library for Approximate Nearest Neighbors (FLANN) [49]. Following this, Lowe’s Ratio Test compares the distances between pairs of matched feature descriptors to reject non-distinct, false matches. Specifically, we check that the Euclidean distance between the closest matched pair,  $m$ , and the second-closest matched pair,  $n$ , is large enough such that  $(m \cdot \text{distance}) < \tau \times (n \cdot \text{distance})$ . In the case of SIFT and AKAZE,  $\tau = 0.7$ . In the case of ORB,  $\tau = 0.8$  to prevent over-rejection of matches. If there are more than four feature descriptors, the perspective transformation,  $H$ , can be found between  $I_C$  and  $I_Q$ . The RANSAC algorithm, with a constant threshold of 15.0, is used to select the best subset of matches. Once the best subset of matches is found, a least-squares estimate method is used to minimize back-projection error and return  $H$ . Any homogeneous point  $[x_i, y_i, 1]$  in  $I_C$  can then be mapped to  $[x'_i, y'_i, 1]$  in  $I_Q$  using  $H$ . A visual example of the homography mapping calculated using this process is shown in Figure 2. Finally, position and orientation information obtained during the homography process is used to compute the coordinate-distance error between the ground-truth coordinates of the CASPR image and the estimated coordinates. The distance error is calcu-




Figure 4. The average maximum number of matches across all images in the query set, per-method, for 15-m and 10-m GRD in SentinelHub and Google Earth Engine, respectively.




Figure 5. The minimum summed distance error across all images in the query set, per-method, for 15-m and 10-m GRD in SentinelHub and Google Earth Engine, respectively.

lated as a summation of the Euclidean distances between four longitudinal and latitudinal coordinates of the proposed region and four ground-truth coordinates that are manually acquired. We consider the difference between the coordinates of the proposed location and the ground-truth. To account for the skew of an image, the error calculated between each corner are then added together. Namely,

$$\text{distance error} = \sum_{i=1}^4 ((x_i - x'_i) + (y_i - y'_i))^2$$

where  $(x_i, y_i)$  is the proposed set of coordinates and  $(x'_i, y'_i)$  are the ground-truth longitude and latitude coordinates. After the mapping procedure is finished, images can then be compared to existing imagery in Google Earth as well as accurate road-network and point-of-interest databases. Since the purpose of this experiment is to pinpoint the location of an image on a map, the longitude and latitude coordinates of the location are considered the primary metric as opposed to physical distance error. A t-test is performed to test significance between datasets.

**Results.** This results section presents key quantitative metrics on alignment accuracy and a comparison of feature-extraction methods. Results from different GRDs of map source imagery are also contrasted. Successes, limitations, and applications of the pipeline are discussed.

*Alignment Accuracy.* To determine the success of the image-registration methods across multiple configurations, we varied the imagery data source and, in the case of SentinelHub, query-image GRD. We conduct an analysis of the performance of each method as a combination of the total number of matches and error as the distance between ground-truth coordinates and coordinates proposed by each model. Both the total number of matches and distance error between proposed corner coordinates and

ground-truth coordinates were recorded for each registration attempt. Imagery from both SentinelHub and Google Earth Engine are tested for each method. SentinelHub data comes from the Sentinel-2 Level-1C data product. Level-1C pixel values are provided as top-of-atmosphere reflectance, and imagery is orthorectified, 15-m GRD. Query data from Google Earth Engine comes from the harmonized Sentinel-2 Level-2A data product, which is a processed version of associated Level-1C imagery. Level-2A pixels are bottom-of-atmosphere reflectance values with orthorectified, 10-m GRD sampling.

The distribution of the maximum number of matches for each query set can be seen in Figure 4. The number of matches is higher for AKAZE, LoFTR, and SIFT with Google Earth Engine’s Sentinel-2 Level-2A imagery. The results in Figure 5 show the median average distance error between the proposed coordinates and the ground-truth coordinates. For the AKAZE algorithm, the distance error for Google Earth Engine is significantly lower than that of the SentinelHub dataset ( $p < 0.05$ ), and thus delivers better accuracy. All other algorithms do not show significant difference in error between the datasets. If the noise in the algorithm is high due to a large number of matches, the mapping accuracy decreases, which highlights the importance of match quality over match quantity. Figures 6 and 7 show the difference in performance of each algorithm when the SentinelHub data is changed from 15-m to 30-m GRD.

The number of outliers and large variance in Figures 4 and 6 can be partially attributed to the large variance in our CASPR dataset. The images range from urban environments, such as Osaka, Japan, to more remote mountainous terrain, such as Seron, Chile. The wide range of imagery types leads to the large variance in number of matches, as all methods find features easily in urban environments with sharp edges and straight lines. LoFTR finds more features than other methods since it has the



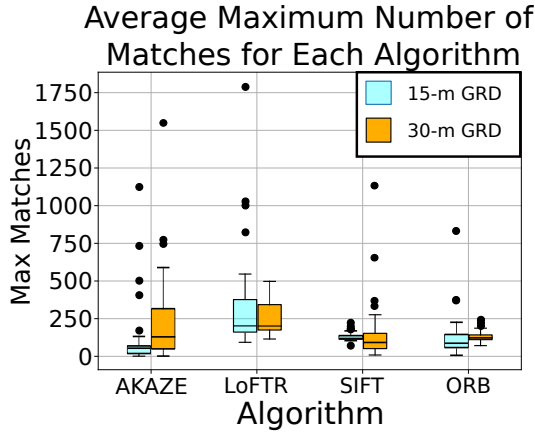


Figure 6. The average maximum number of matches across all images in the query set using SentinelHub, per-method. GRD varies from 15-m to 30-m.

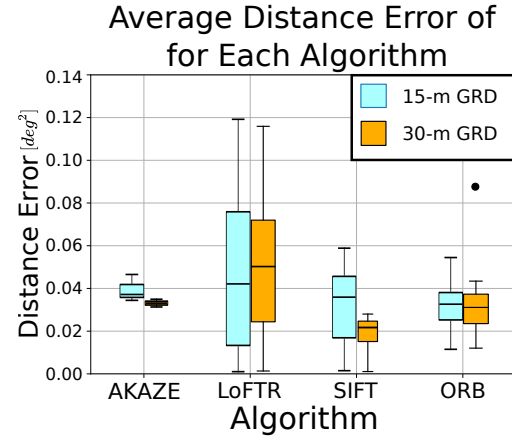


Figure 7. The minimum summed distance error across all images in the query set using SentinelHub, per-method. GRD varies from 15-m to 30-m.

ability to extract features in low-texture regions.

Figure 7 shows the distance error in 30-m GRD imagery and in 15-m GRD imagery. AKAZE, in particular, shows a significant reduction in the distance error, with  $p < 0.001$  for 30-m GRD compared to 15-m GRD. This follows from the higher average amount of features matched at 30-m GRD compared to 15-m GRD in the AKAZE dataset. This is most likely due to the higher-level, more general features that are found at a higher GRD. All other datasets do not demonstrate any statistically significant difference between 15-m GRD and 30-m GRD. More samples would have to be collected to make any claims of significance.

The high number of matches in LoFTR contribute to a larger amount of variance in distance error, but this can be attributed to matching in low-texture regions. Of significant note is our use of LoFTR’s pretrained outdoor model without fine-tuning. That is, we take the LoFTR model and apply it directly to satellite imagery to achieve accurate feature matching on the order of classic feature-matching algorithms. It should be noted that the input for LoFTR is a downsampled 640x480 image, which could be considered an unfair comparison of its capability by reducing the effective GRD. However, the basis of LoFTR is a coarse-to-fine granular approach to feature extraction, which makes the difference in resolution less critical. LoFTR is shown to have a very high variance in Figures 5 and 7, demonstrating that its accuracy is heavily dependent on the input. It is suspected that this high variance is due to the downsampling of the imagery performed in LoFTR’s pipeline. Additionally, the number of matches in low-texture regions is significantly different when using LoFTR compared to other methods. Qualitatively, it can be noted that these low-texture region matches tend to be less conducive to accurate homography.

Overall, each algorithm performed well in mapping CASPR imagery to known datasets. The distance errors



Figure 8. A comparison of two images of Mamutaa from CASPR (left) and Google Earth (right) depicts immense change over nearly six years.

in the coordinates calculated with each algorithm are not significantly different from each other, with median errors occurring around  $0.04 \text{ deg}^2$  for both 15-m and 10-m GRD of SentinelHub and Google Earth, respectively. The lowest observed median error can be seen in Figure 5 by AKAZE for Google Earth Engine at around  $0.02 \text{ deg}^2$ .

**Temporal Analysis.** These efforts have enabled higher quality and more rapid research application of CASPR imagery. An area of immense interest is temporal analysis for change detection and characterization. Perhaps the best example of this thus far is an image CASPR captured of Maamutaa Island in the Maldives. At the time of this writing, the most recent image of Maamutaa Island in Google Earth, seen in Figure 8, was taken on April 7th, 2016.

The CASPR Maamutaa image was taken on February 25, 2022, nearly six years later. Manual geolocation of this

image was rather difficult due to the significant change to the landmass. The once-forested island has been cleared and is now covered with artificial structures. Observing such change is a primary goal of CASPR’s high-resolution, low-GRD, and small form factor iSIM-90 imager. This pipeline enables such observation and evaluation to be automated.

**Future Research.** The capabilities of this pipeline make STP-H7-CASPR imagery more actionable and valuable. With near-immediate mapping of imagery to precise locations, additional applications can be deployed to characterize and evaluate change. These methods could range from simpler, more classical options like delta analysis to more modern techniques like semantic segmentation through deep learning. Stacks of analysis-ready, geolocated imagery at frequently revisited locations could enable long-term assessment of area characteristics.

Since these feature-matching methods show promise of finding features across multiple images, this methodology can also be applied to astronomical imagery. This form of mapping can be used for image registration across multiple astronomical datasets to augment captures with previous data, providing a higher signal-to-noise ratio. By using the pipeline shown in this paper with astronomical imagery, stellar data can be stacked across datasets and across various spectral filters.

The model used to find a homography with the LoFTR method was provided by the authors of the LoFTR paper, and was pretrained on an outdoor scene dataset. This method directly applied this model to satellite imagery with a very different sensor modality and feature space. By fine-tuning the model on satellite imagery, higher degrees of geolocation accuracy could likely be achieved. The absence of fine-tuning in our experiments is likely a reason for the large variation in the measured LoFTR distance error.

**Conclusions.** This research demonstrates a rapid and effective means of satellite-image geolocation. This method increases the capability of STP-H7-CASPR for Earth observation. Upon downlink, an image can be placed within a  $0.04 \text{ deg}^2$  global coordinate accuracy in under one minute using accessible existing tools. LoFTR, with a network pretrained on outdoor scenes, was shown to perform well on satellite imagery with a different feature space, suggesting that space-based LoFTR is feasible for feature-matching. Our pipeline enables further downstream tasks, such as imagery comparison, point-of-interest tracking, and temporal-change analysis. These efforts also open the door for future research in autonomous change characterization powered by deep-learning methods.

**Acknowledgements.** This research was supported by the NSF Center for Space, High-Performance, and Resilient Computing (SHREC) industry and agency members and by the IUCRC Program of the National Science Foundation under Grant No. CNS-1738783. The authors would like to thank the CASPR development team, STP,

the Air Force Research Lab Space Vehicles Directorate for their DoD sponsorship of CASPR, and other SHREC students and member organizations.

## References.

- [1] S. Roffe, T. Schwarz, T. Cook, N. Perryman, J. Goodwill, E. Gretok, A. Phillips, M. Moran, T. Garrett, and A. George, “CASPR: Autonomous sensor processing experiment for STP-H7,” in *Proceedings of the 34th Annual AIAA/USU Conference on Small Satellites*. AIAA, pp. 1–11, 2020.
- [2] R. Guzmán, R. López, E. Ocerin, S. Davis, J. Hernani, R. Brennan-Craddock, N. Kellerman, M. Pastena, N. Mellega, and F. Mariani, “A compact multispectral imager for the MANTIS mission 12U CubeSat,” p. 5, 08 2020.
- [3] A. V. Etten, D. Lindenbaum, and T. M. Bacastow, “SpaceNet: A Remote Sensing Dataset and Challenge Series,” *CoRR*, vol. abs/1807.01232, 2018.
- [4] N. Laban, B. Abdellatif, H. M. Ebeid, H. A. Shedeed, and M. F. Tolba, “Seasonal multi-temporal pixel based crop types and land cover classification for satellite images using convolutional neural networks,” in *2018 13th International Conference on Computer Engineering and Systems (ICCES)*, pp. 21–26, 2018.
- [5] M. Weinmann and U. Weidner, “Land-Cover and Land-Use Classification Based on Multitemporal Sentinel-2 Data,” in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, pp. 4946–4949, 2018.
- [6] M. Hachicha, M. Louati, A. Kallel, and J.-P. Gastellu-Etchegorry, “Prediction of plant growth based on statistical measurements using satellite image time series,” in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, pp. 136–139, 2020.
- [7] P. Mateus, J. Catalão, G. Nico, and P. Benevides, “Mapping Precipitable Water Vapor Time Series From Sentinel-1 Interferometric SAR,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 2, pp. 1373–1379, 2020.
- [8] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, “LoFTR: Detector-Free Local Feature Matching With Transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8922–8931, June 2021.
- [9] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore, “Google Earth Engine: Planetary-scale geospatial analysis for everyone,” *Remote Sensing of Environment*, 2017.
- [10] Sentinel Hub, Sinergise Ltd., “Modified Copernicus Sentinel data [2022]/Sentinel Hub,” 2022.
- [11] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.
- [12] B. Shupp, “Real-Time ISS Location.”
- [13] Google Earth 7.3, “San Carlos, Panama  $8^{\circ}26'34.36''\text{N}$   $79^{\circ}57'12.90''\text{W}$ ,” July 2022.
- [14] E. M. Sims, “The Department of Defense Space Test Program: Come Fly with Us,” in *2009 IEEE Aerospace conference*, pp. 1–6, IEEE, 2009.

- [15] J. G. Masek, M. A. Wulder, B. Markham, J. McCorkel, C. J. Crawford, J. Storey, and D. T. Jenstrom, "Landsat 9: Empowering open science and applications through continuity," *Remote Sensing of Environment*, vol. 248, 2020.
- [16] D. P. Roy, M. A. Wulder, T. R. Loveland, W. C.E., R. G. Allen, M. C. Anderson, D. Helder, J. R. Irons, D. M. Johnson, R. Kennedy, T. A. Scambos, C. B. Schaaf, J. R. Schott, Y. Sheng, E. F. Vermote, A. S. Belward, R. Bind-schadler, W. B. Cohen, F. Gao, J. D. Hipple, P. Hostert, J. Huntington, C. O. Justice, A. Kilic, V. Kovalsky, Z. P. Lee, L. Lyburner, J. G. Masek, J. McCorkel, Y. Shuai, R. Trezza, J. Vogelmann, R. H. Wynne, and Z. Zhu, "Landsat-8: Science and product vision for terrestrial global change research," *Remote Sensing of Environment*, vol. 145, 2014.
- [17] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort, *et al.*, "Sentinel-2: ESA's optical high-resolution mission for GMES operational services," *Remote sensing of Environment*, vol. 120, pp. 25–36, 2012.
- [18] S. Cantrell, J. Christopherson, C. Anderson, G. L. Stensaas, S. N. R. Chandra, M. Kim, and S. Park, "System characterization report on the WorldView-3 Imager," tech. rep., US Geological Survey, 2021.
- [19] W. S. Scott, N. Anderson, and A. Q. Rogers, "Design drivers for a viable commercial remote sensing space architecture," 2020.
- [20] A. E. Frazier and B. L. Hemingway, "A technical review of Planet SmallSat data: Practical considerations for processing and using PlanetScope imagery," *Remote Sensing*, vol. 13, no. 19, p. 3930, 2021.
- [21] A. Collison and A. Jumpasut, "Improving the Radiometric Calibration of the Heterogeneous Planet Dove Fleet," 2021.
- [22] Planet Labs PBC, "PlanetScope," 2021.
- [23] M. Kim, S. Park, A. Sampath, C. Anderson, and G. L. Stensaas, "System characterization report on Planet SkySat," tech. rep., US Geological Survey, 2022.
- [24] M. Longanbach and L. McGill, "Scaling fleet operations: The growth and results of SkySat mission operations," in *2018 SpaceOps Conference*, p. 2706, 2018.
- [25] Planet Labs PBC, "Pelican — Planet," 2021.
- [26] Planet Labs PBC, "Planet Basemaps: Comprehensive, High-Frequency Mosaics for Analysis — Planet," 9 2022.
- [27] Planet Labs PBC, "Planet Monitoring - Satellite Imagery and Monitoring — Planet," 9 2022.
- [28] Q. Zhao, L. Yu, X. Li, D. Peng, Y. Zhang, and P. Gong, "Progress and Trends in the Application of Google Earth and Google Earth Engine," *Remote Sensing 2021, Vol. 13, Page 3778*, vol. 13, p. 3778, 9 2021.
- [29] Planet Labs PBC, "Planet Google Earth Engine Delivery Overview," 9 2022.
- [30] G. Milcinski, M. Batic, M. Kadunc, P. Kolaric, R. Mocnik, *et al.*, "Sentinel-2 Services Library-efficient way for exploration and exploitation of EO data," in *EGU General Assembly Conference Abstracts*, p. 19502, 2017.
- [31] R. Lamb, R. Pieschke, and K. Lemig, "Remotely Sensed Land Imagery and Access Systems: USGS Updates," in *AGU Fall Meeting Abstracts*, vol. 2017, pp. IN23E–05, 2017.
- [32] R. Longhenry, T. Sohre, R. McKinney, and T. Mentele, "USGS Earth Explorer Client for Co-Discovery of Aerial and Satellite Data," in *AGU Fall Meeting Abstracts*, vol. 2011, pp. IN51B–1584, 2011.
- [33] D. Phiri, M. Simwanda, S. Salekin, V. R. Nyirenda, Y. Murayama, and M. Ranagalage, "Sentinel-2 data for land cover/use mapping: a review," *Remote Sensing*, vol. 12, no. 14, p. 2291, 2020.
- [34] T. Zheng, M. H. Bergin, S. Hu, J. Miller, and D. E. Carlson, "Estimating ground-level PM2.5 using micro-satellite images by a convolutional neural network and random forest approach," *Atmospheric Environment*, vol. 230, 2020.
- [35] J. Park, J. Lee, K. Seto, T. Hochberg, B. A. Wong, N. A. Miller, K. Takasaki, H. Kubota, Y. Oozeki, S. Doshi, M. Midzik, Q. Hanich, B. Sullivan, P. Woods, and D. A. Kroodsmas, "Illuminating dark fishing fleets in North Korea," *Science Advances*, vol. 6, no. 30, 2020.
- [36] R. Hänsch, J. Arndt, D. Lunga, M. Gibb, T. Pedelose, A. Boedihardjo, D. Petrie, and T. M. Bacastow, "SpaceNet 8 - The Detection of Flooded Roads and Buildings," 2022.
- [37] D. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, pp. 91–, 11 2004.
- [38] K.-Z. Liu, P.-J. Lee, G.-C. Xu, and B.-H. Chang, "SIFT-Enhanced CNN Based Objects Recognition for Satellite Image," pp. 1–2, 2020.
- [39] C. Joshi and S. Mukherjee, "Empirical analysis of SIFT, Gabor and fused feature classification using SVM for multispectral satellite image retrieval," pp. 1–6, 2017.
- [40] R. Mehrotra, K. Namuduri, and N. Ranganathan, "Gabor filter-based edge detection," *Pattern Recognition*, vol. 25, no. 12, pp. 1479–1494, 1992.
- [41] Y. Xiang, F. Wang, and H. You, "OS-SIFT: A Robust SIFT-Like Algorithm for High-Resolution Optical-to-SAR Image Registration in Suburban Areas," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 6, pp. 3078–3090, 2018.
- [42] P. Fernández Alcantarilla, "Fast explicit diffusion for accelerated features in nonlinear scale spaces," 09 2013.
- [43] M. Elantcev, I. Arkhipov, and R. Gafarov, "The modified method of statistical differentiation for matching of an aerial photograph and a satellite image," in *2020 International Conference on Information Technology and Nanotechnology (ITNT)*, pp. 1–5, 2020.
- [44] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," pp. 2564–2571, 2011.
- [45] L. Li, L. Han, M. Ding, Z. Liu, and H. Cao, "Remote sensing image registration based on deep learning regression model," *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.



- [46] D. Demchev, V. Volkov, E. Kazakov, and S. Sandven, "Feature tracking for sea ice drift retrieval from SAR images," in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 330–333, 2017.
- [47] J. Zhao, D. Yang, Y. Li, P. Xiao, and J. Yang, "Intelligent matching method for heterogeneous remote sensing images based on style transfer," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 6723–6731, 2022.
- [48] Google Earth 7.3, "Maamutaa, Maldives 0°50'53.49"N 73°26'29.43"E," July 2022.
- [49] M. Muja and D. G. Lowe, "Scalable nearest neighbor algorithms for high dimensional data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 11, pp. 2227–2240, 2014.