The Generation of TGO-CaSSIS Stereo Products after six years of scientific mission: The Complete Stereo Chain of INAF

Adriano Tullo [1], Cristina Re [1], Emanuele Simioni [1], Silvia Bertoli [1], Riccardo La Grassa [1], Gabriele Cremonese [1], Nicolas Thomas [2]

[1] INAF-Astronomical Observatory of Padova, Vicolo dell'Osservatorio 5, 35122 Padova, Italy (adriano.tullo@inaf.it)
[2] Physikalisches Institut, Sidlerstr. 5, University of Bern, CH-3012 Bern, Switzerland

Keywords: Mars, CaSSIS, Photogrammetry, 3DPD, 3D Reconstruction, Software updates.

Abstract

CaSSIS (Colour and Stereo Surface Imaging System) on board the ExoMars Trace Gas Orbiter, which has been capturing images of Mars since April 2018, is a multispectral stereo camera and its stereo and colour data are used for generating high-resolution, three-dimensional maps and orthorectified images of the Martian surface. The instrument has already taken more than 50.000 images and covering >7.6% of Mars' surface, producing a total of 23.8 Tbit of data, with a small percentage (4.8%) taken using colour filters. These data are processed with the 3DPD software, which creates DTMs and orthorectified images with a ground resolution of 13.5 m per pixel and vertical accuracy of around 8 meters. All the stereo products are available to the scientific community through the INAF-hosted repository. The 3DPD software pipeline includes various stages: pre-processing to remove optical distortions and mosaicking, a matching core and DTM interpolation routines. A recent improvement through the integration of Bundle Adjustment, using the jigsaw tool from the USGS ISIS suite, has enhanced the precision of the data, reducing errors in the final 3D products. Overall, the work on CaSSIS and 3DPD plays a significant role in advancing Mars exploration, producing high-quality topographic maps and spatial data critical for scientific research.

1. Introduction

The Colour and Stereo Surface Imaging System (CaSSIS) is an innovative stereo camera designed for the ExoMars Trace Gas Orbiter mission, orbiting around Mars since April 2018.
The CaSSIS camera can produce detailed 3D maps and capture images in four different colour bands, thus providing crucial data for the analysis of the Martian surface and its composition.
The CaSSIS stereo configuration is possible thanks to an innovative telescope rotation system, which creates a convergence angle between images of the stereo pairs of approximately 22°. Acquiring from a circular orbit about 400 km from the Martian surface, the CaSSIS ground resolution averages 4.5 m/px, resulting in digital stereogrammetric digital terrain models (DTMs) at about 13.5 m/px (Fig.1).

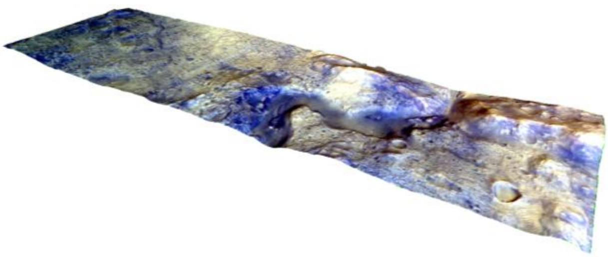

Figure 1. 3D representation of a CaSSIS DTM and composite RGB image of related orthorectified colour filters, in detail NIR-PAN-BLU. The stereopair, MY37_028515_017_1/2, depicts the western boundary of the Jezero Crater, in the region of Syrtis Major and Isidis Planitia, Mars.

The 3DPD software, introduced by Re at al., (2019), is the baseline for the generation of the Digital Terrain Models (DTMs) and the orthorectified images, and it is the official CaSSIS DTMs creation tool.
Since the definition of the framework, described in detail by Simioni et al. (2021), the software continues to be updated to improve the performance and quality of the stereo data product generation (Re et al., 2022).
The development team located at the Astronomical Observatory of Padova, INAF, is responsible for generating and archiving the stereo products from CaSSIS data and, since the beginning of the scientific mission, making them available to the entire team and scientific community.
To date, more than 2200 stereo couples are available on a total of 50.000 images acquired, and more than 400 DTMs and orthoimages have been produced and made available in the OAPD-hosted repository (https://cassis.oapd.inaf.it/archive/) (Fig.2).

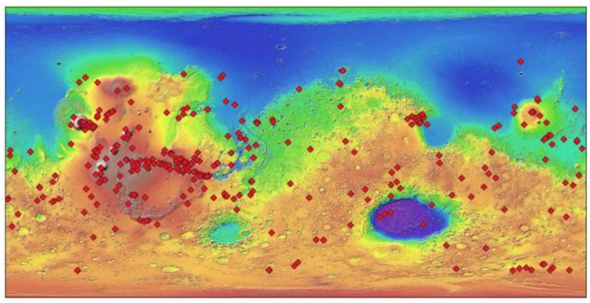

Figure 2. Global distribution of DTMs generated so far, available at the INAF-OAPD archive (https://cassis.oapd.inaf.it/archive/).

This paper presents recent developments and the pipeline update with the additional techniques adopted that enable the generation of highly accurate stereo products. The quality and accuracy of the DTMs are also discussed to make the possible sources of errors comprehensive and mitigable.

https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-1471-2025

2. The 3DPD pipeline

The developed pipeline permits the generation of stereo products from radiometrically calibrated CaSSIS frames by comprehending several processing steps, summarised in the diagram in Fig.3 and described in detail below.

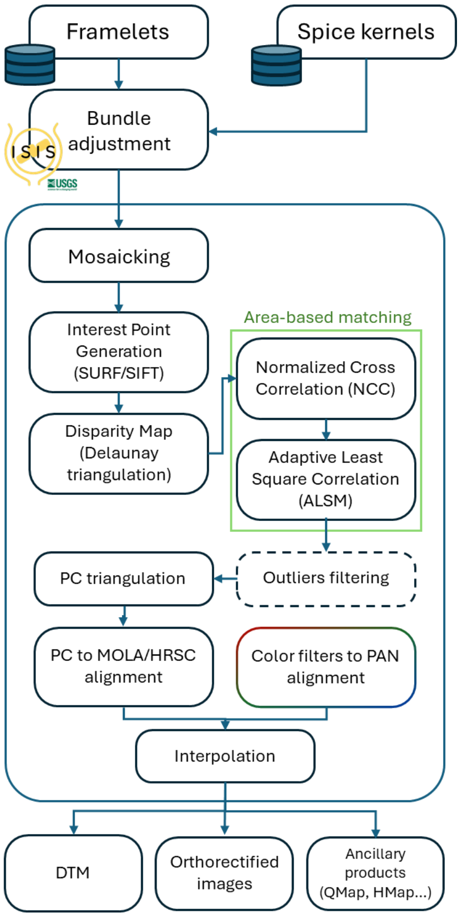

Figure 3. The 3DPD workflow, starting from the radiometrically calibrated CaSSIS framelets and kernels to the final products (updated from Re et al., 2022).

2.1 Bundle adjustment and pre-processing

The three-dimensional positioning and attitude data of the individual framelets are extracted from SPICE kernels (SK) as well as the interior camera parameters. The projection matrices required for the triangulation process are then estimated.
In the last implementation of 3DPD, the possibility to take advantage also from a Bundle Block Adjustment has been included, employing the *jigsaw* routines of the USGS ISIS platform (Laura et al., 2023).
To do this, a compatibility layer was created to integrate the outputs of the bundle already in the mosaic phase. Thanks to *jigsaw*, the projection matrices can be refined from the misalignments introduced by jitter effects or more generally from errors in pointing accuracy, which would otherwise affect the DTMs with distortions and geometrical artefacts.
For example, the misalignments between each push-frame acquisition, which have not been adequately modelled and resolved by the Bundle Adjustment could commonly result in steps in the DTM of even hundreds of m in the worst cases (Fig.4).

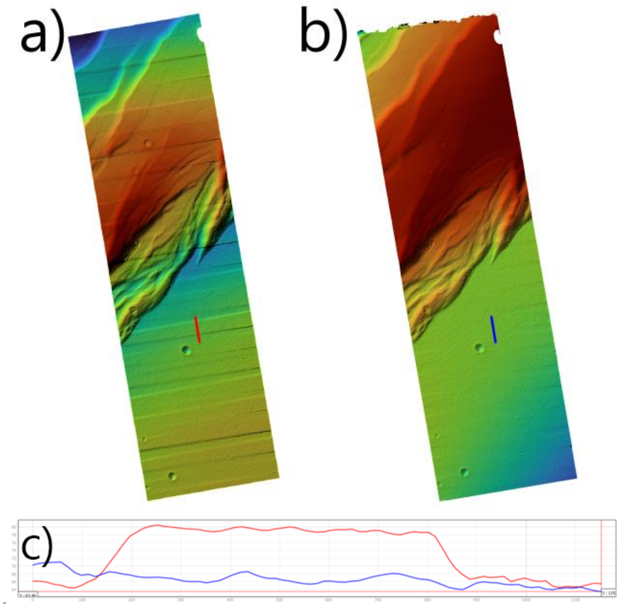

Figure 4. The mosaicking of the CaSSIS data includes also the improvement of the framelets camera pointing information through ISIS *Jigsaw* Bundle Adjustment. This can prevent the formation of elevation artifacts that can reach up to tens or hundreds of meters in elevation. DTMs extracted from MY36_020366_190 stereo pair a) without and b) with bundle adjustment. In c) an elevation profile of an artefact by comparison.

Subsequently, the pre-processing phase includes the various steps to make the radiometric-corrected framelets suitable for ingestion into the matching core.
This encompasses the geometrical correction phases, by removing any optical distortions inherent to CaSSIS, and then the mosaicking. Any residual misalignments between framelets are further corrected with a Least Squares adjustment that takes advantage of the overlapping portion between them. A Gaussian low-pass filter is also applied to improve the quality of the image textures, resulting in a more consistent matching further reducing the possible mismatches.
Mosaic images are, therefore, pre-rectified by projecting them on a plane that is initialized from the elevation of a MOLA-based DTM and then corrected exploiting the minimizing of the Normalized Cross Correlation (NCC) (Ching, 1995) in the overlapping portions of the images where parallax is negligible. The minimization is done using a nonlinear square approach based on Levenberg-Marquardt (Marquardt, 1963).
Although not used directly for photogrammetric processing, the same procedure is also repeated for the CaSSIS colour filters (BLUE, RED and NIR), so that they can subsequently be aligned to the panchromatic (PAN) and orthorectified.

2.2 Feature identification and disparity map

Feature detection for the search of interest points is approached using two different algorithms. The Speeded Up Robust Feature (SURF) algorithm (Bay et al., 2008) is first applied by dividing the images by tiles, to distribute more evenly the net across the images. Subsequently, a further search, recently introduced, using the Scale-Invariant Feature Transform (SIFT) (Lowe, 1999) allows the network of points to be further densified, if necessary, focusing on areas that are less well covered by SURF.

The result of the application of this process is a list of matched interest points on both images of the stereo pair, which represents the first sparse disparity map. The information is then expanded to the whole image using Delaunay triangulation, resulting in a TIN surface of parallaxes that is used as initialization for the subsequent matching process.

2.3 Dense Matching phase

The dense matching phase is employed through a multi-scale system, dividing the process into several image pyramids of increasing resolution. This develops a coarse-to-fine approach to reduce the search area and to limit blunders.
The correlation initialization is executed with the NCC algorithm starting from the first level of the image pyramid. Then the matched points of the lower level are transferred to the higher resolution level where the disparity of the additional points is predicted from the neighbourhoods with a bilinear interpolation. This is repeated until the last level of the pyramid (higher resolution), where a sub-pixel refinement by area-based matching (Least-Squares Matching) is finally applied (Gruen, 1985).
Based on the positions and attitudes of the sensor and the geometric properties of the camera, the high dense set of corresponding points can be converted from the image coordinates (in pixels) to 3D object coordinates by means of forward ray intersection, that derive a contiguous model of the surface topography. The resulting point cloud (PC), at heterogeneous density due to the perspective view, is then interpolated on a regular grid of height values in a reference coordinate system.
The process achieves a vertical accuracy up to 8 m (Fig.5), in accordance with the validation of Re et al. 2022 and the best horizontal accuracy is considered up to 13.5 m/px, i.e. 3 CaSSIS GSD (Ground Sample Distance).
Each 3D point of the PC is referenced with a set of coordinates, Latitude/Longitude/Height, which are further refined in the subsequent alignment phase with a global reference DTM before the interpolation of the final products.

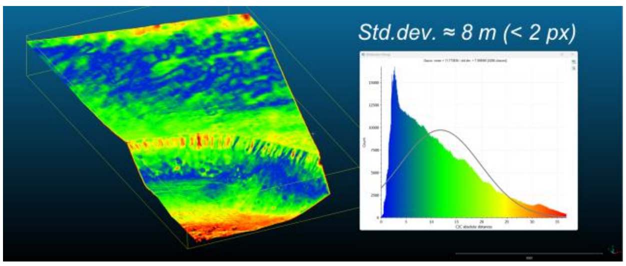

Figure 5. Cloud-to-cloud distance between a CaSSIS DTM (MY34_003673_018) and a HiRISE DTM (DTEEC_005533_1975_005388_1975) at 1 m/px, demonstrating a vertical accuracy in the overlapping portion of about 8 m.

2.4 MOLA/HRSC alignment and interpolation

The ground projection of remote sensing data, particularly in the planetary domain, still maintain a certain degree of uncertainty for several reasons, including inaccuracies in satellite data, jittering, and difficulties in attitude and calibration. Such uncertainties tend to lead to inaccuracies in projection and absolute height information. To mitigate this effect, an alignment step (Fig.6) to a lower resolution global DTM was introduced in the 3DPD, which is the second version of the hybrid MOLA/HRSC product introduced by Fergason et al., 2018 at 200 m/px.

Of the global DTM, the portion of interest is first extracted, estimated from the initial extent of the CaSSIS PC plus a surrounding margin. An initial coarse alignment is made by correcting the PC elevation coordinates by fitting a plane equation on the vertical differences between the PC and the reference DTM using a least-squares solution. This mitigates the presence of any generalised elevation errors or tilt trends, ensuring the intersection between the two data for the subsequent refinement. The fine co-registration is achieved through an implementation of the Iterative Closest Point (ICP) algorithm (Besl and McKay, 1992; Glira et al., 2015). The ICP iteratively estimates a 6-degree-of-freedom transformation (as combination of translations and rotations) to minimise errors. These errors are estimated as the sum of squared differences between the coordinates of the matched pairs.

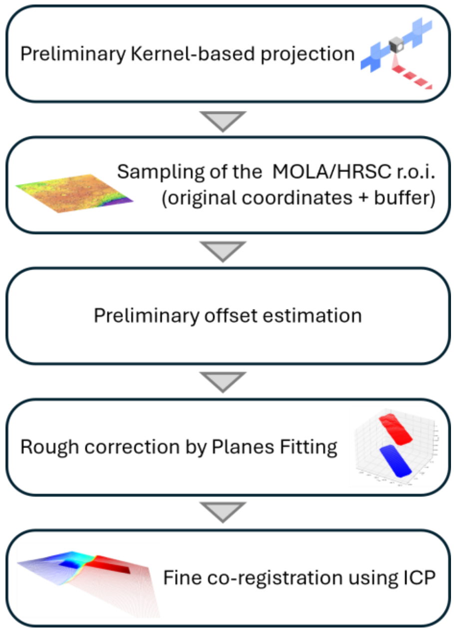

Figure 6. The 3DPD alignment workflow using a two-step (coarse-to-fine) co-registration approach.

The alignment resulting from this process results in an overall accurate vertical positioning within 50 m of the global elevation (below the single spatial resolution pixel of the reference DTM) (Fig.7). This process was also extended to the entire database of DTMs already produced as one major update.

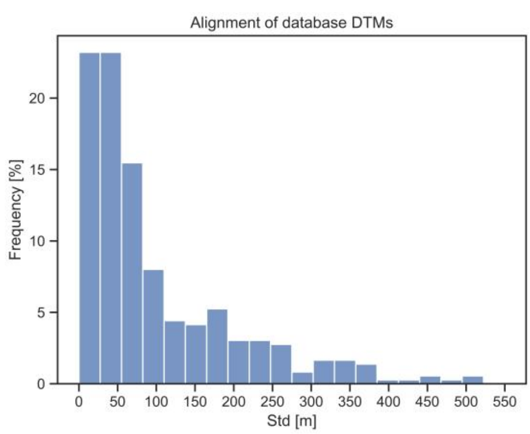

Figure 7. The standard deviation, expressed in meters, regarding the vertical co-registration accuracy of the CaSSIS DTMs database with the MOLA-HRSC.

2.5 Orthorectification of CaSSIS filters

The corrected point cloud is finally interpolated into DTM according to different raster formats through a bilinear interpolation. The spectral component is also interpolated on the elevation data, resulting in an accurate orthorectification of the PAN pair and the other colour filters available aligned to them (Fig.8). The alignment of these takes place prior to orthorectification through a homographic projection estimated by matched interest points identified with SIFT.
Bilinear interpolation is also adopted for filters for the minimum impact of alteration of the spectral component, as well as for performance due to the simplicity of calculation.

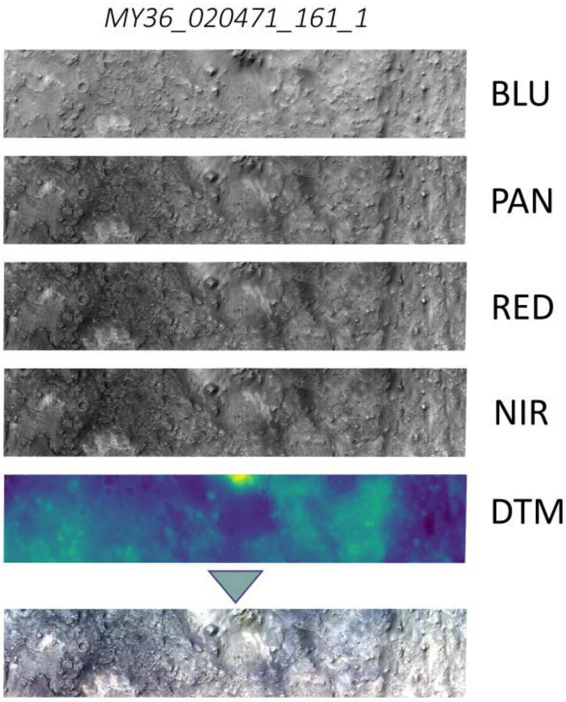

Figure 8. The individual colour filters of the MY36_020471_161_1 acquisition orthorectified on the DTM (in coloured scale). On the right a final composite RGB (NPB) of the image.

The recent update of orthorectification of filters enables a new level in the colour exploitation of CaSSIS.
Alignment with topographic information simplifies and makes photometric correction more accurate (Munaretto et al., 2021) and enables its application in spectral analysis studies (Tornabene et al., 2018; Rangarajan et al., 2023).

3. Georeferencing/alignment accuracy

Absolute spatial referencing accuracy was measured by comparison with orthorectified HRSC mosaics, which can be up to 12.5 m/px (Gwinner et al., 2010).
For this analysis, a series of points of interest were identified and matched using the SIFT computer vision algorithm (Fig.9). The Euclidean distance between the coordinates of the common points between the images is uniquely influential and indicative of the georeferencing accuracy resulting from the alignment.

The comparison study was conducted on a sample of 10 CaSSIS orthoproducts, from different areas of Mars including Oxia Planum in Arabia Terra, Jezero Crater in Syrtis Major and some areas in the Noachis Terra region. The HRSC images used are the orthorectified level 5 mosaics (referred as nd5) as described in Gwinner et al., 2016 and Michael et al. 2016.
The results, summarized in Fig.10, show that the projection accuracy reach values even lower than a CaSSIS pixel, and on average about 2-3 pixels.

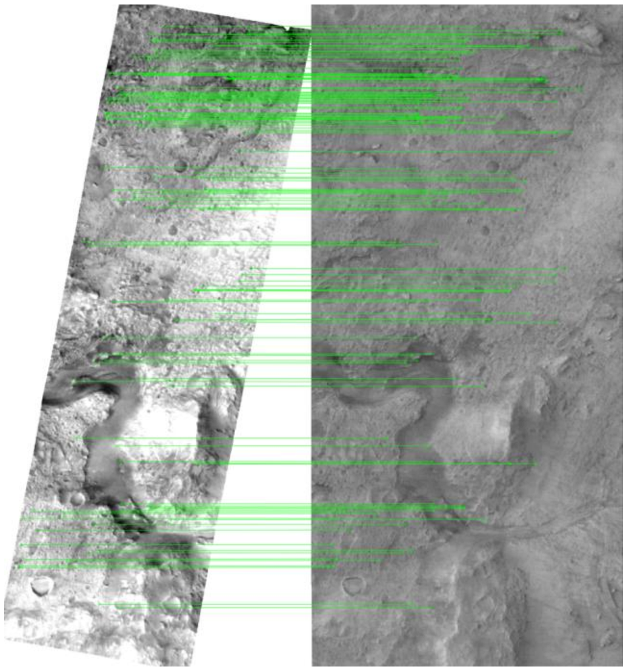

Figure 9. Interest point matching between orthorectified CaSSIS (left) and part of a HRSC nd5 mosaic (Michael et al. 2016) (right). The euclidean distance between the matching points gives us an indication of the projection accuracy overall.

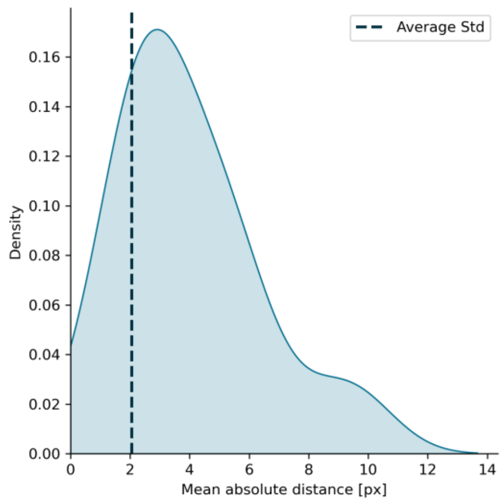

Figure 10. Mean and standard deviation of the absolute euclidean distance of matching interest points between a sample of 10 CaSSIS orthorectified images and the orthorectified HRSC nd5 basemap, expressed in HRSC pixels (12.5 m/px).

https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-1471-2025

4. Conclusions

The CaSSIS camera on board the ExoMars TGO allows the creation of four-band stereo products, which are exploited for 3D reconstruction thanks to the 3DPD software operated by the INAF OAPd team. The software, which manages the various parts of the creation of DTMs, has been under continuous development since 2019. Recent developments include the integration of Bundle Adjustment which allows for the correction of inaccuracies related to kernel errors that can result in artefacts of up to several tens of meters vertically. In addition, the introduction of a second computer vision algorithm (SIFT) makes it possible to densify the search for tie points preparatory to the creation of the initial disparity maps, resulting in a further increase in stability in the creation of DTMs and overall quality, particularly in conditions where SURF approach find fewer interest points. Finally, the alignment approach with the MOLA/HRSC through two co-registration steps allows for improved absolute quota information and georeferencing of products. The process has recently been extended to the entire processed database available online, resulting in a vertical difference to MOLA/HRSC in terms of standard deviation generally less than 100 m and in a very high georeferencing accuracy, averaging in range of 1 to 3 pixels. In addition, products include the addition of orthorectified colour filters. These represent new potential in the exploitation of CaSSIS data for geological and spectral analysis.

5. Future developments

Future developments include updates to the integration of Bundle Adjustment with elevation data from MOLA at the same preliminary stage, with the goal of a three-dimensional correction of positioning information and, thus, further increase in vertical accuracy.

Currently in development and validation is also the introduction of AI implementations at different stages of the Pipeline. These include the production of accurate dense image points, which would result in an evolution of current disparity maps and their completeness. In addition, modern applications of Monocular Depth Estimation (La Grassa et al., 2022) and an ever-growing database of DTMs suitable for training open the door to different possibilities for use, from more advanced preliminary co-registration to new types of DTM extraction.

Acknowledgements

The study has been supported by the Italian Space Agency (ASI-INAF agreement no. 2024-40-HH.0). Thanks to the CaSSIS Team for the precious advice and feedback.

References

Bay, H., Ess, A., Tuytelaars, T., Van Gool, L., 2008. Speeded-Up Robust Features (SURF). Computer Vision and Image Understanding, Similarity Matching in Computer Vision and Multimedia 110, 346–359. https://doi.org/10.1016/j.cviu.2007.09.014

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14, 239–256. https://doi.org/10.1109/34.121791

Ching, W.-S., 1995. Normalized cross-correlation: the contrast-dependent problem and its solution. JEI 4, 278–282. https://doi.org/10.1117/12.210718

Fergason, R.L., Hare, T.M., Laura, J., 2018. HRSC and MOLA blended digital elevation model at 200m v2. Astrogeology PDS Annex, U.S. Geological Survey.

Glira, P., Pfeifer, N., Briese, C., Ressl, C., 2015. A Correspondence Framework for ALS Strip Adjustments based on Variants of the ICP Algorithm. Photogrammetrie-Fernerkundung-Geoinformation 2015, 275–289.

Gwinner, K., Jaumann, R., Hauber, E., Hoffmann, H., Heipke, C., Oberst, J., Neukum, G., Ansan, V., Bostelmann, J., Dumke, A., Elgner, S., Erkeling, G., Fueten, F., Hiesinger, H., Hoekzema, N.M., Kersten, E., Loizeau, D., Matz, K.-D., McGuire, P.C., Mertens, V., Michael, G., Pasewaldt, A., Pinet, P., Preusker, F., Reiss, D., Roatsch, T., Schmidt, R., Scholten, F., Spiegel, M., Stesky, R., Tirsch, D., van Gasselt, S., Walter, S., Wählisch, M., Willner, K., 2016. The High Resolution Stereo Camera (HRSC) of Mars Express and its approach to science analysis and mapping for Mars and its satellites. Planetary and Space Science 126, 93–138. https://doi.org/10.1016/j.pss.2016.02.014

Gwinner, K., Scholten, F., Preusker, F., Elgner, S., Roatsch, T., Spiegel, M., Schmidt, R., Oberst, J., Jaumann, R., Heipke, C., 2010. Topography of Mars from global mapping by HRSC high-resolution digital terrain models and orthoimages: Characteristics and performance. Earth and Planetary Science Letters 294, 506–519.

La Grassa, R., Gallo, I., Re, C., Cremonese, G., Landro, N., Pernechele, C., Simioni, E., Gatti, M., 2022. An Adversarial Generative Network Designed for High-Resolution Monocular Depth Estimation from 2D HiRISE Images of Mars. Remote Sensing 14, 4619. https://doi.org/10.3390/rs14184619

Laura, J., Acosta, A., Addair, T., Adoram-Kershner, L., Alexander, J., Alexandrov, O., Alley, S., Anderson, D., Anderson, James, Anderson, Jeff, Annex, A., Archinal, B., Austin, C., Backer, J., … Young, A., 2023. Integrated Software for Imagers and Spectrometers. https://doi.org/10.5281/zenodo.7644616

Lowe, D.G., 1999. Object recognition from local scale-invariant features, in: Proceedings of the Seventh IEEE International Conference on Computer Vision. Presented at the Proceedings of the Seventh IEEE International Conference on Computer Vision, pp. 1150–1157 vol.2. https://doi.org/10.1109/ICCV.1999.790410

Marquardt, D.W., 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. Journal of the Society for Industrial and Applied Mathematics 11, 431–441. https://doi.org/10.1137/0111030

Michael, G.G., Walter, S.H.G., Kneissl, T., Zuschneid, W., Gross, C., McGuire, P.C., Dumke, A., Schreiner, B., van Gasselt, S., Gwinner, K., Jaumann, R., 2016. Systematic processing of Mars Express HRSC panchromatic and colour image mosaics: Image equalisation using an external brightness reference. Planetary and Space Science 121, 18–26. https://doi.org/10.1016/j.pss.2015.12.002

Munaretto, G., Pajola, M., Lucchetti, A., Re, C., Cremonese, G., Simioni, E., Cambianica, P., Thomas, N., 2021. Topographic correction of HiRISE and CaSSIS images: Validation and application to color observations of Martian albedo features. Planetary and Space Science 200, 105198. https://doi.org/10.1016/j.pss.2021.105198

Rangarajan, V.G., Tornabene, L.L., Osinski, G.R., Conway, S.J., Seelos, F.P., Silvestro, S., Salese, F., Pajola, M., Lucchetti, A., Munaretto, G., Bickel, V.T., Thomas, N., Cremonese, G., 2023. Change detection and monitoring of active Martian surface phenomena with the Colour and Stereo Surface Imaging System (CaSSIS) onboard the ExoMars Trace Gas Orbiter (TGO). Icarus 394, 115443. https://doi.org/10.1016/j.icarus.2023.115443

Re, C., Fennema, A., Simioni, E., Sutton, S., Mège, D., Gwinner, K., Józefowicz, M., Munaretto, G., Pajola, M., Petrella, A., Pommerol, A., Cremonese, G., Thomas, N., 2022. CaSSIS-based stereo products for Mars after three years in orbit. Planetary and Space Science 219, 105515. https://doi.org/10.1016/j.pss.2022.105515

Re, C., Tulyakov, S., Simioni, E., Mudric, T., Cremonese, G., Thomas, N., 2019. Performance evaluation of 3DPD, the photogrammetric pipeline for the CaSSIS stereo images. Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. XLII-2/W13, 1443–1449. https://doi.org/10.5194/isprs-archives-XLII-2-W13-1443-2019

Simioni, E., Re, C., Mudric, T., Cremonese, G., Tulyakov, S., Petrella, A., Pommerol, A., Thomas, N., 2021. 3DPD: A photogrammetric pipeline for a PUSH frame stereo camera. Planetary and Space Science 198, 105165. https://doi.org/10.1016/j.pss.2021.105165

Thomas, N., Cremonese, G., Ziethe, R., Gerber, M., Brändli, M., Bruno, G., Erismann, M., et al., 2017. The Colour and Stereo Surface Imaging System (CaSSIS) for the ExoMars Trace Gas Orbiter. Space Sci Rev 212, 1897–1944. https://doi.org/10.1007/s11214-017-0421-1

Tornabene, L.L., Seelos, F.P., Pommerol, A., Thomas, N., et al., 2017. Image Simulation and Assessment of the Colour and Spatial Capabilities of the Colour and Stereo Surface Imaging System (CaSSIS) on the ExoMars Trace Gas Orbiter. Space Sci Rev 214, 18. https://doi.org/10.1007/s11214-017-0436-7

Appendix

The CaSSIS DTM archive is currently available at: https://cassis.oapd.inaf.it/archive