Planetary and Space Science xxx (2017) 1–10

Contents lists available at ScienceDirect

Planetary and Space Science

journal homepage: www.elsevier.com/locate/pss

Automatic Coregistration and orthorectification (ACRO) and subsequent mosaicing of NASA high-resolution imagery over the Mars MC11 quadrangle, using HRSC as a baseline

Panagiotis Sidiropoulos [a,*], Jan-Peter Muller [a], Gillian Watson [a], Gregory Michael [b], Sebastian Walter [b]

[a] Mullard Space Science Laboratory, University College London, UK
[b] Freie Universitat Berlin, Germany

ARTICLE INFO

Keywords:
Mars
High-resolution imagery
Mars orbiters
Coregistration
Orthorectification
Mosaicing
MC11 quadrangle

ABSTRACT

This work presents the coregistered, orthorectified and mosaiced high-resolution products of the MC11 quadrangle of Mars, which have been processed using novel, fully automatic, techniques. We discuss the development of a pipeline that achieves fully automatic and parameter independent geometric alignment of high-resolution planetary images, starting from raw input images in NASA PDS format and following all required steps to produce a coregistered geotiff image, a corresponding footprint and useful metadata. Additionally, we describe the development of a radiometric calibration technique that post-processes coregistered images to make them radiometrically consistent. Finally, we present a batch-mode application of the developed techniques over the MC11 quadrangle to validate their potential, as well as to generate end products, which are released to the planetary science community, thus assisting in the analysis of Mars static and dynamic features. This case study is a step towards the full automation of signal processing tasks that are essential to increase the usability of planetary data, but currently, require the extensive use of human resources.

1. Introduction

The systematic high-resolution mapping of Mars that started more than 4 decades ago with Viking Orbiters (Soffen and Snyder, 1976) has generated more than 500 thousand image strips with resolution finer than 100 m/pixel, which in total cover an area that is larger than 6 times the area of Mars (Sidiropoulos and Muller, 2015a). These products are accompanied by position and orientation metadata, in the form of SPICE kernels (Acton, 1996), which are the best estimate of the true image coordinates with a limited accuracy, which can produce misalignments that may reach up to several kilometres (Sidiropoulos and Muller, 2016). This spatial accuracy has improved over time but it is still the case that for images coming from different instruments and also for images acquired by the same instrument, at high-resolution, Mars imagery consists of images that are geometrically inconsistent, a problem that is amplified by the possible use of different geodetic coordinate reference frames from different Mars missions.

This geometric misalignment of available high-resolution imagery prohibits pairwise image comparisons at pixel-level, while generating artifacts in the border of overlapping images. The former impedes the use of stacks of images mapping the same region to track small-scale surface changes (Wagstaff et al., 2012), (Di et al., Liu), (Xin et al., 2017) (Sidiropoulos and Muller, 2017), that may be associated with Mars dynamic features (such as new impact craters (Byrne et al., 2009), recurring slope lineae (McEwen et al., 2011), slope streaks (Schorghofer et al., 2007), etc.) while the latter significantly reduces the quality and functionality of Mars image mosaics. This misalignment is alleviated by the use of image coregistration, which is a family of computer vision techniques that identify common properties of the image content in order to establish a fully-defined geometric correspondence, i.e. the mapping of the one image (named target or input image hereafter) to the coordinates of the other (named reference or baseline image hereafter). While automatic coregistration has been extensively used for the geometric alignment of several types of images, in the case of high-resolution planetary images usually the most common approach is a tedious manual or semi-automatic coregistration (often using photogrammetric

* Corresponding author.
E-mail addresses: p.sidiropoulos@ucl.ac.uk (P. Sidiropoulos), j.muller@ucl.ac.uk (J.-P. Muller), g.watson@ucl.ac.uk (G. Watson), gregory.michael@fu-berlin.de (G. Michael), sebastian.walter@fu-berlin.de (S. Walter).

https://doi.org/10.1016/j.pss.2017.10.012
Received 18 July 2017; Received in revised form 5 September 2017; Accepted 17 October 2017
Available online xxxx

techniques such as bundle adjustment (Cooper and Cross, 1988)) which requires the extensive use of human resources, therefore being impractical for large volumes of data (Sidiropoulos and Muller, 2018).

In order to overcome volume limitations, as well as to minimize the required human resources, we have developed and implemented a pipeline that can perform batch-mode automatic coregistration and orthorectification of multi-instrument high-resolution imagery to a common baseline. The employed algorithm is fully automatic; firstly, it does not require any human involvement in any stage of the software execution, and secondly, it does not require any parameter tuning, i.e. it runs with the same set of parameters for all images from a single instrument. This property, along with the algorithm's speed and reliability allows the systematic coregistration in realistic time of large chunks of Martian imagery, potentially even the complete high-resolution dataset.

The production of two HRSC high-resolution mosaics from the HRSC team (Gwinner et al., 2016), covering the MC11 Mars quadrangle with a 12.5 m/pixel panchromatic nadir image and a 50 m/pixel DTM, has provided the opportunity to validate the novel algorithm in a substantially challenging setup, which includes more than 11,500 high-resolution orbital images from 3 NASA instruments, namely, Context Camera (CTX) (Bell et al., 2013) onboard Mars Reconnaissance Orbiter (MRO), Thermal Emission Imaging System - Visual (THEMIS-VIS) (Christensen et al., 2004) onboard Mars Odyssey and Mars Orbiter Camera - Narrow Angle (MOC-NA) (Malin et al., 2010) onboard Mars Global Surveyor (MGS). Apart from the coregistration and the orthorectification of these products, two mosaics were produced using as an input the CTX coregistered images and a radiometric correction technique (Michael et al., 2016).

The main objectives of this work are the following: firstly to demonstrate that using the novel automated coregistration and orthorectification algorithm the batch-mode fully-automatic, geometric alignment of large volumes of multi-instrument planetary data is within reach today; secondly, to carry out the processing all the way from data ingestion to product release in order to contribute georeferenced products that do not require any further processing to the planetary science community. This paper is structured so as to address both objectives. More specifically, Section 2 summarises the MC11 quadrangle, which is the region of interest in this work, before the input data and the baseline are briefly presented in Section 3. The methods used for this work are discussed in Section 4, the released products are described in Section 5 and validated in Section 6, while Section 7 concludes this work.

2. MC11 quadrangle

The quadrangle Mars Chart 11 "Oxia Palus" (MC-11) is one of the 16 equatorial quadrangles and is located to the west of the Martian prime meridian. As with the other equatorial quadrangles it covers 45° in the east-west and 30° in the north-south direction. Its extent in the east-west direction is about 2,660 km at the equator, about 2,295 km at 30° N and about 1780 km in the north-south direction. Located at the dichotomy boundary it covers both Chryse Planitia in the north-east and Arabia Terra in the west. (see Fig. 1). The quadrangle hosts the Mars Pathfinder landing site in Ares Vallis and two candidate landing sites for the ExoMars and the Mars 2020 rover missions (Oxia Planum and Mawrth Vallis). The Mawrth Vallis region is an ancient outflow channel containing extensive occurrences of phyllosilicates on the ancient plateaus (Loizeau et al., 2012). The quadrangle hosts early findings of surface changes of aeolian features. With its dark material deposits on bright floors of many craters in western Arabia Terra it offers ideal opportunities for examining bright/dark spectral differences at local scale (Ruff and Christensen, 2002). On the other hand, Chryse Planitia being one of the darkest areas on Mars with its high dust contents stands in strong albedo contrast to the bright highlands of Arabia Terra on a regional scale (Dobrea et al., 2010). Both of these aeolian dust deposit types are dynamic – (Geissler, 2005) describes changes of the albedo dichotomy between Viking and MGS observations, while changes of the dark streaks have already been observed between Mariner 6 (1969), Mariner 9 (1972) and Viking 2 (1976–1980) (Thomas and Veverka, 1979).

The eastern part of MC-11 was chosen as the first prototype half-quadrangle for the HRSC mosaicing efforts due to its complete

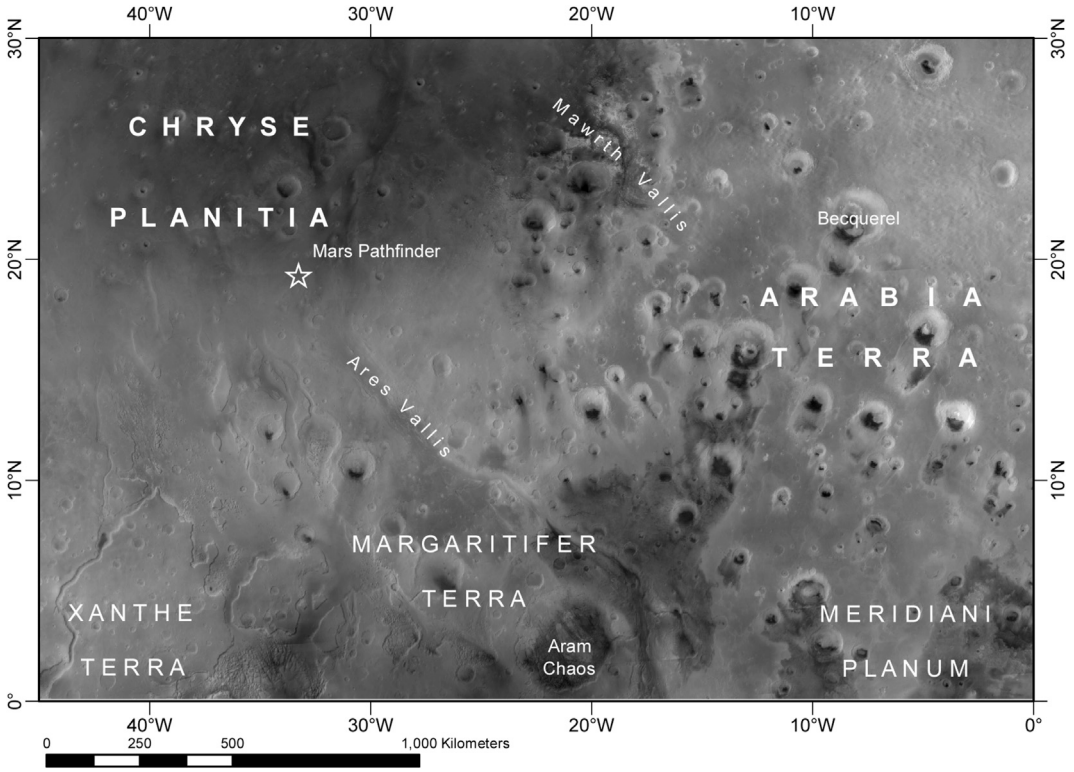

Fig. 1. Overview map of the MC-11 quadrangle with annotations. Background image: MOC WAC mosaic (Caplinger and Malin, 2001). Scale-bar only valid at the equator.

coverage of HRSC images and because of the generally good image quality of the single strips (Gwinner et al., 2016). The HRSC mosaics (both DTM and images) are available on-line for viewing in a dynamic map environment at http://imars.planet.fu-berlin.de and http://i-mars.eu/webGIS (Walter et al., 2017).

3. Input and baseline data

The input data originate from 3 instruments that have acquired high-resolution images over the last 2 decades, i.e. MOC-NA, THEMIS-VIS and CTX. MOC-NA acquired images from selected regions of interest with a spatial resolution varying from 1.5 to 12 m/pixel. By the end of the MGS mission, the MOC-NA images covered 5.45% of Mars' surface at resolutions finer than 12 m/pixel, and approximately 0.5% at better than 3 m/pixel (Malin et al., 2010). However, in the first mission phase MOC-NA was out of focus, thus producing images where the majority have diminished scientific and mapping value. If these are excluded, the 12 m/pixel coverage falls to 5.27%. The MGS orbit was nearly circular with an average altitude of about 378 km, the spacecraft orbiting the planet 12 to 13 times per day. The MOC-NA camera was a pushbroom type of camera. The camera's IFOV was 3.7μ rad, and when coupled with the 2048 element line array, narrow angle images could cover about 3 km in the cross-track dimension (Malin et al., 2010).

THEMIS is the longest surviving imaging system in orbit around Mars, acquiring images in the visual (THEMIS-VIS) and infrared (THEMIS-IR) spectrum since 2001. According to the baseline observing plan, THEMIS-VIS and THEMIS-IR compete for the same data volume, the nominal tradeoff being 62% for THEMIS-IR and 38% for THEMIS-VIS products. However, due to the mission longevity, THEMIS-VIS surpassed the mapping target of 60% of the planet in 2013. THEMIS-VIS has 1024 cross-track pixels with an 18 m IFOV covering a 18.4 km swath. The visible images are acquired in framelets that are 1024 samples crosstrack by 192 lines downtrack in size. Framelets are typically taken at 1-sec intervals, resulting in 26 rows of downtrack overlap at a nominal orbit velocity of 3.0 km/s. SNR is larger than 100 at 0.25 albedo and 60° solar incidence angle (Christensen et al., 2004).

CTX is a pushbroom camera designed to acquire 30-km-wide, 40-km-long, 5 − 6.5 metres/pixel views of the Martian surface from the nominal 255 − 320 km altitude at 15 : 00 local time (Malin et al., 2007). The CTX nominal goal was to cover 9% of Mars at 6 m/pixel (Malin et al., 2007). However, due to MROs longevity, the camera has mapped more than 90% of the Martian surface, with 34.6% covered more than once for stereo, change monitoring, or clearer atmosphere repeat imaging (Bell et al., 2013) (the latter number being from 2013). In PDS only CTX Experimental Data Record (EDR) products are released. These need to be de-noised and radiometrically calibrated through ISIS (https://isis.astrogeology.usgs.gov/) before being employed in any image processing pipeline.

The methods presented in this work have been used in order to automatically coregister, orthorectify and subsequently mosaic high-resolution images that map areas within the MC11 quadrangle. The main reason for selecting MC11 as the region of interest is the baseline availability, since the HRSC team has internally released to members of the HRSC CoI team, two high-resolution mosaics, mapping the East (MC11-E) and the West (MC11-W) part of this quadrangle, respectively (Gwinner et al., 2015), (Michael et al., 2016). The baseline was produced by the bundle block adjustment and the radiometric calibration of tens of HRSC single strips (level-4 products (Gwinner et al., 2015)). Its resolution is 12.5 m/pixel for the panchromatic image and 50 m/pixel for the DTM, while the employed datum is a sphere with radius 3396.0 km. The product is in equidistant projection, and not in sinusoidal projection as are the typical single-strip HRSC products (Gwinner et al., 2015).

Even after radiometric calibration, the HRSC image strips show inconsistent brightness levels due to the differing illumination and atmospheric conditions at the times of image capture. The images are brought into mutual consistency by means of an external brightness reference image: the 7.5 km/pixel global albedo map produced by the TES instrument (Christensen et al., 2001) is used for this purpose. The local brightness average within an HRSC strip is tied continuously to the values of the reference image. The technique eliminates brightness jumps at image boundaries in the mosaic whilst maintaining high resolution detail. Feathering is also applied along the image edges, yielding a mosaic which is close to seamless (Michael et al., 2016).

4. Methods

In this section, the novel techniques that have been developed so as to achieve a fully automatic processing of large volumes of high-resolution Mars orbital data are presented. These can be divided into two categories, aiming at the geometrical alignment of the input imagery (Section 4.1) and its radiometric calibration. (Section 4.2), respectively.

4.1. Automatic coregistration and orthorectification (ACRO)

The coregistration and orthorectification of high-resolution planetary images is typically conducted using approaches that are not fully automatic, since they require human intervention at one or more stages of the processing. For example, most commercial off-the-shelf software (COTS) (such as ArcMap (http://desktop.arcgis.com/en/arcmap/) and ENVI (http://www.harrisgeospatial.com/docs/RegistrationImageToImage.html)) require the user to select a set of tie-points in a pre-processing step, before the geometric alignment proceeds automatically, while in (Cheng et al., 2013) a manual crater detection step is employed to impose geometrical constraints for the automatic coregistration pipeline. Moreover, the few automatic techniques reported in the literature (Kim et al., 2001) (Troglio et al., 2012), depend on distributions of surface features (such as craters and rocks), which are not trivial to be theoretically modelled, while they typically exhibit severe parameter sensitivity (Sidiropoulos and Muller, 2018).

Both computer-assisted and automatic algorithms that exhibit severe parameter sensitivity are not suitable for large-scale image processing, since they require extensive human resources, the former for guiding the processing, one image at a time, while the latter for re-tuning the algorithm parameters every time they fail to produce high-quality output. The main obstacle for establishing a fully automatic algorithm, which could process a large volume of multi-instrument images in one pass is the fact that the employed image matching techniques (Lowe, 2004), (Mount et al., 1999) (Bay et al., 2006), are the ones developed for a “generic” image type, i.e. they ignore the specific properties of high-resolution planetary images (such as the large image size, the lack of distinctive features, the imaging setup, etc.). Contrary to the use of generic image matching techniques, which exhibit limitations on their transferability to planetary image applications, we have developed a novel image matching technique that is tailored to high-resolution planetary images and demonstrated its potential through the coregistration of large volumes of multi-instrument planetary images to a common baseline (Sidiropoulos and Muller, 2018).

The pipeline input is a raw high-resolution image in NASA PDS3 format, as well as a nadir image and a DTM baseline, which in the case of Mars are provided by HRSC map-projected Level-4 products (Gwinner et al., 2015). During the initial stages of the pipeline the input data are ingested into the pipeline and pre-processed using ISIS, i.e. de-noised, map-projected, transformed to tiff format, and analysed to retrieve the georeferencing information, along with relevant projection properties (such as the central latitude and longitude, the datum and the employed projection). At this stage, feature points are extracted using the Scale-Invariant Feature Transform (SIFT) algorithm (Lowe, 2004) for both the input image and the HRSC nadir image. However, instead of proceeding with SIFT matching, as is typically done in the literature, a more elaborate matching approach is followed.

Theoretically speaking, the main novelty of the matching technique is the use of geometrical constraints in the matching stage (and not merely

in a post-processing, outlier detection stage), in order to limit the number of candidate matches, thus speeding up the process while decreasing both the false acceptance (errors of commission) and the false rejection (errors of omission) ratios. The image matching algorithm, named coupled decomposition (Sidiropoulos and Muller, 2015b), alternates between a step in which a small number of common features are estimated and a step in which the common features are used to decompose the image into corresponding sub-images, which are further processed independently. In coupled decomposition matching the list of tie-points is augmented exponentially, while the full-image matching step is delayed until the small size of the corresponding sub-images secure a fast and accurate matching.

Coupled decomposition is a generic algorithm, thus not specifying either the method of identifying the common features or the decomposition process. In order to focus on high-resolution Mars images a variation called ring matching has been developed (Sidiropoulos and Muller, 2018). Ring matching uses the fact that in high-resolution remote sensing images the misregistration error type with the largest magnitude is translation, i.e. that:

$$E(i,j) = E_s(i,j) + C, \quad \|C\| \gg \|E_s\| \tag{1}$$

where $E(i,j)$ is the total misregistration error of the input image pixel p with (pixel) coordinates (i,j), C is the misregistration error corresponding to global translation (which is constant for all p) and $E_s(i,j)$ the misregistration residuals after global translation is modelled out. Assuming that prior georeference information is available for the input image (as is the case in Mars high-resolution imagery), if $E = 0$ then the input image pixel $p(i,j)$ would correspond to the baseline image pixel $q(i',j')$ with the same world coordinates. The misalignment implies that the correct position of the corresponding pixel is in $q(i',j') + C + E_s(i,j)$. Therefore, if $||C||$ is known, then the corresponding baseline pixel would be in a ring of radius $||E_s||$. If the baseline image is decomposed into rings of radius M (where $||M|| \gg ||E_s||$) and point p is matched with the points of each ring independently, then most correct matches would be found in one specific ring, which depends on $||C||$, and all other rings would have only false positive matches. The correct ring can be found by repeatedly searching for geometrically consistent matches through RANSAC (Fischler and Bolles, 1981).

The ring matching algorithm is described in Algorithm 1. Because SIFT point matching is conducted sequentially, one-point-at-a-time, and the algorithm is terminated once the correct ring is found, the required computational time is limited. The output of the algorithm is a small number of matched SIFT points, which are used to initialise coupled decomposition. The coupled decomposition sub-images are produced by the grid that ring matching tie-points imply, i.e. by using the lines between SIFT points to decompose the images into corresponding non-overlapping patches, which are subsequently matched independently.

Algorithm 1 Ring matching algorithm.

Input: (a) SIFT points of input image v_t and their georeferencing information (b) SIFT points of the baseline image v_b and their georeferencing information (c) A structure S of N empty cells which is populated by matched SIFT points.

Step 1: Select randomly a SIFT point $v \in v_t$, with pixel coordinates $p(i,j)$.

Step 2: Estimate the baseline pixel $q(i',j')$ that corresponds to the same world coordinates with $p(i,j)$.

Step 3: Decompose set v_b into N non-overlapping subsets v_{bk}, $k \in \{1,2,...N\}$. A SIFT point with coordinates $q'(i'',j'')$ would be assigned to subset v_{bk}, if and only if $(k-1)N < ||q'-q|| \leq kN$.

Step 4: Match v independently with each set of SIFT points v_{bk} and append the structure S with all matches.

Step 5: Examine independently all N cells of S using a RANSAC-based outlier detection technique [13], [36]. If a cell is found with geometrically consistent matches then return the coordinates of the matched points (after discarding the outliers), otherwise go to Step 1.

Matching using coupled decomposition produces as an output, a set of point correspondences in pixel coordinates. Subsequently, using the georeferencing information of the HRSC nadir image, which is the coregistration baseline, along with the corresponding DTM (in order to estimate the height value), these can be transformed into a set of correspondences W between pixels $p(i,j)$ in the input image and world coordinates (X,Y,Z). Using w as an input, a camera model is estimated, i.e. two functions $f_x(X,Y,Z)$ and $f_y(X,Y,Z)$ that determine the projection of Mars' surface onto the input image plane. In this work, we have used a hybrid camera model; a rigid linear pushbroom camera model (Gupta and Hartley, 1997) produces the initial coregistration, while the residuals are suppressed using a polynomial model.

The coregistered and orthorectified images are produced by estimating the image coordinates of world points that belong to a rectangular grid with cell size equal to the input image resolution. For each point with world coordinates (X,Y) a DTM is used to generate the vector (X,Y,Z), which is then given as input to f_x and f_y to estimate the input image corresponding pixel. When this is within the image limits, the corresponding grey value is copied to the coregistered image, interpolating if necessary. The output image is saved in a geotiff format, along with metadata related to the raw image properties, the image position and the coregistration processing. A more complete list of properties is saved into a separate text file, while the image footprint is also produced. More details about the automatic coregistration and orthorectification (ACRO) algorithm can be found in (Sidiropoulos and Muller, 2018).

4.2. Radiometric correction and mosaicing

The brightness equalisation is carried out based on the technique developed for radiometrically correcting HRSC images (Michael et al., 2016) and used for producing the HRSC baseline mosaics. In the case of CTX mosaics, firstly, an intermediate resolution brightness reference map is created. More specifically, individual CTX images are divided into cells and the local mean brightness of a cell is tied to the local mean of the equivalent area in the reference image, using a continuous interpolation. Images adjusted this way are placed into a mosaic at moderate resolution with the goal of obtaining a higher resolution brightness reference with the same average brightness characteristics as the source reference image. Remaining image edge artifacts are eliminated by applying a Gaussian blur, which leaves the average spatial brightness characteristic unchanged.

During the next step, a full resolution mosaic is created and tied to an intermediate brightness reference map. In this second iteration, the images are processed at full resolution, tying locally and continuously to the intermediate brightness reference map, using a smaller cell size to find the local mean. The images are placed into the final mosaic, feathering (i.e. fading from 0% to 100% transparency) over a narrow range of pixels. Note that feathering helps to hide the image boundary, but is only effective if the image brightnesses are well matched already.

Then the image sequence is optimised. Images are put into the mosaic in order of best ground sampling resolution (lowest to highest). In cases where this is not the optimal sequence, for reasons of quality or consistency, we can manually construct a list of ordering relations. This is done by comparison of the assembled mosaic with individual image strips. Finally, the contrast is recovered. Many images show reduced contrast caused by increased atmospheric scattering, appearing in the processed mosaic as relatively flat bands. Contrast is recovered by stretching the histogram width for individual images by a factor of typically $1.5-6$, the amount being adjusted iteratively after inspection of the assembled mosaic.

In this way, we are able to produce a visually consistent image mosaic for geomorphological studies in the absence of a full correction of atmospheric effects. More information about this brightness equalisation technique can be found in (Michael et al., 2016).

5. Released product description

The baseline resolution imposes the most significant constraint regarding the input imagery that can be automatically coregistered. It has been experimentally found that ring matching exhibits robustness to resolution differences up to 1 order of magnitude, thus the HiRISE dataset (McEwen et al., 2007) is not included in the processing. Moreover, the use of a pushbroom camera model excludes Viking Orbiters data, which uses an obsolete frame camera technology. However, experiments have provided evidence that the ring matching algorithm is completed successfully in the case of Viking Orbiters images. More research is required for the use of the extracted tie-points in order to establish high-quality camera models, because the imaging setup causes large-magnitude local deformations (e.g. at one corner of the image the misregistration error may be less than 1 pixel while at the opposite larger than 50 pixels). As a result, Viking Orbiters images are not included in this work.

It was established that other setup characteristics (such as the point spread function, the spectral responsivity, the bit depth, etc.) of Mars orbital cameras do not deter the algorithm performance, thus the remaining instruments can be coregistered to the HRSC MC11 mosaics. This consists of images acquired by 3 instruments, namely, CTX, with nominal resolution $6m/pixel$, THEMIS-VIS, with resolution $17.5-75m/pixel$ and MOC-NA, with resolution $1.5-12m/pixel$. The total number of images per half-quadrangle is reported in Table 1.

The automatically coregistered images follow the georeference setup of the HRSC mosaics, i.e. they are projected using equidistant projection to a spherical datum of radius 3396.0 km (during the image ingestion the same datum is used, so as to avoid distortions). Other than that, the coregistration process does not change any other feature of the input image, such as the bit depth and the spectral response. The output format is a single-band geotiff image, in which 0 is used as a no-data value. A small part of the THEMIS-VIS dataset consists of multi-band images. In this case the coregistration refers to the red channel, while all other channels are ignored for consistency reasons.

Table 1
The available (top row) and the processed (bottom row) high-resolution orbital images (*Res.* $\leq$ 100 m/pixel) of Mars MC11 quadrangle, excluding HiRISE and Viking Orbiters images. This list does not include any images acquired after 28 February 2017.

Dataset	MC11-E			MC11-W		
	CTX	MOC-NA	THEMIS	CTX	MOC-NA	THEMIS
# Images	1365	1558	3629	913	1220	3152
# Proc. Im.	1303	1043	3323	868	695	2860

In a post-processing stage the image visual quality is examined, using an in-house automatic planetary image quality assessment algorithm (Sidiropoulos and Muller, 2015c). The output is a binary value, with 0 corresponding to poor visual quality and 1 to good visual quality image. The current version of the software does not discriminate between poor quality caused by errors in the processing chain and inherent poor quality. Since the latter may be associated with interesting atmospheric phenomena, it should be noted that the assessment refers only to visual quality (i.e. quality of the raw visual signal) and not to information quality (quality of the content of the signal). However, this flag can be useful in applications related to planetary image processing, e.g. when the coregistered images are to be the input of a stereo processing chain. Hence, this is stored in the metadata file, which is a text file having the same name as the coregistered image (apart of course from the suffix). Other information stored in the metadata file includes the camera model, processing statistics (e.g. the start and end time), and the original label of the raw image, as extracted using ISIS software.

For each automatically coregistered image a footprint is created and stored in a shapefile format. The shapefile uses the same datum and projection as the input image and includes mostly geometric information, i.e. the corrected lat/lon position of the high-resolution image. The footprint perimeter and bounding box are described following the format used by ODE (http://ode.rsl.wustl.edu/mars/coverage/ODE_Mars_shapefile.html) i.e. cells X and Y give the lat/lon coordinates of the endpoints of consecutive perimeter linear segments, with the only difference being the use of 3, instead of 2, decimal digits.

The ACRO processing that generated these products is fully automatic, i.e. it didn't include any human intervention or any parameter tuning. The input is given to the pipeline with a single set of parameters and the algorithm is executed in a batch-mode, while the images for which the algorithm failed to produce any results (the most common issue being the failure of ring matching to retrieve the correct ring) are copied to a separate repository to be reprocessed. The latter is conducted in a second pass, in which the automatically coregistered CTX images was the baseline (using the same HRSC DTM, though). This post-processing significantly reduced the failure rate (e.g. from 9.3% to 4.54% in CTX MC11-E images, from 8.77% to 4.93% in CTX MC11-W images, from 54.6% to 33.1% in MOC-NA MC11-E images, etc.), thus demonstrating that failure rate depends on the quality of both the input image (e.g. the failure rate of MOC-NA dataset is much larger than that of CTX) and the baseline. The low failure rate of the algorithm when using either CTX as an input or as a baseline provides evidence for the potential of the novel (ring matching) algorithm. If high-quality planetary stereo products are consistently achieved then ring matching can successfully handle the automatic coregistration of high-resolution planetary images, without requiring the extensive use of human resources. The MC11 coregistered products, as well as the footprints and metadata are available upon request. Two examples of 7 coregistered CTX and 7 coregistered MOC-NA

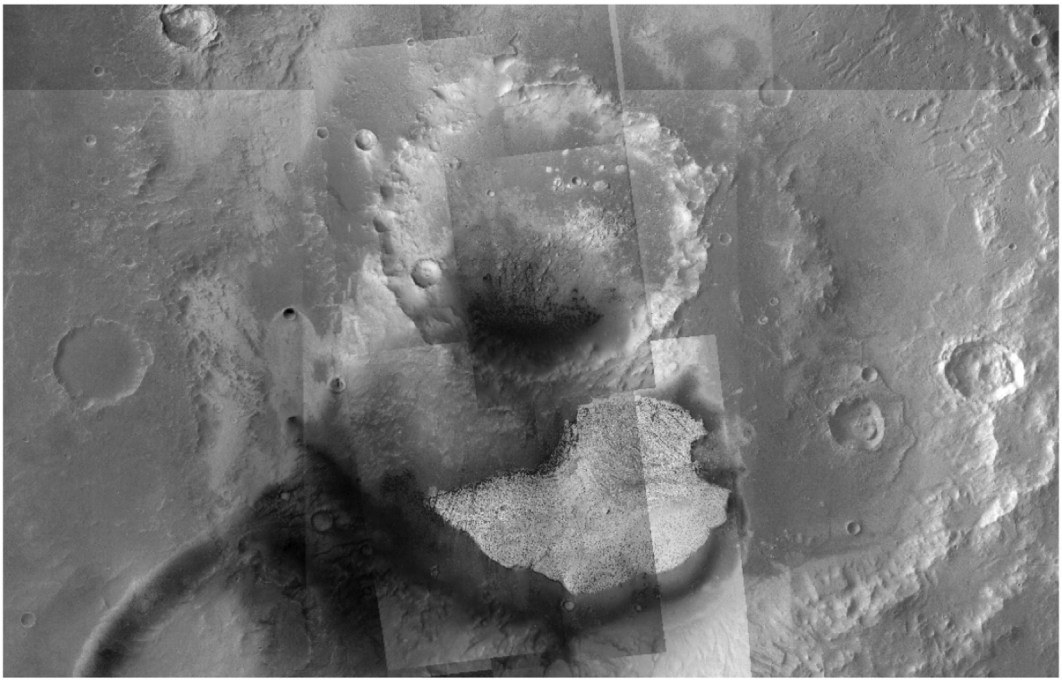

Fig. 2. An example of 7 CTX images coregistered to the HRSC MC11-E mosaic (which is the background image). The centre coordinates are 22.1 N, 352 E (Becquerel crater).

Fig. 3. An example of 7 MOC-NA images coregistered to the HRSC MC11-E mosaic (which is the background image). The centre coordinates are 1.5 N, 359 E.

images are shown in Figs. 2 and 3, respectively.

Finally, a mosaic is assembled from the coregistered CTX images (Fig. 4). As seen from Fig. 2, the CTX images are not mutually consistent in brightness, so that the image boundaries remain visible. We therefore apply the same equalisation technique to the CTX as was done for the construction of the HRSC mosaic, referencing the local average brightness to an external calibration image. Once again, we use the TES global albedo map at 7.5 km/pixel. This allows us to produce a near-seamless mosaic.

Because of the lower imaging altitude of the CTX instrument compared to HRSC, we also see some across-strip brightness variation, likely resulting from the varying optical path length through the atmosphere in combination with its changing conditions. This effect often appears in CTX images as an apparent curvature of the image strip. Use of the external brightness reference alleviates this problem to a degree, but does not fully eliminate it. A dedicated correction before mosaicing would be preferable.

An additional issue is that the CTX strip width (≈30 km) encompasses fewer pixels of the brightness reference image (7.5 km/pix) than does the typical HRSC strip width (50–75 km). This constrains the calculation of

Fig. 4. An overview of the CTX MC11 mosaic, which was produced by ACRO'd CTX images, as displayed in the iMars website (where the mosaic is available in full resolution). The background image is the HRSC MC11 image mosaic (Gwinner et al., 2016).

local average brightness to a smaller region, and makes it harder for the technique to retain brightness information at intermediate scales. In future work, we will consider the use of a higher resolution brightness reference.

Regarding the release status, currently all processed images, as well as the CTX mosaics are available for viewing in a dynamic map environment at the iMars website (http://i-mars.eu/webGIS). Anyone interested to access the images can contact the authors.

6. Accuracy evaluation

The MC11 coregistered and orthorectified products that were generated using the techniques of Section 4.1 are examined so as to estimate the achieved geometric accuracy, as well as to confirm the potential of the introduced technique. The main issue with the geometric evaluation is that on Mars there is no ground control network that could be used to assess the position of landmark features in the coregistered images. The current global geodetic reference frame comes from Mars Orbiter Laser Altimeter (MOLA) (Zuber et al., 1992), which is used as the baseline in most mosaics produced by instrument teams (including the HRSC MC11 mosaics (Gwinner et al., 2016)). Hence, the analysis is limited to a consistency check, between corresponding pixels in the coregistered and the (consistent with the MOLA sphere) HRSC mosaic.

More specifically, the set of tie-points estimated in the first stage of the algorithm (i.e. before the camera model estimation) are split into two halves, the first one used to determine the camera model while the second used for validation. Following the introduced method, the mean geometric accuracy both along and across track is computed for every coregistered image. The results, given in Table 2, show that the achieved accuracy is consistently less than 1 baseline pixel, regardless of the input dataset, thus confirming that in general sub-pixel accuracy in both directions is achieved. The comparison of the 25%, 50% and 75% percentile imply that for all datasets the misregistration errors are rather evenly distributed around the median error. The lack of a strong tail in the error distribution signifies that only a small number of processed images exhibit large misregistration error. The bottom line of Table 2, which reports the coregistered image percentage with average errors less than 1 HRSC pixel in both directions, provides further evidence for the geometric accuracy of the technique.

Table 2
Statistics of the achieved misregistration error. The top 6 lines correspond to the median (50% error), 25% quantile and 75% quantile of the misregistration error (counted in metres). The bottom line reports the percentage of images that the mean misregistration error was less than 1 baseline pixel for both X and Y directions.

Dataset	MC11-E			MC11-W		
	CTX	MOC-NA	THEMIS	CTX	MOC-NA	THEMIS
X 50% Error (m)	5.33	5.19	7.83	5.55	4.82	7.89
Y 50% Error (m)	5.29	4.66	7.22	5.83	4.68	7.2
X 25% Error (m)	3.84	4.27	6.14	4.3	3.91	6.6
Y 25% Error (m)	4.22	3.96	5.87	4.79	3.64	6.17
X 75% Error (m)	7.76	6.49	9.88	7.61	6.18	9.9
Y 75% Error (m)	6.76	5.63	9.27	7.52	6.57	9.13
Err. $\leq$ 1 HRSC pix. (%)	87.39	96.7	87.54	88.93	91.51	86.15

While in general the automatic coregistration and orthorectification pipeline exhibits good performance across all 3 datasets, the inter-dataset comparison of the achieved geometric accuracy and failure rate leads to interesting conclusions regarding the pipeline potential, as well as its limitations. More specifically, there is a contrast between CTX, which has a low failure rate but high rate of images with misregistration error higher than 1 pixel, and MOC-NA, which has a high failure rate but a low rate of images with misregistration error higher than 1 pixel. This could be explained by the data acquisition problems that MOC-NA faced during the MGS mission (Malin et al., 2010), a large part of which originated from adverse atmospheric conditions (for example, between July and October 2001 Mars was obscured by a planet-encircling dust cloud (Malin et al., 2010)). This contributed to data quality issues, which affect ACRO performance. MOC-NA imagery is either of sufficient quality to achieve sub-pixel coregistration or of (low) quality that prevents establishing point matches between the input image and the HRSC baseline, thus causing the pipeline to fail (for example, the failure rate between July and October 2001 is 76.5%). On the other hand, CTX (and to a lesser degree, THEMIS-VIS) includes a number of images for which sub-pixel accuracy was not achieved, but approximate coregistration (with an

Table 3
Automatic coregistration failure rate and geometric accuracy versus the 1-to-7 image rating assigned to THEMIS-VIS imagery from the instrument team. The reported accuracies are the median scores over the corresponding THEMIS-VIS subsets.

THEMIS Image Rating	1	2	3	4	5	6	7
Failure Rate MC11-E (%)	77.8	40	29.7	22	16.8	10.9	12.9
Failure Rate MC11-W (%)	100	66.7	34.9	20.8	17.7	0	12.5
X Accuracy MC11-E (m)	8.4	10.9	8.7	8.1	7.7	5.8	5.8
Y Accuracy MC11-E (m)	8.2	10.3	7.7	7.5	7.2	5.6	5.8
X Accuracy MC11-W (m)	N/A	11.3	8.8	8.3	7.7	6.7	6.2
Y Accuracy MC11-W (m)	N/A	9.9	7.7	7.7	7	6.7	6.1

accuracy between 1 and 5 pixels) was possible.

It would be useful to have a quality index for each input image, in order to fully examine the correlation of failure rate/geometric accuracy and image quality. Unfortunately, such an image rating exists only for THEMIS-VIS products, for which a 1-to-7 scale (1 denoting the worse quality) is used by the instrument team. The achieved scores are summarised in Table 3. Note that approximately 99% of the MC11 THEMIS-VIS images are assigned an image rating between 3 and 6, therefore the statistics corresponding to classes 1, 2 and 7 are estimated over a small number of samples. Even when these are ignored, the correlation between image quality and automatic coregistration performance is apparent. The substantially lower failure rate of THEMIS-VIS images of good quality (including a 0% failure for images of quality 6, for MC11-W), along with the progressive reduction of the misregistration residuals, signify that the presented technique can be systematically used to automatically coregister large volumes of planetary high-resolution images, with the proviso that the image acquisition should not cause a severe image quality degradation.

Another important factor that is associated with the coregistration performance is the off-nadir angle of the input image. Images that are acquired with a large off-nadir angle are expected to follow the terrain morphology, a property that is useful for stereo applications but is a source of noise in coregistration and orthorectification pipelines. In this case too, we restrict the analysis to a single dataset, namely, CTX, because THEMIS-VIS acquires by default nadir images, while MOC-NA has to deal with serious quality issues when attempting to acquire high-resolution stereo images (Malin et al., 2010). More specifically, the CTX images of both MC11-E and MC11-W were clustered according to their off-nadir angle into bins of 2.5° and the median along-track and across-track accuracies were estimated for each bin (Fig. 5). As expected, the former remains rather constant for all off-nadir angles. On the other hand, the latter starts to increase progressively when the off-nadir angle exceeds approximately 10°, reaching 10.4 m for off-nadir angles higher than 25°. This is explained by the failure of the polynomial model to suppress the rigid linear pushbroom camera model residuals. The dashed lines, corresponding to the median accuracy that is achieved solely by the rigid linear pushbroom camera model, i.e. without employing a polynomial model to suppress the residuals, imply that the rigid camera model accuracy is rather independent of the off-nadir angle. The Y accuracy deterioration in high off-nadir angles is caused by the fact that the Y residuals, being increasingly influenced by the terrain morphology, become progressively more difficult to be modelled by low-order polynomial functions.

On the other hand, we found only a weak correlation between the pipeline performance and the image size. The size/performance correlation hypothesis in this case would be that large-sized images would tend to accumulate non-linear deformations caused by inconsistencies in the spacecraft trajectory, thus increasing the coregistration residual magnitude. The size/performance analysis was conducted only on the CTX dataset because CTX is the only (out of the 3 examined) instrument that has so large a size variation, with the ratio of the largest over the smallest image (the size counted in "total number of pixels") being more than a factor of 26, while the top-quality consistency of the CTX imagery facilitates modelling-out other image properties correlated with the coregistration accuracy. Limiting the stereo angle between 0° to 10° was the only pre-processing that was necessary, before clustering the images to 10 (equal-sized) bins according to the input image total pixel number and estimating the median X and Y accuracy (Fig. 6). While both lines indicate a slight deterioration with image size, the X and the Y error increasing by approximately 2 and 1 m, respectively, the performance correlation with size seems weak, especially if we take into account the fact that some of the bins corresponding to large sizes had a small number of samples (e.g. only 1.1% of processed CTX images was between 210 – 235 Mpixels).

We did not conduct a quantitative analysis of the performance dependence on the resolution. Qualitatively speaking, the fact that imagery from 3 high-resolution instruments with distinct nominal resolutions were systematically coregistered to a common HRSC baseline gives evidence about the algorithm's robustness to input image resolution. However, even though sub-pixel accuracy was achieved for MC11 images with resolution varying from 1.5 m/pixel to 70 m/pixel, the dependence of the residuals on the resolution is not easily quantitatively analysed. The main reason is that this would require instruments capable of acquiring images of similar SNR in a large range of resolutions. However, none of the instruments included in this analysis (nor any other high-resolution Mars orbital camera) has this capability. Instead, each of the

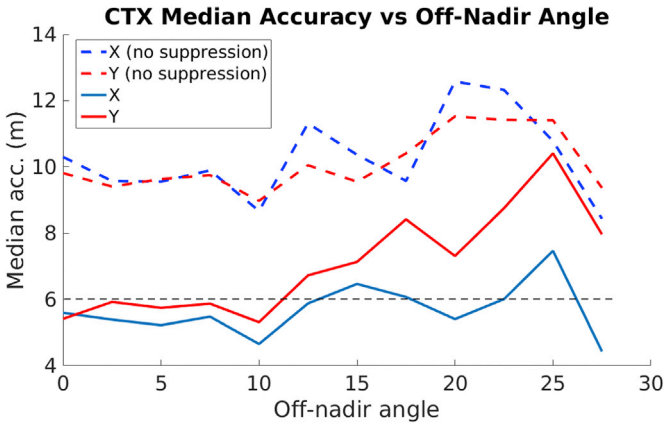

Fig. 5. A plot of the CTX misregistration error versus the off-nadir angle. The individual results were clustered to 12 bins of 2.5° each and the median error for each bin was stored. The dotted black line show the nominal CTX resolution (6 m/pixel).

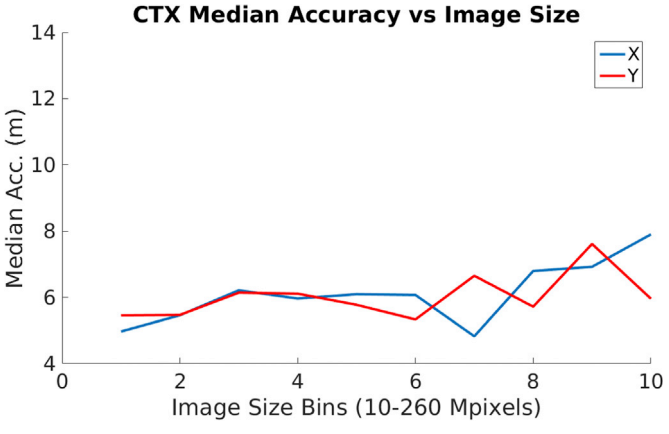

Fig. 6. A plot of the CTX misregistration error versus the image size, counted in number of Mpixels. The minimum size of the CTX images coregistered in this work is 10 Mpixels and the maximum 260 Mpixels. The imagery was clustered to 10 bins of 25 Mpixels size each and the median score for each bin was stored.

3 instruments use a fixed nominal resolution. Images with resolution different than the nominal one signify the presence of severe noise in the original acquisition, which imposed a subsampling (e.g. averaging) to suppress the noise. As a result, more than 99.9% of MC11 CTX images and 92% of MC11 THEMIS-VIS images have resolution approximately equal to the nominal one (6 m/pixel for CTX and 17.5 m/pixel for THEMIS-VIS). On the other hand, while MOC-NA images have been released with various resolutions (multiples of 1.5 m/pixel), according to the input noise level, the input resolution was not found to be correlated with neither the accuracy or the failure rate.

7. Conclusions and future work

The planetary science community lacks techniques that would successfully process large volumes of data without requiring the use of human resources, as well as training in sophisticated computer-vision and data-science software. This work constitutes a significant step towards this direction, by providing strong evidence that the batch-mode fully-automatic coregistration and orthorectification of large volumes of Mars high-resolution images is feasible. Far from being a mere demonstration of the theoretical properties of the pipeline, this work carries through the processing all the way from the input up to the release of large volumes of high-quality coregistered data, being acquired by 3 NASA Mars orbital cameras, presenting and discussing in detail the main characteristics of the actual released products. The released products cover the MC11 quadrangle, which is of particular scientific interest, while including a number of rover landing sites. The further processing of CTX coregistered images generated two radiometrically calibrated mosaics. Even though the produced CTX mosaics are currently the most extensive mosaics available to the community, their quality can be improved by further adjusting the radiometric correction pipeline to the specific characteristics of CTX imagery. In total, the overall pipeline assessment confirms the validity of this analysis, thus making it suitable for the processing of more Mars quadrangles, and possibly (when a high-resolution baseline becomes available) the pre-existing Mars imagery. Our future work will follow this direction.

Acknowledgements

This work has received funding from the STFC "MSSL Consolidated Grant" ST/K000977/1 and partial support from the European Union's Seventh Framework Programme (FP7/2007–2013) under i-Mars grant agreement no 607379. Greg Michael and Sebastian Walter were supported by the German Space Agency (DLR Bonn), grant 50QM1702 (HRSC on Mars Express), on behalf of the German Federal Ministry of Economics and Technology. The core of the software for ACRO is available through the iMars GitHub account https://github.com/UCL-iMars/ACRO/blob/master/UCL_ring_matching.m.

References

Acton, C.H., 1996. Ancillary data services of NASAs navigation and ancillary information facility. Planet. Space Sci. 44 (1), 65–70.

Bay, H., Tuytelaars, T., Gool, L.-V., 2006. SURF: speeded up robust features. In: Proc. European Conf. On Computer Vision, pp. 404–417.

Bell III, J.F., Malin, M.C., Caplinger, M.A., Fahle, J., Wolff, M.J., Cantor, B.A., James, P.B., et al., 2013. Calibration and performance of the Mars Reconnaisance Orbiter Context Camera (CTX). Mars 8, 1–14.

Byrne, S., Dundas, C., Kennedy, M., Mellon, M., McEwen, A., Cull, S., Daubar, I., et al., 2009. Distribution of mid-latitude ground ice on Mars from new impact craters. Science 325 (5948), 1674–1676.

Caplinger, M.A., Malin, M.C., 2001. Mars Orbiter Camera geodesy campaign. J. Geophys. Res. Planets 106 (E10), 23595–23606.

Cheng, L., Ma, L., Yang, K., Liu, Y., Li, M., 2013. Registration of Mars remote sensing images under the crater constraint. Planet. Space Sci. 85, 13–23.

Christensen, P.R., Jakosky, B.M., Kieffer, H.H., Malin, M.C., McSween, H.Y.J., Nealson, K., Mehall, G.L., et al., 2001. Mars Global Surveyor Thermal Emission Spectrometer experiment: investigation description and surface science results. J. Geophys. Res. 106 (E10), 23823–23871.

Christensen, P.R., Jakosky, B.M., Kieffer, H.H., Malin, M.C., McSween, H.Y.J., Nealson, K., Mehall, G.L., et al., 2004. The thermal emission imaging system (THEMIS) for the mars 2001 Odyssey mission. Space Sci. Rev. 110 (1), 85–130.

Cooper, M.A.R., Cross, P.A., 1988. Statistical concepts and their application in photogrammetry and surveying. Photogramm. Rec. 12 (71), 637–663.

K. Di, Y. Liu, W. Hu, Z. Yue, Z. Liu, Mars surface change detection from multi-temporal orbital images, IOP Conf. Ser. Earth Environ. Sci. 17.

Dobrea, E.Z.N., Bishop, J.L., McKeown, N.K., Fu, R., Rossi, C.M., Michalski, J.R., Heinlein, C., et al., 2010. Mineralogy and stratigraphy of phyllosilicate-bearing and dark mantling units in the greater Mawrth Vallis/West Arabia Terra area: constraints on geological origin. J. Geophys. Res. Planets 115 (E7), E00D19.

Fischler, M.A., Bolles, R.C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM 24 (6), 381–395.

Geissler, P.E., 2005. Three decades of Martian surface changes. J. Geophys. Res. Planets 110 (E2), E02001.

Gupta, R., Hartley, R., 1997. Linear pushbroom cameras. IEEE Trans. Pattern Analysis Mach. Intell. 19 (9), 963–975.

Gwinner, K., Jaumann, R., Bostelmann, J., Dumke, A., Elgner, S., Heipke, C., Kersten, E., et al., 2015. The first quadrangle of the Mars Express HRSC multi orbit data products (MC11-E). In: European Planetary Science Congress (EPSC).

Gwinner, K., Jaumann, R., Hauber, E., Hoffmann, H., Heipke, C., Oberst, J., Neukum, G., et al., 2016. The High Resolution Stereo Camera (HRSC) of Mars Express and its approach to science analysis and mapping for mars and its satellites. Planet. Space Sci. 126, 93–138.

Kim, J.R., Muller, J.-P., Morley, J.G., Mitchell, K.L., 2001. Automated registration of MDIM with MOLA tracks. In: 32nd Annual Lunar and Planetary Science Conference (LPSC).

Loizeau, D., Werner, S.C., Mangold, N., Bibring, J.-P., Vago, J.L., 2012. Chronology of deposition and alteration in the Mawrth Vallis region, mars. Planet. Space Sci. 72 (1), 31–43.

Lowe, D., 2004. Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vis. 60 (2), 91–110.

Malin, M.C., Bell, J.F., Cantor, B.A., Caplinger, M.A., Calvin, W.M., Clancy, R.T., et al., 2007. Context Camera investigation on board the Mars Reconnaissance Orbiter. J. Geophys. Res. Planets 112 (E5), 1–60.

Malin, M.C., Edgett, K.S., Cantor, Caplinger, M.A., Danielson, E., Jensen, E.H., Ravine, M.A., Sandoval, J.L., Supulver, K.D., 2010. An overview of the 1985-2006 Mars Orbiter Camera science investigation. Mars 5, 1–60.

McEwen, A., Ojha, L., Dundas, C., Mattson, S., Byrne, S., Wray, J., Cull, S., et al., 2011. Seasonal flows on warm Martian slopes. Science 333 (6043), 740–743.

McEwen, A., Eliason, E., Bergstrom, J., Bridges, N., Hansen, C., Delamere, W., Grant, J., et al., 2007. Mars Reconnaissance Orbiter's high resolution imaging science experiment (HiRISE). J. Geophys. Res. Planets 112 (E5).

Michael, G.G., Walter, S.H.G., Kneissl, T., Zuschneid, W., Gross, C., McGuire, P.C., Dumke, A., et al., 2016. Systematic processing of Mars Express HRSC panchromatic and colour image mosaics: image equalisation using an external brightness reference. Planet. Space Sci. 121, 18–26.

Mount, D.M., Netanyahu, N.S., Moigne, J.L., 1999. Efficient algorithms for robust feature matching. Pattern Recognit. 32 (1), 17–38.

Ruff, S.W., Christensen, R.R., 2002. Bright and dark regions on Mars: particle size and mineralogical characteristics based on Thermal Emission Spectrometer data. J. Geophys. Res. Planets 107 (E12), 5119.

Schorghofer, N., Aharonson, O., Gerstell, M., Tatsumi, L., 2007. Three decades of slope streak activity on Mars. Icarus 191 (1), 132–140.

Sidiropoulos, P., Muller, J.-P., 2015. On the status of orbital high-resolution repeat imaging of Mars for the observation of dynamic surface processes. Planet. Space Sci. 117, 207–222.

Sidiropoulos, P., Muller, J.-P., 2015. Matching of large images through coupled decomposition. IEEE Trans. Image Process. 24 (7), 2124–2139.

Sidiropoulos, P., Muller, J.-P., 2015. Automatic quality assessment of planetary images. In: European Planetary Science Congress.

Sidiropoulos, P., Muller, J.-P., 2016. Large-scale co-registration of Mars high-resolution NASA images to HRSC: a case-study of the MC11-E quadrangle. In: Lunar and Planetary Science Conference.

Sidiropoulos, P., Muller, J.-P., 2017. Automatic detection of changes on Mars surface from high-resolution orbital images. In: EGU General Assembly.

Sidiropoulos, P., Muller, J.-P., 2018. A systematic solution to multi-instrument co-registration of high-resolution planetary images to an orthorectified baseline. IEEE Trans. Geoscience Remote Sens. 56 (1), 1–15.

Soffen, G.A., Snyder, C.W., 1976. The first Viking mission to Mars,. Science 193 (4255), 759–766.

Thomas, P., Veverka, J., 1979. Seasonal and secular variation of wind streaks on Mars: an analysis of Mariner 9 and Viking data. J. Geophys. Res. Solid Earth 84 (B14), 8131–8146.

Troglio, G., Moigne, J.L., Benediktsson, J.A., Moser, G., Serpico, S.B., 2012. Automatic extraction of ellipsoidal features for planetary image registration. IEEE Trans. Geoscience Remote Sens. Lett. 9 (1), 95–99.

Wagstaff, K.L., Panetta, J., Ansar, A., Greeley, R., Hoffer, M.P., Bunte, M., Schorghofer, N., 2012. Dynamic landmarking for surface feature identification and change detection. ACM Trans. Intelligent Syst. Technol. 3 (3).

Walter, S.H.G., Muller, J.-P., Sidiropoulos, P., Tao, Y., Gwinner, K., Putri, A.R.D., Kim, J.-R., Steikert, R., van Gasselt, S., Michael, G.G., 2017. The iMars web GIS - an interactive online mapping tool for the spatio-temporal visualization of topography

data and dynamic time-series of single image layers. ISPRS J. Photogrammetry Remote Sens. (Submitted).
Xin, X., Di, K., Wang, Y., Wan, W., Yue, Z., 2017. Automated detection of new impact sites on martian surface from HiRISE images. Adv. Space Res. 60 (7), 1557–1569.
Zuber, M.T., Smith, D.E., Solomon, S.C., Muhleman, D.O., Head, J.W., Garvin, J.B., Abshire, J.B., Bufton, J.L., 1992. The Mars Observer Laser Altimeter investigation. J. Geophys. Res. 97 (E5), 7781–7797.