

A Systematic Solution to Multi-Instrument Coregistration of High-Resolution Planetary Images to an Orthorectified Baseline

Panagiotis Sidiropoulos¹, *Member, IEEE* and Jan-Peter Muller²

Abstract—We address the problem of automatically coregistering planetary images to a common baseline, introducing a novel generic technique that achieves an unprecedented robustness to different image inputs, thus making batch-mode coregistration achievable without requiring the usual parameter tweaking. We introduce a novel image matching technique, which boosts matching performance even under the most strenuous circumstances, and experimentally demonstrate validation through an extensive experimental multi-instrument setup that includes images from eight high-resolution data sets of the Mars and the Moon. The technique is further tested in a batch-mode processing, in which approximately 1.6% of all high-resolution Martian imagery is coregistered to a common baseline.

Index Terms—High-resolution imaging, image matching, image registration, multi-instrument coregistration, planetary images, remote sensing.

I. INTRODUCTION

SINCE Luna 3 reached the far side of the Moon in 1959, acquiring the first orbital images of a planetary body, a large number of missions from several space agencies have retrieved millions of orbital images from the rocky planetary bodies of our solar system, from the Moon [2], [1] and Mars [3] up to the dwarf planets of Pluto and Charon [4]. This data volume has significantly increased our knowledge about planetary geomorphology [5], [6], while a number of previously unknown dynamic surface features have been identified and thoroughly examined [7]–[9], demonstrating that Earth is far from being the only active planet of the solar system.

It is commonly understood that both geomorphological analysis and change detection tasks would greatly benefit from the integration of all data into a common geospatial framework. However, even though most space missions append temporal and spatial metadata to their images (typically following the Spacecraft Planet Instrument C-matrix Events

standard [10]) that allow an approximate position estimation, small errors in the instrument position and orientation can cause spatial inconsistencies that may reach the order of kilometers for high-resolution images (note that in this paper we are using the high-resolution planetary image definition of [3], i.e., planetary images with resolution finer than 100 m/pixel). As a matter of fact, these inconsistencies are present not only on images acquired by different instruments but also in single-instrument images, thus resulting in each image being actually projected onto its own coordinate system.

Coregistration of planetary images, i.e., the projection of images onto a single coordinate system, is traditionally conducted independently for each input image, using manual or semiautomatic pipelines based on software such as ENVI¹ and DaVinci² [11]. Such approaches determine the input parameters manually and individually for each image, thus making it impossible to automate the (batch-mode) processing of a large set of images with a single set of processing parameters to be applied irrespective of the image source or the image location. Moreover, typical planetary image coregistration pipelines generally have high computational cost and require a good knowledge of complex and elaborate software with a steep learning curve, a feature that becomes more irritating as this software is often used by people who are not computer experts.

As a result, most planetary image data sets remain uncoregistered, even decades after the end of the corresponding space mission. For example, the 95 966 high-resolution images that Mars Orbiter Camera—Narrow Angle (MOC-NA) acquired of the surface of Mars between 1997 and 2006 [12] are still not coregistered to a common baseline, Lunar Reconnaissance Orbital Camera (LROC) [2] has acquired hundreds of terabytes of high-resolution images that are still not coregistered, while the Mercury Dual Imaging System—Narrow Angle Camera (MDIS-NAC) team coregistered 88 746 narrow angle images to create a global Mercury mosaic [13], ignoring the rest of the 277 928 high-resolution images that MDIS-NAC has acquired.

Multi-instrument coregistration, i.e., the coregistration of images from several different orbiter instruments to a single baseline, is a much more challenging task, since the images acquired by different instruments have generally different

Manuscript received December 29, 2016; revised May 16, 2017 and July 18, 2017; accepted July 24, 2017. Date of publication October 20, 2017; date of current version December 27, 2017. This work was supported in part by STFC under MSSL Consolidated Grant ST/K000977/1 and in part by the European Union's Seventh Framework Programme (FP7/2007-2013) under i-Mars Grant 607379. (Corresponding author: Panagiotis Sidiropoulos.)

The authors are with the Mullard Space Science Laboratory, University College London, London RH56NT, U.K. (e-mail: p.sidiropoulos@ucl.ac.uk; j.muller@ucl.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TGRS.2017.2734693

¹<http://www.harrisgeospatial.com/docs/RegistrationImageToImage.html>

²<http://davinci.asu.edu>

resolution ranges, distinct point spread functions, different spectral filters, different sensor techniques, and unique camera setups. The fact that they are generally separated in time (i.e., images were acquired over different time periods) causes additional problems both due to technological advances that limit multi-instrument compatibility and due to surface changes that may have happened in the meantime.

By focusing on Mars high-resolution cameras, the instrument diversity becomes apparent: Viking Orbiter (VO) acquired images between 1976 and 1980 in a resolution range from 8 to 1000 m/pixel with a frame camera that was using vidicon technology [14]; MOC-NA was a pushbroom camera that was active between 1997 and 2006 had a resolution of 1.5–12 m/pixel and was using charge-coupled device (CCD) technology [12]; the Thermal Emission Imaging System—Visible (THEMIS-VIS) is a thermal imaging system that acquires framelets of 17.5–75 m/pixel resolution [15]; High-Resolution Stereo Camera (HRSC) is a stereo pushbroom CCD camera that is mapping Mars since 2004 with a resolution 12.5–100 m [16]; Context Camera (CTX) builds upon MOC-NA legacy acquiring 6 m/pixel images from large areas of Mars [17]; and finally, High-Resolution Imaging Science Experiment (HiRISE) uses a specialized time delay and integration setup of 14 pushbroom cameras that acquire since 2006 images of 0.25–0.5 m/pixel resolution [18]. Such complexities have prevented the development of systematic multi-instrument planetary image coregistration algorithms. As far as we know, systematic multi-instrument coregistration is not supported by any of the current coregistration pipelines and algorithms.

The main obstacle that a coregistration algorithm has to overcome is the automatic estimation of a sufficient number of tie-points in the target (or input) and the baseline (or reference) image. Efficient, reliable, and fast tie-point estimation is difficult even in single-instrument coregistration. This challenge is further amplified when dealing with multi-instrument images, as has been extensively reported both in earth observation (see [19]) and planetary science literature (see [20]). However, if tie-point estimation is achieved, then several powerful techniques exist in the literature, which could convert the estimated tie-points to rigid, nonrigid, or hybrid camera models to produce coregistration results of high accuracy.

Accordingly, we have focused our attention on developing a novel point matching approach that uses *a priori* position estimations and standard geometric properties of planetary image acquisition setup to significantly increase the matching potential of local image descriptors. The *a priori* position estimations are the georeference information of both the reference and the target image, which are assumed to generate approximately correct position estimations, in order to automatically determine the reference image that will be used as a baseline. Moreover, in this paper, scale invariant feature transform (SIFT) [21] features are employed as image descriptors. SIFT features are extracted using the original algorithm and parameters introduced in [21]. The new algorithm is fast, does not require parameter tuning, and is robust to non-affine deformations that are common in planetary images. Hence, it is the

ideal basis for developing and implementing multi-instrument planetary image pipelines capable of performing batch-mode coregistration in a single pass and within a realistic time (with the main exclusion of surfaces with large slopes, which are quite rare on most planetary surfaces). As far as we are aware, this is the first systematic approach on multi-instrument coregistration of high-resolution planetary images.

This paper is organized as follows. Some basic principles of remote sensing image coregistration are revisited in Section II, before focusing on planetary image coregistration in Section III. The ring matching algorithm, which is the major novelty of this paper, is introduced in Section IV and further discussed in Section V, before being experimentally validated in several data sets coming from the Mars and the Moon in Section VI. Finally, Section VII concludes this paper.

II. REMOTE SENSING IMAGE COREGISTRATION

Over the last several years, remote sensing coregistration has achieved some progress by addressing several of its initial shortcomings. The relevant research spans from the theoretic analysis of the lower boundary of registration accuracy [22], [23] to *ad hoc* techniques that are tailored to specific types of imagery (e.g., optical to synthetic aperture radar images [24]). Despite this multitude, there are still a number of unresolved issues and a lack of consensus about the best strategies for coregistration applications.

In this section, we classify remote sensing coregistration techniques according to: 1) the user involvement; 2) the abstraction level of the image features used for coregistration; and 3) the inputs each could process, respectively. For a more thorough presentation of state-of-the-art remote sensing coregistration techniques, please refer to some of the literature reviews such as [25] and [26].

A. User Involvement

Despite the significant effort that has been put into the automation of remote sensing image coregistration, there are still many applications where the coregistration process is done manually, i.e., through some specialized software [27] that allows the user full control of the coregistration pipeline. However, manual coregistration is a tedious task that requires human resources, cannot be employed for batch-mode processing to coregister large volumes of input data, and, last but not least, it may lead to inaccurate tie-points when coregistering multisource images [28], [29].

As a consequence, a number of techniques have emerged, which employ a limited user input on the initial stages of the coregistration pipeline, before this proceeds automatically until the completion of the task. For example, in [30], semiautomatic image segmentation is used as a preprocessing step to geometrically constrain image matching, while in [29], manual crater detection generates geometrical constraints for planetary image coregistration. Additionally, a number of coregistration techniques in which parameter estimation is performed through least squares estimation (see [31]–[33]) require the approximate input image position

TABLE I
NNDR OPTIMAL VALUES THAT ARE REPORTED
IN THE RECENT LITERATURE

Method	[27]	[44]	[40]	[62]	[4]	[45]	[16]
NNDR	0.8	0.9	0.6-0.7	0.9	0.75	0.85	0.7-0.8

to be known (i.e., the images being originally approximately aligned) or provided by the user [28].

Nevertheless, the bulk of recent remote sensing image coregistration techniques are fully automatic, meaning that they do not require any user involvement apart from determining the input and the processing parameter values [34]–[37]. However, parameter sensitivity often limits the potential of an automatic method, since in the worst case scenario, parameter tuning should take place independently for each input image. An example of a sensitive coregistration parameter is the nearest neighbor distance ratio (NNDR) threshold, which determines if a point match is declared [38]. Table I is a (far from being exhaustive) list of the NNDR values suggested in the recent remote sensing coregistration literature, demonstrating the large range of the optimal NNDR value. It is clear that coregistration techniques with a performance sensitive to NNDR imply an extensive offline human interaction. Such techniques are not suitable for applications on which batch-mode image coregistration is required, i.e., automatic coregistration of large volumes of input imagery with minimum user involvement.

In order to emphasize the importance of remote sensing image coregistration techniques to perform batch-mode coregistration, we introduce another class of coregistration techniques, named *systematic*. Systematic techniques, apart from being robust to parameter tuning, are computationally efficient, do not require extensive computer resources, and are robust to the input image characteristics (i.e., resolution, data quality, and bit depth). Despite the progress in automatic image coregistration, currently there are hardly any systematic image coregistration techniques for remote sensing data. This point is further confirmed by the small-size data sets that are used in the relevant literature. To the best of our knowledge, [41] uses one of the largest experimental validation data sets in earth observation coregistration, consisted of 30 satellite images. Regarding planetary image coregistration, the state-of-the-art literature lacks any systematic techniques, with the few automatic techniques [42], [35] depending on the presence of surface features (such as craters) that are not universally found on all planetary surfaces.

B. Abstraction Level Used for Image Matching

The most common criterion that is used for the grouping of remote sensing image coregistration techniques is the primitive features employed in image matching. The resulting groups are labeled “area-based” and “feature-based” [25], “pixel-based” and “ground control points (GCPs)-based” [28], “intensity-based” and “feature-based” [34], and so on. In this paper, we introduce a new three-class (instead of the typical two-class) grouping of the coregistration techniques, according to the information content of the primitive features used

for matching. This grouping is more complete, encompassing methods that would be in the borderline of the previously employed coregistration classes.

More specifically, the related literature can be classified into the following.

- Low-level image matching techniques use noninformative (i.e., raw) data, often the original pixel values of the images. Matching is performed through cross correlation [31], mutual information [43], optimization [44], and so on. While such techniques have been extensively used over the last decades, they suffer from high computational complexity and are not suitable for large images, which are common in remote sensing in general and planetary science in particular. An alternative is to conduct matching in a transformation space, employing techniques such as shearlet transform [28], wavelet transform [45], and radon transform [46]. However, in this case, the gain in computational efficiency is attenuated by the inability to model complex geometric deformations, which are generally present in remote sensing images. Brigot *et al.* [47] model original misplacement as motion and proceed to coregister the images using a novel optical flow variation. This algorithm achieves fast and reliable coregistration of earth observation images; however, the mean residual magnitude is 4.2 pixels, an accuracy that is not deemed satisfactory for planetary image applications.
- Medium-level image matching techniques initially detect simple local features, before using them to achieve image matching. The employed features include well-known detectors/descriptors such as SIFT [21] and speeded up robust features (SURF) [48], as well as their recent variations that are tailored to remote sensing images [34], [39]. Medium-level techniques are the most common in the literature, since they combine accuracy and computational efficiency, even though they are still computationally demanding when dealing with large-size images [20]. Fast variations [49], [50] typically reduce the number of matched points, which may have a detrimental effect on accuracy, especially when the input images are acquired by different instruments [24], [20].
- High-level image matching techniques, in which the algorithm input is a set of semantically meaningful features such as roads in earth observation images [51], and craters and rocks in planetary images [35]. Such techniques are expected to be robust to noise and different acquisition characteristics, such as the resolution, the instrument point spread function, and the viewing angle. Unfortunately, not only is the construction of a set of universally present high-level features far from trivial, but also the density and distribution of high-level features depend on qualitative parameters that cannot be generally modeled (e.g., crater distribution on planetary images depends on the surface geological context).

C. Multisource and Unisource Methods

The majority of remote sensing image coregistration techniques are unisource, i.e., can process images only coming from the same instrument, due to algorithm properties that are

TABLE II

SIZE AND TIMES IN RECENT (EARTH OBSERVATION) COREGISTRATION LITERATURE. THE REPORTED SIZES ARE THE ONES OF THE LARGEST IMAGE PAIRS TESTED IN EACH PUBLICATION

Method	[40]	[55]	[39]	[16]	[24]
Largest Image Size (Mpixels)	4	5.8	4	0.25	0.25
Time (Minutes)	8	53.3	8	1.7	0.03

not consistent with multisource coregistration. For example, it has been extensively reported in the literature that the image matching method that was originally proposed in [21] fails in remote sensing image pairs from multiple instruments or/and image pairs from different bands of the same instrument [37], [52], [24], [20], [19]; therefore, a coregistration pipeline that adopts this method is limited to unisource input. Other factors that can adversely affect the multisource coregistration performance of a technique are the sensitivity to low-level image characteristics such as the point spread function, the bit depth, and the resolution, and the use of simple transformation models that fail to express the generally nonaffine deformations between multisource images.

On the other hand, there is a growing literature of multi-source coregistration techniques [41], [35], [45], even though their experimental validation is often conducted in small data sets that hinder the extraction of safe conclusions about their performance and usability. As a matter of fact, multisource coregistration has a number of unique characteristics and requirements that further complicate a task that is already very challenging; hence, its evaluation needs to be as thorough as possible, including multiple setups and a rich set of test images. In this paper, we conduct such an experimental analysis of our novel multisource systematic approach, which will be introduced after some theoretic issues of high-resolution planetary image coregistration are presented and discussed.

III. PLANETARY IMAGE COREGISTRATION

Coregistration of high-resolution planetary images has evolved to a domain that is rather distinct from the original remote sensing image coregistration, since planetary images (even though strictly speaking belonging to the general class of remote sensing data) share certain features that distinguish them from data collected by Earth Observation satellites. In this section, we present some of the most important planetary image features that should be taken into account by image coregistration algorithms. Additional complications arise from the fact that multi-instrument planetary images typically have different resolutions, SNRs, bit depths, and spectral ranges [16].

A. Size

The size and the computational time of some recent earth observation coregistration methods, presented in Table II, provide evidence that these techniques are developed to process images of size that is no more than 10 Mpixels, i.e., several orders of magnitude lower than high-resolution planetary images, which is typically larger than 30 Mpixels and may




Fig. 1. (a) Example of a planetary surface with lack of informative features. The area is from Mars and has coordinates $(182.81^{\circ}E, 48.14^{\circ}N)$. (b) Examples of a planetary surface with patterned terrain. The area is from Mars and has coordinates $(178.27^{\circ}E, -1.62^{\circ}N)$.

reach up to 10 Gpixels [20]. Additionally, most of them report results in image pairs of various sizes, which show an approximate linear computational complexity in relation to the total number of pixels. By linearly extrapolating the computational times of Table II to planetary image sizes, it becomes apparent that many earth observation coregistration algorithms, if used in planetary image coregistration applications, would fail to be executed in realistic time periods. Therefore, planetary image coregistration algorithms need to be meticulous about their associated computational complexity.

B. Lack of Informative Surface Features

Earth satellite images are often abundant with local surface (texture) features, such as buildings and roads, which enhance their information content. As a result, a number of powerful coregistration techniques have emerged that benefit from these features [53], [28], [51] to achieve high-accuracy performance. Similar to the most featureless earth regions (e.g., Arabian Desert), most planetary surfaces lack any informative surface features [Fig. 1(a)], thus making techniques based on informative features not suitable for planetary image coregistration.

C. Patterned Terrain

On earth, “patterned” terrain (i.e., with a local texture that is repeated for a large area) is quite rare, due to the large diversity of its geomorphology. On the contrary, patterned terrain is quite common on planetary surfaces [Fig. 1(b)]. Since most local image detectors/descriptors, such as SIFT [21] and SURF [48], model the local texture, techniques that perform point matching are expected to present a higher rate of false negatives (errors of omission) and false positives (errors of commission) rates in comparison with those of many earth observation applications (especially those related to urban landscapes), thus further complicating the accurate coregistration of planetary images.

D. Limited Knowledge of Planetary Processes

Planetary coregistration input consists of a pair of images that are generally acquired at different seasons, under different




Fig. 2. Example of two ARPs, $\text{ARP}(P_1, \delta, 6)$ and $\text{ARP}(P_2, \delta, 6)$. Both the matching points Q_1 and Q_2 lie in the same ring ($k = 5$).

weather conditions and with a time distance on the order of years or even decades. Furthermore, any active surface–atmospheric interaction processes on planetary bodies, which affect the image content, are not as well understood as the corresponding processes on earth. As a result, planetary images that map the same area at different times may be exposed to geometric and photometric differences that cannot be necessarily modeled. A systematic planetary coregistration algorithm should therefore take into account the presence of this “spatial noise,” thus avoiding rigid geometric models, such as those based upon rotation–scaling–translation and affine transformations (see [37], [54], [46]).

IV. RING MATCHING FOR PLANETARY IMAGE COREGISTRATION

A fact that is typically adopted by the community, even with a lack of theoretic confirmation, is that the translation error is the error with the largest magnitude in remote sensing mis-registration. Based on this, in this section, a ring matching algorithm is proposed, which is based on the assumption that the total mis-registration error $E(x, y)$ in a pixel $P(x, y)$ is the sum of a global translation error C and a local mis-registration error $E_l(x, y)$, of smallest magnitude

$$E(x, y) = C + E_l(x, y), \quad \|E_l(x, y)\| \ll \|C\|. \quad (1)$$

It is assumed that both the reference and the target image are georeferenced (before coregistration) and the point $P_t(x, y)$ in the target image has the same coordinates with the point $P_r(x, y)$ in the reference image. A number of M concentric circles with their center on $P_r(x, y)$ and radius $k\delta$, $k = 1, 2, \dots, M$ partition the reference image into M rings, each bounded by a larger circle of radius $k\delta$ and a smaller circle of radius $(k-1)\delta$ (Fig. 2). We call this partition an adaptive ring partition $\text{ARP}(P, \delta, M)$. Note that the ARP is distinct for each pixel of the target image, i.e., $\text{ARP}(P, \delta, M) \equiv \text{ARP}(P', \delta, M)$ if- f $P = P'$. The coregistration potential of ARP is based on the following proposition.

Proposition 1: If a point $P(x, y)$ of the target image matches to a point $Q(x', y')$ in the reference image that lies in the k th ring of $\text{ARP}(P, \delta, M)$, then a point $P'(x, y)$ of the target image matches to a point in the reference image that also lies in the k -th ring (or in the $k-1$ -th and $k+1$ -th ring).

Proposition 1 derives from (1), if δ is selected sufficiently large so that $\delta > \|E_l(x, y)\|$. In this case, since $Q(x', y')$ lies

in the k th ring of $\text{ARP}(P, \delta, M)$, $(k-1)\delta < E_P(x, y) < k\delta$, from which it follows that $((k-1)\delta - \|E_l(x, y)\|) < C < (k\delta + \|E_l(x, y)\|)$ and finally that $(k-2)\delta < E_{P'}(x, y) < (k+1)\delta \forall P'(x, y)$.

Proposition 1 implies that image matching can be performed independently in each ring of the ARP, generating in total M sets of matched points (one for each ring). All correct matches (along with outliers) are contained in l sets ($1 \leq l \leq 3$), corresponding to l adjacent rings. The rest $(M-l)$ sets of matched points are composed exclusively of outliers. The rings containing all true positive matches can be detected from the fact that for all other rings the outlier rate would be 100%.

More specifically, if two points P_1 and P_2 in the target image are, respectively, matched with two points Q_1 and Q_2 in the reference image, then $\|Q_2 - Q_1\| = \|P_2 - P_1\|$. The distance $\|Q_2 - Q_1\|$ can be straightforwardly estimated, since the reference image is orthorectified; hence, the ratio between $\|P_2 - P_1\|$ and the pixel distance between P_1 and P_2 would determine the average resolution, s_a in the profile connecting P_1 and P_2 . s_a is expected to be approximately equal to the nominal image resolution s_n , estimated from the ephemeris data of the target image, i.e., $s_a/s_n = 1 \pm \epsilon$, where ϵ is the tolerance. On the other hand, if at least one of the matches is incorrect, then s_a is random, taking values in a much larger range. Consequently, the set of correct matches can be identified by initially estimating s_a for all pairs of matched points in the target image and subsequently identifying points for which the equation $s_a/s_n = 1 \pm \epsilon$ stands for a large point subset.

This approach is implemented in Algorithms 1 and 2. In the first stage (Algorithm 1), SIFT points in the target image are matched one at a time and using an ARP with SIFT points in the reference image until a ring with sufficient number of correct matches is identified. At that time, the identified ring along with the set of preliminary tie-points constitutes the input to the second stage of the algorithm, in which the set of tie-points is augmented.

Algorithm 1 First Stage of the Ring Matching Algorithm

Input: Two sets of (SIFT) features P and Q in the target and reference image; prior georeferencing information of the reference and the target image, including the target image nominal resolution s_n ; parameter values δ, M, ϵ , and X .

- 1: Select a point P_i in the target image feature set
 - 2: Estimate the $\text{ARP}(P_i, \delta, M)$ partition of the reference image and assign each point of Q to a single ring
 - 3: Perform independently feature matching in each ring. For each ring k , identify the point Q_{ik} that matches to P_i .
 - 4: For each pair of $(Q_{ik}, Q_{jk}), i \neq j$ estimate the implied average resolution s_{kij} of the line segment $P_i P_j$.
 - 5: If there is a ring k and a point i for which $(1 - \epsilon) \leq (s_{kij}/s_n) \leq (1 + \epsilon)$ for more than X values of j , then the matching is over; the output is a set W consisted of pair (P_i, Q_{ik}) as well as all pairs (P_j, Q_{jk}) for which $(1 - \epsilon) \leq (s_{kij}/s_n) \leq (1 + \epsilon)$. Otherwise, the algorithm continues from step 1.
-

Algorithm 2 Second Stage of the Ring Matching Algorithm

Input: Two sets of (SIFT) features P and Q in the target and reference image; prior georeferencing information of the reference and the target image, including the target image nominal resolution s_n ; parameter value ϵ ; identified ring k ; and a set of matches W from Algorithm 1.

- 1: For each point in P , P_i , declare the area in the reference image that is bounded by the $(k - 1)$ th and the $(k + 1)$ th ring as the “valid area” A_i .
- 2: Identify the point Q_i , from the subset of Q that lie in A_i , that match with P_i .
- 3: Estimate the average resolution s_{ij} where P_j is selected from the set of preliminary tie-points W and check whether $(1 - \epsilon) \leq (s_{ij}/s_n) \leq (1 + \epsilon)$ for all P_j in W , in which case store the pair (P_i, Q_i) to a new set of matches W' .
- 4: After the process is terminated, concatenate matches W and W' .

More specifically, in the second stage (Algorithm 2), the points in the target image are matched to the points in the reference image that lie in the area that is bounded by the $(k - 1)$ th and $(k + 1)$ th rings, according to Proposition 1. After declaring a match, the average resolutions of all profiles with one endpoint to the new match and the other endpoint to a tie-point estimated by Algorithm 1 are compared with the nominal resolution and the match is discarded or accepted, accordingly. Finally, the sets of preliminary matches and new matches are concatenated, thus containing the set of tie-points used in the subsequent coregistration steps.

Ring matching algorithm generates a set of tie-points that are further processed, depending on the application setup and the data availability. For example, if a digital terrain model (DTM) is available along with the orthorectified baseline (as is the case in HRSC level-4 data [16]), then the 3-D georeferencing information can be used to generate a rigorous pushbroom camera model [55], a rational polynomial model [56], etc., while if not, coregistration can be performed through the estimation of transformation parameters (e.g., rotation, translation, scaling, etc.) using least squares. As a matter of fact, there are several coregistration techniques that can achieve subpixel coregistration performance if supplied with large and accurate tie-point sets. However, the analysis of such models is beyond the scope of this paper, as is the comparison of SIFT with recently introduced local detectors/descriptors, such as the MUR-SIFT [34] and the adaptive binning SIFT [39]. Ring matching algorithm does not depend on the employed local feature, and thus we consider this research orthogonal to ours.

V. TECHNIQUE PROPERTIES

A. Parameter Sensitivity

The parameters that ring matching insert to the automatic coregistration process are the ARP size M , the ring range δ , the tolerance ϵ , and the correct match threshold, required for the algorithm to pass from the first to the second stage, X . In this section, we examine all of them in relation to their sensitivity,

since parameter sensitivity prohibits the systematic coregistration of large volumes of multi-instrument images without using human resources, even when the employed algorithm is fully automatic.

The ARP size M is used to reduce the coregistration computational cost, since the accuracy of the position and orientation ephemeris data limits the magnitude of the mis-registration error. As a matter of fact, the largest mis-registration that we have experimentally found (in the coregistration of a VO high-resolution image to HRSC baseline) was 14.6 km. In order to nullify the impact of even double mis-registration errors, M has been selected so as the outer ring radius to be 30 km, a hard-coded value suggested for the systematic coregistration of planetary images, even of low position and orientation accuracy.

The ring range δ is more difficult to optimize, since it should compromise between high outlier ratio (for large δ values) and detection of points in a very limited mis-registration range, which may lead to severe terrain distortions (for small δ values). Optimizing δ is easier in single-instrument coregistration applications, since in this case, the outlier ratios are expected to be rather low. Even in multi-instrument coregistration applications, the experimental section will demonstrate that the algorithm achieves a near-optimality for large ranges of δ , even though there is no universal optimal value for δ .

The tolerance ϵ also affects the outlier ratio, as well as the uniformity of the tie-point data set. More specifically, the average resolution ratio criterion is more strict in neighboring Q_{ik} and Q_{jk} than in distant ones. A side effect of this approach is that a small δ combined with a high tolerance ϵ could lead to the acceptance of outliers as correct matches in cases where k is small and points Q_{ik} and Q_{jk} distant. On the other hand, a much too rigorous tolerance ϵ would increase the false rejection rate, especially in cases where significant geometric distortions are expected (e.g., if an image was acquired with high emission angle).

This issue can be resolved by selecting a high value of the threshold X , thus minimizing both the false acceptance and false rejection ratio and relaxing the tuning of δ and ϵ . The forfeit in this case is the increase in computational cost. However, as will become apparent in the subsequent analysis, the introduced algorithm is not computationally demanding, and hence, a mild increase in the computational cost is not expected to be of significant importance in most planetary image coregistration applications.

It should be noted that most of the parameter sensitivity discussion is conducted for theoretic reasons, since one of the key advantages of the introduced technique is its parameter robustness. As a matter of fact, it has been experimentally found that a single parameter set ($M = 30$, $\delta = 2$, $\epsilon = 0.02$, and $X = 15$) could be used to successfully coregister more than 60% of each data set used in this work (apart from VO, for which the success rate is higher than 50%). In practice, the parameter selection stage of the processing is reduced to a brute-force search in a low-dimensional discrete parameter space, using a small-sized validation set. More specifically, in all the experiments included in this paper, $M = 30$ and the other three parameters are selected between 27 parameter

combinations (three for each parameter), using a set of five images. This causes an overload of 135 coregistrations, which is minimal in comparison with the typical planetary image data set size (e.g., it is 0.07% of THEMIS-VIS volume, 0.14% of MOC-NA volume, etc.). As it will become apparent in Section VI, this simple parameter selection strategy is enough to accurately coregister large volumes of multi-instrument data to a common baseline.

B. Computational Complexity

Sidiropoulos and Muller [20] have demonstrated that the computational complexity of matching feature points from large images may become prohibitive (exceeding several days per image pair) if it is conducted as originally suggested in [21]. The introduced technique ameliorates the image matching computational complexity through the following features.

- 1) The finite maximum radius size of ARP (equal to $M\delta$) employs the prior georeferencing information to impose distance constraints on the points that can constitute a valid match, thus reducing the number of candidate matches (i.e., points that need to be checked).
- 2) The first (and more computationally demanding) stage of the ring matching algorithm is performed at one-match-at-a-time basis, until the number of correct matches implies the valid ring k and a preliminary set of tie-points W is established. Subsequently, the matching becomes even faster by ignoring most rings from further processing.
- 3) The outlier detection step (Algorithm 1, step 5), even though it is quadratic in relation to the number of preliminary matches, is very fast because it involves only simple computations (e.g., subtractions and multiplications). As a matter of fact, it can be shown that the estimation of s_{kij} requires 15 elementary computations, i.e., approximately 5.8% of the computations needed to estimate the distance of two 128-D SIFT points.

Although a formal computational complexity analysis depends on *ad hoc* image characteristics, such as the number of estimated feature points, the image sizes, and the outlier ratio, an example using typical values could help clarify the computational gain that is achieved through ring matching. In this example, it is postulated that a CTX high-resolution image with 510^5 SIFT points is matched to a HRSC baseline with 210^6 SIFT points. Moreover, it is postulated that for each CTX SIFT point, 10% of the HRSC points are within a 30-km radius. A simple image matching as suggested in [21] would include 10^{12} comparisons of 128-D SIFT points. An ARP with $M = 30$ km, $\delta = 500$ m, $X = 20$, and 20% of the target points being matched in the first stage of ring matching with an outlier rate of 90% would require 210^8 SIFT-point comparisons and 310^7 s_{kij} estimations (which correspond to approximately 210^6 SIFT-point comparisons). In the second stage of ring matching, all CTX points are compared on average with 10^4 HRSC points, leading to 210^9 SIFT-point comparisons. In total, the 2.202 billion SIFT-point comparisons are approximately 450 times less than the 1 trillion comparisons required if image matching is conducted according to [21].

Further experimental results about the coregistration computational time are given in Section VI.

C. Robustness

A prerequisite for a systematic coregistration technique is to be robust against common input-image deficiencies, such as mediocre georeferencing accuracy and low visible image quality. It is noted that georeferencing sensitivity is scarcely addressed in the remote sensing coregistration literature, even though its performance often depends on a rather accurate prior position estimation [28]. On the other hand, the only georeferencing hard constraint of ring matching is the initial positioning error not being larger than tens of kilometers. Other than that, ring matching performance in poorly georeferenced images is affected by the fact that the area of the outer rings is larger than the one of the inner rings; therefore, the outlier rate is typically a more important issue. However, as will be demonstrated in the experimental section, the overall performance is usually only slightly deteriorated. Additionally, a simple adjustment that may resolve this issue, when extra diligence is required, is to define δ as a decreasing function of the ring index k .

An increase in the outlier rate, caused by the decline of true positive matches, is the main ramification of low-quality data, regardless if the quality degradation origins are external (e.g., atmospheric dust) or internal (e.g., instrument malfunction). While the performance deterioration when coregistering low-quality data is unavoidable, ring matching seems nevertheless exceptionally robust to low-quality data, due to the fact that the average resolution comparison criterion can produce correct results even when the outlier rate is much higher than 90%. Actually, the outlier rate may prompt the coregistration failure either by decisively raising the (quadratic in relation to the point matches) computational time of the outlier detection step of the first stage of the algorithm, or by not establishing enough true positives in the full set of target-image feature points. Empirical evidence has shown that the new coregistration technique can regularly overcome both obstacles for outlier rates up to 98%, while the failure probability exceeds 50% for outlier rate approximately around 99%. This extreme outlier robustness plays central role in the algorithm's potential to systematically coregister large volumes of multifarious data, as will be demonstrated in the next section.

Finally, the Achilles heel of the new technique is its sensitivity to the target image nominal resolution s_n . It is understood from the algorithmic description that the successful coregistration of a planetary image requires rather accurate prior knowledge of its resolution. This condition may inhibit the transfer of this method to close-range coregistration applications (e.g., the coregistration of aerial images). However, is almost always satisfied in planetary images, since it mainly depends on the spacecraft trajectory, which in general is very precisely known.

VI. EXPERIMENTAL VALIDATION

Before passing to the actual experimental validation, some basic properties of the experimental setup need to be clarified.

TABLE III
DATA SETS USED FOR ALGORITHM VERIFICATION

Id	Camera	Resolution (m)	Years	Baseline	Base Resolution (m)	Target	No. Images
1	CTX	6	2006-2016	HRSC ORI	12.5	Mars	100
2	CTX	6	2006-2016	HRSC Mosaic	12.5	Mars	50
3	THEMIS-VIS	17.5-75	2002-2016	HRSC ORI	12.5	Mars	75
4	THEMIS-IR	100	2002-2016	THEMIS-VIS	17.5-35	Mars	75
5	MOC-NA	1.5-12	1997-2006	HRSC ORI	12.5	Mars	50
6	MOC-NA	1.5-12	1997-2006	HRSC Mosaic	12.5	Mars	50
7	Viking Orbiter	8-100	1976-1980	HRSC Mosaic	12.5	Mars	50
8	LROC-NAC	0.5-1	2009-2016	LROC-NAC	0.5-1	Moon	50

First, it should be noted that the systematic coregistration of planetary images necessitates considering the employed parameter set as hard-coded for all images of a single data set (but not for different data sets). As a result, parameter tuning is not conducted at the image level but rather at the data set level. Second, the computational time allocated to the coregistration algorithm execution is finite, instead of allowing the coregistration to continue uninterrupted. This property is fundamental for the batch-mode processing of planetary images, in which the overall processing time is an important aspect of the application. Additionally, planetary image coregistration, unlike earth observation image coregistration, suffers from a lack of GCPs for which the coordinates are precisely known. Therefore, coregistration accuracy is estimated by randomly dividing the set of estimated tie-points (after outlier detection) into two equal parts: the first part is used for coregistration and the second for assessing its accuracy.

Finally, in all of our experiments, we have employed three measures of performance:

- 1) the failure rate, i.e., the percentage of images that the algorithm failed to coregister within a specific processing time and using a predetermined offline parameter set;
- 2) the mean accuracy in both X and Y directions, measured in meters;
- 3) the overall computational time, including preprocessing, feature (SIFT) point extraction, ring matching, and coregistration.

A. Algorithm Verification

This section deals with the experimental validation of the presented algorithm. Three key issues are examined: 1) the coregistration accuracy, including an assessment of the assumption that global translation is the most severe mis-registration error; 2) the number of tie-points that are estimated by ring matching, which is fundamental for the further processing of a coregistration algorithm; and 3) the computational cost of the algorithm. This series of experiments was conducted on a multifarious set of 500 images, selected from eight high-resolution planetary image scenarios (Table III). VO images with resolution coarser than 100 m/pixel are ignored, since for them, the resolution difference with the HRSC baseline (with resolution 12.5 m/pixel) makes their coregistration impractical.

Additionally, in order to be consistent, coregistration does not involve either DTM information or laser altimeter height

TABLE IV
RANK STATISTICS OF THE ORIGINAL AVERAGE TIE-POINTS' MIS-ALIGNMENTS (IN METERS) FOR THE EIGHT COREGISTRATION SCENARIOS OF TABLE III. THREE RANK STATISTICS ARE REPORTED: THE 25% PERCENTILE, THE 50% PERCENTILE (I.E., THE MEDIAN VALUE), AND THE 75% PERCENTILE. THE OUTLIERS WERE IDENTIFIED AND DISCARDED BEFORE MIS-ALIGNMENT WAS ESTIMATED

Id	Camera	25% Perc.	50% Perc.	75% Perc.
1	CTX	207.26	242.89	376.11
2	CTX	201.09	261.14	383.25
3	THEMIS-VIS	569.29	617.83	681.44
4	THEMIS-IR	189.01	258.95	406.47
5	MOC-NA	158.53	184.44	920.63
6	MOC-NA	125.35	231.25	902.26
7	Viking Orbiter	1233.26	1928.57	2571.29
8	LROC-NAC	30.63	34.59	48.26

measurements, since neither of these are available in all data sets. The reported experimental results refer to image-to-image coregistration, in which a polynomial function is used to transform the coordinates of the target image to the coordinate system of the reference image.

1) *Accuracy*: Table IV summarizes the original mis-alignment between the reference and the target image in the eight included coregistration scenarios. The reported rank statistics for each data set are the first and third quartile error values, as well as the median error, which were estimated from the average geometric distance of the reference and target image matched points (tie-points), after discarding the outliers.

A first conclusion that can be drawn from Table IV is that, regardless of the application setup, planetary image coregistration is far from being considered of trivial importance. As a matter of fact, spacecraft housekeeping progress has reduced the mis-alignment for approximately one order of magnitude over the last three to four decades (from VO up to the latest orbiters). However, the current mis-registration error is expected to be much larger than 10 pixels for images acquired by different instruments (e.g., 40.1 pixels for CTX/HRSC) and significant even for images acquired by the same instrument (e.g., 2.6 pixels for THEMIS-IR/THEMIS-VIS); hence, the housekeeping data are still far from making image coregistration obsolete (while the planetary data volume signifies that coregistration needs to be resolved from the development of systematic coregistration algorithms).




Fig. 3. Example of a large mis-alignment between two CTX images before their coregistration to an HRSC baseline. (Top) CTX images P01_001559_1889 superimposed on the CTX image P03_002047_1889 with 50% transparency. The fact that features are repeated twice (e.g., the dark-sand leaf-type feature in the center of the image is repeated near the left edge) demonstrates the clear misalignment. (Bottom) Same images are superimposed with 50% transparency after coregistration, showing the actual appearance of the site. Note that the diagonal line in the left edge of the top panel (center of the bottom panel) designates the boundary of the two images.

Additionally, the comparison of the median with the first and third quartile indicates that the error distribution is positively skewed. The right tail of the error distribution can be mainly attributed, first, to the sporadic release of poor position and orientation estimates from instrument teams [Fig. 3 (top)] and, second, to the oblique images that some instruments acquire to allow stereo processing. While the poor quality of the housekeeping data cannot be counterbalanced, the misalignment of a stereo target image to an orthorectified baseline can be roughly modeled out from the viewing angle and the spacecraft trajectory, which are both typically known with an acceptable accuracy. This minor adjustment to the original ring matching algorithm has been applied to 14%, 12%, 8%, and 8% of data sets 1, 2, 5, and 6, respectively.

Regardless of the preprocessing, ring matching was applied to examine the validity of (1). Table V reports three median errors per data set; the original median mis-alignment, i.e., the 50% percentile value of Table IV; the median mis-alignment if only global translation is used for coregistration; and the median residuals after full ring matching-based coregistration is performed [Figs. 4 and 3 (bottom)]. It becomes apparent from Table IV that in all eight data sets, the global trans-

TABLE V
MEDIAN MIS-ALIGNMENT ERRORS (IN METERS) FOR THE EIGHT COREGISTRATION SCENARIOS OF TABLE III. “ORIGINAL” COLUMN REFERS TO TABLE IV, “TRANSLATION” TO THE MIS-ALIGNMENT ERROR WHEN ONLY GLOBAL TRANSLATION IS APPLIED TO THE TARGET IMAGE, AND “COREGISTRATION” TO THE RING MATCHING RESIDUALS

Id	Camera	Original	Translation	Co-Registration
1	CTX	242.89	38.56	8.2
2	CTX	261.14	29.61	7.73
3	THEMIS-VIS	617.83	37.04	12.47
4	THEMIS-IR	258.95	83.23	58.1
5	MOC-NA	184.44	35.71	6.89
6	MOC-NA	231.25	33.13	6.58
7	Viking Orbiter	1928.57	99.75	81.17
8	LROC-NAC	34.59	1.87	0.64

lation term dominates the mis-alignment error, as expected from the analysis of Section IV. However, the accuracy of a technique that would ignore all other types of mis-alignments would usually be larger than 1 baseline pixel. For example, given that HRSC (which is the baseline in five coregistration cases) has a resolution of 12.5 m/pixel, the global translation reduces the residuals to 2–3 (HRSC) pixels for more than half of the images, but subpixel accuracy is typically achieved only through a more elaborate approach that would take into account additional mis-alignments terms, such as the introduced ring matching algorithm. It should be noted that subpixel accuracy is reached despite the poor *a priori* position estimation, thus demonstrating ring matching robustness to initial positioning error.

2) *Number of Tie-Points*: It is widely accepted that a larger number of tie-points is associated with a higher coregistration quality, mainly because coregistration robustness to outliers is increased with tie-point quantity [28] but also because averaging attenuates spatial noise. The number of identified tie-points clearly depends on the image size; therefore, it is sensible not to compare absolute numbers but the “tie-points per megapixel” rate, instead. Additionally, following the rationale of the previous section, rank statistics are used to analyses the tie-point abundance for each data set.

The number of tie-points for each data set can be found in Table VI. Considering the lack of informative surface features of planetary images (in comparison with earth observation images), as well as the state-of-the-art tie-point rate of earth observation coregistration (e.g., 15–108 tie-points per megapixel for MUR-SIFT [34]), it can be concluded that ring matching coregistration generates an abundant set of tie-points for the great majority of planetary images.

An exception is the case of 12, out of 75, THEMIS-IR images (coregistered to THEMIS-VIS baseline), for which the number of tie-points is less than 10. While a definitive conclusion would require an analysis that is not within the scope of this paper, it seems that this failure mainly originates from the low quality of a part of the THEMIS products as well as the distinctly different surface reflectance values in infrared and visible spectra.

Additionally, Table VI shows an improvement in tie-point quantity when an HRSC mosaic is used (data sets 2 and 6),




Fig. 4. 1700 × 1200 m detail of two LROC-NAC Moon surface images, after ring matching coregistration. (a) M129941166LE. (b) M118548416LE.

TABLE VI

RANK STATISTICS OF THE NUMBER OF TIE-POINTS PER MILLION PIXELS IDENTIFIED THROUGH RING MATCHING, FOR THE EIGHT COREGISTRATION SCENARIOS OF TABLE III. THREE RANK STATISTICS ARE REPORTED: THE 25% PERCENTILE, THE 50% PERCENTILE (I.E., THE MEDIAN VALUE), AND THE 75% PERCENTILE

Id	Camera	25% Perc.	50% Perc.	75% Perc.
1	CTX	413.05	1311.76	2308
2	CTX	458.64	1562.31	2696.73
3	THEMIS-VIS	264.08	447.51	584.12
4	THEMIS-IR	13.18	45.06	56.33
5	MOC-NA	61.2	102.97	130.54
6	MOC-NA	160.48	204.62	592.66
7	Viking Orbiter	55.81	135.12	202.48
8	LROC-NAC	1068.42	1607.9	2518.19

instead of single HRSC strips (data sets 1 and 5, respectively), which implies a dependence from the baseline image quality, since the HRSC single strips include both good and poor quality images, while the released HRSC mosaic of Mars Chart 11—East (MC11-E) [57] generally exhibits excellent quality. The induced coregistration accuracy improvement is approximately 5% (Table V).

Apart from the number of tie-points, the coregistration performance depends on the (uniform) distribution of the tie-points on the image. An algorithm that generates tie-point sets that are confined to a small part of the target

TABLE VII
COMPARISON OF TIE-POINT DISTRIBUTION USING FULL MATCHING AND RING MATCHING

Id	Camera	Q_D (full matching)	Q_D (ring matching)
3	THEMIS-VIS	0.13	0.29
5	MOC-NA	0.07	0.42
6	MOC-NA	0.08	0.36

image may cause large mis-registration errors on other parts of the image, especially those far from the tie-point location. The distribution uniformity of N tie-points may be modeled by the average pairwise distance between tie-points, d_t , in comparison with the average pairwise distance of N points that are uniformly selected in the image, d_u . The latter is not trivial to be generically estimated, since the swath orientation and size as well as the spacecraft trajectory usually impose an image plane that is similar to a rotated rectangle, with at least three degrees of freedom. In order to overcome this, we follow a Monte Carlo approach; several instances of N random points are selected and the average d'_u is returned as the d_u value. Finally, the ratio $D_Q = d_t/d_u$ is used to measure the uniformity of the tie-point distribution.

Using the aforementioned evaluation scheme, the presented algorithm was compared with the tie-point distribution of an algorithm that performs full-image matching, i.e., it uses the (full image) point matching approach of [21] instead of ring matching. Due to the extended computational time of full-image matching, only data sets 3–7 are considered. Moreover, as thoroughly described in [20], full-image matching fails to produce any tie-points in most multi-instrument pairs. These are ignored from the evaluation, including all image pairs of data sets 4 and 7. The mean ratio Q_D for the rest of the images, demonstrated in Table VII, confirms that ring matching not only increases the number of tie-points but also produces tie-points that are more evenly distributed on the image plane.

3) *Computational Complexity*: The computational time required for image coregistration generally depends on the input image size (usually through a linear relationship between the total pixel number and the computational time) as well as the machine that the processing is conducted. In order to nullify the effect of the latter, the computational time unit that we employ is the “mean computational time for an image of size 1 megapixel” T_1 . All the reported computational times T thence are actually rates $T = T_s/T_1$, where T_s is the coregistration computational time, counted in seconds.

Additionally, in the computational complexity analysis, all eight datasets have been jointly analyzed, since the processing time of ring matching has been experimentally found to be much more affected by the image sizes than any other image properties (such as the bit depth, the point spread function, the illumination angle, etc.). Finally, because ring matching computational time is additionally related to the correct matches threshold X , the processing has been repeated several times with an increasing X . The stored computational time T was the one associated with the run that employed the smallest X value that achieved subpixel accuracy.

Based on the above, the computational times were binned according to the image size, to 1 Mpixel bins, and the




Fig. 5. Computational complexity of ring matching coregistration as a function of pixel size.

average value of each bin is plotted in Fig. 5. The linear regression estimate of the computational time T as a function of image size N (counted in Mpixels) is $T = 0.08N + 0.38$. This equation shows that ring matching computational time increases mildly with image size, thus making it suitable for most planetary image sizes. For example, the coregistration of an image of 20 Mpixels (typical MOC-NA image size) takes double the time of T_1 , while the computational time of an image of 250 Mpixels (typical CTX image size) requires approximately $20T_1$ time. Considering that T_1 is on the order of minutes even for common workstations, the ring matching algorithm can be concluded as completing in a realistic time for the vast majority of high-resolution planetary images.

B. (Batch-Mode) Systematic Coregistration

As already said, the introduced ring matching technique is a systematic coregistration technique, since it is capable of automatically processing large volumes of multifarious data without requiring an offline parameter tuning stage for each image. This property, which makes ring matching coregistration unique among the available planetary coregistration techniques, was established through a realistic coregistration scenario, involving a large region of Mars, where 7056 images have been acquired.

More specifically, we have coregistered all images acquired from CTX, THEMIS-VIS, MOC-NA, and VO in Martian MC11-E half-quadrangle. MC11-E is the east half of the Oxia Palus quadrangle, extending between 0° – 30° north and 0° – 22.5° west (or 337.5° – 360° east), and contains a number of the most geologically interesting regions of Mars, such as Chryse Planitia, Xanthe Terra, Mawrth Vallis, and Meridiani Planum. Lately, the HRSC team released a mosaic of MC11-E, with a panchromatic resolution of 12.5 m/pixel, while the corresponding DTM has a resolution of 50 m/pixel [57]. This mosaic used 89 single-strip high-resolution HRSC nadir images [57] and tie-points generated from individual strips as an input, which were subsequently passed by an elaborate processing chain, including bundle block adjustment [58] and radiometric correction [59].

Apart from the recently released HRSC mosaic, the MC11-E quadrangle has been extensively imaged by all the NASA high-resolution visible cameras. The statistics of MC11-E image

TABLE VIII
HIGH-RESOLUTION ($\text{RES} \leq 100$ M/PIXEL) ORBITAL IMAGES
OVER MC11-E HALF-QUADRANGLE OF MARS

Camera	CTX	MOC-NA	THEMIS-VIS	VO
MC11-E Images	1,365	1,558	3,629	504

coverage with images of resolution finer than 100 m/pixel are shown in Table VIII, excluding HiRISE images, for which the resolution difference from the HRSC baseline is too large for the coregistration to take place.

The 7056 high-resolution orbital images coregistered to the HRSC base mosaic constitute approximately 1.6% of all high-resolution images of Mars. This percentage is further increased to 3.3% if we ignore the regions of Mars for which there is currently no HRSC 3-D-model, i.e., there is no basemap available. This imagery was coregistered using ring matching with a single set of parameters for each data set that was selected by optimizing the coregistration performance in a subset of 10 images per dataset. Additionally, since the realistic computational time is central to the systematic processing of large volumes of data, both stages of the ring matching algorithm were limited to 1 h each. If the 1-h computation time expired in the first stage of ring matching, then the processing was aborted and the image coregistration is flagged “failed,” while if it expired in the second stage, the processing continues with the set of tie-points that were already identified.

In this application, the tie-point set estimated through ring matching is the input in a processing pipeline that makes use of the HRSC areo-reference information, as well as the corresponding DTM to optimize the coregistration accuracy. More specifically, the HRSC coordinates of the tie-points (which are originally in pixel coordinates both for the target image and for the HRSC baseline image) are transformed to 3-D world coordinates using the areo-reference information and the corresponding HRSC DTM. Thus, a number of correspondences between the target image pixels and their position in 3-D world coordinates are established. These are used to estimate a camera model for the target image, which subsequently determines the projection of the target image onto the common coordinate system. The employed camera model is a combination of a rigorous camera model [55] and a polynomial model. The estimated camera model is the one that is finally used to produce the output of the processing chain, i.e., the coregistered version of the target image. It should be noted that, by comparing the accuracy of this method with the accuracy of an image-only polynomial coregistration scheme in data sets 2 and 6 of the previous section, we have found an improvement of 20% in the achieved accuracy.

The average accuracy (ErrX and ErrY) and the failure rate are demonstrated in Table IX. In total, 5356 out of 7056 images were successfully coregistered to a common baseline with an average subpixel accuracy, using a single parameter set and requiring a limited amount of computational time. This performance confirms that ring matching can be used for the systematic processing of large volumes of high-resolution planetary images, combining high accuracy, limited computational time, and minimum use of




Fig. 6. Detail of Becquerel crater on Mars (22.1° N, 352.0° E), as imaged from four different (fully automatically) coregistered images. (a) CTX image P01_001546_2016. (b) CTX image B19_017211_2015. (c) CTX image G19_025637_2021. (d) HRSC MC11-E mosaic [57].

TABLE IX
STATISTICS OF THE SYSTEMATIC COREGISTRATION OF HIGH-RESOLUTION IMAGES TO THE HRSC MC11-E MOSAIC

Camera	CTX	MOC-NA	THEMIS-VIS	VO
ErrX (m)	6.487	5.334	7.012	12.211
ErrY (m)	6.081	4.859	6.849	11.71
Fail. Rate (%)	7.25	34.51	19.84	43.06

human resources. An example of three CTX coregistration results and the corresponding HRSC baseline is shown in Fig. 6. All coregistered images can be accessed through i-Mars WebGIS.³ i-Mars WebGIS hosts also a 6 m/pixel mosaic of MC11-E that was produced by photometrically correcting the set of coregistered CTX images, without altering their geometry [60] (Fig. 7). This product (which will be thoroughly discussed in a future publication) further validates the accuracy of ring matching coregistration, while demonstrating its potential to significantly reduce the effort that is required for the release of high-resolution planetary image mosaics.

The failure rate in MC11-E data set, which is processed without requiring an offline parameter tuning step, is in line with the potential of the presented algorithm to be used for the systematic processing of high-resolution planetary images. In order to provide more evidence, the failed images were reprocessed using the coregistered CTX imagery as a baseline, without changing any of the algorithm parameters. Note that

typical CTX imagery is of higher quality than HRSC, having a higher SNR and bit depth. This processing reduced the failure rate from 7.25% to 1.31% for CTX images, from 19.84% to 2.88% for THEMIS-VIS images, from 34.51% to 4.55% for MOC-NA images, and from 43.06% to 23.17% for VO images. Therefore, it can be deduced that the algorithm failure is mostly determined by the quality of the baseline, while being robust to parameter selection. Since a single set of parameters is sufficient to coregister a large-volume data set, the algorithm parameters can be hard-coded within the implementation, thus removing the need for parameter tweaking and significantly reducing the human resources required for large-scale planetary image coregistration.

C. Limitations

While ring matching is a very powerful coregistration technique, it may suffer from low performance when dealing with images of low visible quality, since it depends on the potential of local image detectors/descriptors to express the visual content that is matched. Evidence supporting this claim can be found in the failure rate of THEMIS-VIS MC11-E data set, since THEMIS is the only high-resolution orbital camera that incorporates an image rating, with a range of values from one (worse quality) up to seven (best quality). By comparing the failure rate of THEMIS-VIS MC11-E images with rating ≤ 4 with the corresponding value of images with rating ≥ 5 , it is found that the low-quality image failure rate was 32.19% while

³<http://www.i-mars.eu/web-gis>




Fig. 7. Detail of the CTX mosaic of MC11-E half-quadrangle, which was produced using as input images the CTX images coregistered to the HRSC MC11-E mosaic [57]. Note that mosaicing included only photometric correction, while the geometry was directly acquired by the coregistered products. (a) Detail of the CTX MC11-E mosaic, in 100-m/pixel resolution. (b) Same detail in the HRSC MC11-E mosaic [57]. (c) Close-up of the CTX MC11-E mosaic, in 6-m/pixel resolution. (d) Same area in the HRSC MC11-E mosaic, in 6-m/pixel resolution. In (d), visible quality is low because the nominal resolution of the HRSC MC11-E mosaic (12.5 m/pixel) is coarser than the one used for the display.

the high-quality image failure rate was only 10.68%. This divergence implies a strong connection between the image quality and the ring matching performance, even though the introduced algorithm still achieves a subpixel coregistration for two-thirds of the low-quality THEMIS-VIS images.

Another limitation of the algorithm is related to the resolution difference between the reference and the target image. The reported results imply that the achieved coregistration accuracy is subpixel in relation to the reference image pixel size; hence, in cases of large resolution difference between the reference and the target image, this accuracy may correspond to residuals of several (target image) pixels. For example, if a HiRISE image (of resolution 0.25 m/pixel) is coregistered to HRSC (of resolution 12.5 m/pixel) and the achieved accuracy is 6 m, then the coregistration residuals would be 24 HiRISE pixels. It is clear that such an accuracy is not adequate and the coregistration needs to be repeated using a baseline with a resolution nearer to the target image, such as CTX or color and stereo surface imaging system [61].

VII. CONCLUSION

In this paper, we have presented a novel fully automatic planetary image coregistration technique, called ring matching, that has been shown to be fast, accurate, and robust to parameter tuning. We have shown how ring matching can be

used to systematically coregister large volumes of input data, making their geometry consistent, thus allowing their fusion as well as their straightforward comparison. Experimental validation is given by testing the method in a challenging experimental setup that included thousands of high-resolution planetary images, acquired from six orbital instruments.

We are currently employing this method to coregister as much of Martian high-resolution imagery as possible, using HRSC as a baseline, the results being released through the i-Mars WebGIS. Moreover, in the near future, a similar approach is planned for the Moon, in order to geometrically align the vast volume of data that LROC has acquired since 2009, and for Mercury, in order to geometrically align MDIS high-resolution imagery.

On the other hand, we have already started to examine the use of ring matching to coregister further types of data (apart from visible spectrum images). As a matter of fact, in this paper, experimental results regarding the coregistration of infrared images (more specifically, THEMIS-IR) were presented, with promising results. In the future, we plan to study the coregistration of infrared to visible planetary data, in order to establish a similar processing paradigm. Finally, we are expanding the systematic coregistration to imaging spectrometer data, mainly focusing on Compact Reconnaissance Imaging Spectrometer for Mars data acquired of the Martian surface.

ACKNOWLEDGMENT

The authors would like to thank Dr. G. Michael of Freie Universität Berlin, Berlin, Germany, for the mosaicing.

REFERENCES

- [1] B. Zhao *et al.*, "Overall scheme and on-orbit images of Chang'E-2 lunar satellite CCD stereo camera," *Sci. China Technol. Sci.*, vol. 54, pp. 2237–2242, Sep. 2011.
- [2] M. S. Robinson *et al.*, "Lunar Reconnaissance Orbiter Camera (LROC) instrument overview," *Space Sci. Rev.*, vol. 150, nos. 1–4, pp. 81–124, 2010.
- [3] P. Sidiropoulos and J.-P. Muller, "On the status of orbital high-resolution repeat imaging of Mars for the observation of dynamic surface processes," *Planetary Space Sci.*, vol. 117, pp. 207–222, Nov. 2015.
- [4] H. A. Weaver *et al.*, "The small satellites of Pluto as observed by New Horizons," *Science*, vol. 351, no. 6279, p. aae0030, 2016. [Online]. Available: <https://www.semanticscholar.org/paper/The-small-satellites-of-Pluto-as-observed-by-New-H-Weaver-Buie/a4c1f8bd23a1076c74a32b56ae70ecb9f155c7c9>
- [5] V. R. Baker *et al.*, "Fluvial geomorphology on Earth-like planetary surfaces: A review," *Geomorphology*, vol. 245, pp. 149–182, Sep. 2015.
- [6] T. J. Parker, D. S. Gorsline, R. S. Saunders, D. C. Pieri, and D. M. Schneeberger, "Coastal geomorphology of the Martian northern plains," *J. Geophys. Res., Planets*, vol. 98, no. E6, pp. 11061–11078, 1993.
- [7] A. S. McEwen *et al.*, "Seasonal flows on warm Martian slopes," *Science*, vol. 333, no. 6043, pp. 740–743, 2011.
- [8] S. Byrne *et al.*, "Distribution of mid-latitude ground ice on Mars from new impact craters," *Science*, vol. 325, no. 5948, pp. 1674–1676, 2009.
- [9] J. W. Barnes *et al.*, "Cassini/VIMS observes rough surfaces on Titan's Punga Mare in specular reflection," *Planetary Sci.*, vol. 3, p. 3, Dec. 2014.
- [10] C. H. Acton, Jr., "Ancillary data services of NASA's navigation and ancillary information facility," *Planetary Space Sci.*, vol. 44, no. 1, pp. 65–70, 1996.
- [11] C. S. Edwards, K. J. Nowicki, P. R. Christensen, J. Hill, N. Gorelick, and K. Murray, "Mosaicking of global planetary image datasets: 1. Techniques and data processing for Thermal Emission Imaging System (THEMIS) multi-spectral data," *J. Geophys. Res., Planets*, vol. 116, no. E10, pp. 1–21, 2011.
- [12] M. C. Malin *et al.*, "An overview of the 1985–2006 Mars Orbiter Camera science investigation," *Mars*, vol. 5, pp. 1–60, 2010. [Online]. Available: <http://www.marsjournal.org/contents.shtml>
- [13] N. L. Chabot *et al.*, "Mapping mercury: Global imaging strategy and products from the MESSENGER mission," in *Proc. 47th Lunar Planetary Sci. Conf.*, 2016, p. 1256.
- [14] G. A. Soffen and C. W. Snyder, "The first Viking mission to Mars," *Science*, vol. 193, no. 4255, pp. 759–766, 1976.
- [15] P. R. Christensen *et al.*, "The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey mission," *Space Sci. Rev.*, vol. 110, no. 1, pp. 85–130, 2004.
- [16] R. Jaumann *et al.*, "The high-resolution stereo camera (HRSC) experiment on Mars Express: Instrument aspects and experiment conduct from interplanetary cruise through the nominal mission," *Planetary Space Sci.*, vol. 55, nos. 7–8, pp. 928–952, 2007.
- [17] J. F. Bell, III, *et al.*, "Calibration and performance of the Mars Reconnaissance Orbiter Context Camera (CTX)," *Mars*, vol. 8, pp. 1–14, 2013. [Online]. Available: <http://www.marsjournal.org/contents.shtml>
- [18] A. S. McEwen *et al.*, "Mars reconnaissance orbiter's High Resolution Imaging Science Experiment (HiRISE)," *J. Geophys. Res., Planets*, vol. 112, no. E5, 2007. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1029/2005JE002605/full>
- [19] A. A. Sima *et al.*, "Compact Hyperspectral Imaging System (COSI) for small Remotely Piloted Aircraft Systems (RPAS)—System overview and first performance evaluation results," in *Proc. Int. Arch. Photogramm., Remote Sens. Spatial Inf. Sci.*, vol. XLI-B1, 2016, pp. 1157–1164. [Online]. Available: <https://www.int-arch-photogramm-remote-sens-spatial-inf-sci.net/XLI-B1/1157/2016/isprs-archives-XLI-B1-1157-2016.pdf>
- [20] P. Sidiropoulos and J.-P. Muller, "Matching of large images through coupled decomposition," *IEEE Trans. Image Process.*, vol. 24, no. 7, pp. 2124–2139, Jul. 2015.
- [21] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [22] D. Robinson and P. Milanfar, "Fundamental performance limits in image registration," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1185–1199, Sep. 2004.
- [23] M. L. Uss, B. Vozel, V. A. Dushepa, V. A. Komjak, and K. Chehdi, "A precise lower bound on image subpixel registration accuracy," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 6, pp. 3333–3345, Jun. 2014.
- [24] J. Senthilnath and R. Prasad, "A new SIFT matching criteria in a genetic algorithm framework for registering multisensory satellite imagery," in *Proc. Indian Conf. Comput. Vis. Graph. Image Process.*, 2014, Art. no. 21.
- [25] B. Zitová and J. Flusser, "Image registration methods: A survey," *Image Vis. Comput.*, vol. 21, pp. 977–1000, Oct. 2003.
- [26] J. Le Moigne, N. S. Netanyahu, and R. D. Eastman, Eds., *Image Registration for Remote Sensing*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [27] L. Barazzetti, M. Gianinetti, and M. Scaioni, "Automatic registration of multi-source medium resolution satellite data," in *Proc. Int. Arch. Photogramm., Remote Sens. Spatial Inf. Sci.*, vol. XL-7, 2014, pp. 23–28. [Online]. Available: <https://www.int-arch-photogramm-remote-sens-spatial-inf-sci.net/XL-7/23/2014/isprsarchives-XL-7-23-2014.pdf>
- [28] J. M. Murphy, J. Le Moigne, and D. J. Harding, "Automatic image registration of multimodal remotely sensed data with global shearlet features," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 3, pp. 1685–1704, Mar. 2016.
- [29] L. Cheng, L. Ma, K. Yang, Y. Liu, and M. Li, "Registration of Mars remote sensing images under the crater constraint," *Planetary Space Sci.*, vol. 85, pp. 13–23, Sep. 2013.
- [30] H. Goncalves, L. Corte-Real, and J. A. Goncalves, "Automatic image registration through image segmentation and SIFT," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 7, pp. 2589–2600, Jul. 2011.
- [31] A. W. Gruen, "Adaptive least squares correlation: A powerful image matching technique," *South African J. Photogramm., Remote Sens., Cartogr.*, vol. 14, no. 3, pp. 175–187, 1985.
- [32] L. Barazzetti, M. Scaioni, and M. Gianinetti, "Automatic co-registration of satellite time series via least squares adjustment," *Eur. J. Remote Sens.*, vol. 47, no. 1, pp. 55–74, 2014.
- [33] W. Zhao, Z. Tian, L. Yang, W. Yan, and J. Wen, "Image registration using a kernel partial least squares based mismatches removal method," *AEU-Int. J. Electron. Commun.*, vol. 70, no. 4, pp. 427–435, 2016.
- [34] S. Paul and U. C. Pati, "Remote sensing optical image registration using modified uniform robust SIFT," *IEEE Trans. Geosci. Remote Sens. Lett.*, vol. 13, no. 9, pp. 1300–1304, Sep. 2016.
- [35] G. Troglia, J. Le Moigne, J. A. Benediktsson, G. Moser, and S. B. Serpico, "Automatic extraction of ellipsoidal features for planetary image registration," *IEEE Geosci. Remote Sens. Lett.*, vol. 9, no. 1, pp. 95–99, Jan. 2012.
- [36] M. Vakalopoulou, K. Karatzalos, N. Komodakis, and N. Paragios, "Simultaneous registration and change detection in multitemporal, very high resolution remote sensing data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2015, pp. 61–69.
- [37] B. Kupfer, N. S. Netanyahu, and I. Shimshoni, "A SIFT-based mode-seeking procedure for efficient, accurate registration of remotely sensed images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2013, pp. 4142–4145.
- [38] K. Mikołajczyk and C. Schmid, "A performance evaluation of local descriptors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 10, pp. 1615–1630, Oct. 2005.
- [39] A. Sedaghat and H. Ebadi, "Remote sensing image matching based on adaptive binning SIFT descriptor," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 10, pp. 5283–5293, Oct. 2015.
- [40] A. Sedaghat and H. Ebadi, "Very high resolution image matching based on local features and k -means clustering," *Photogramm. Rec.*, vol. 30, no. 150, pp. 166–186, 2015.
- [41] M. Scaioni, L. Barazzetti, and M. Gianinetti, "The concept of reliability applied to co-registration of multiple satellite time-series of images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2015, pp. 4420–4423.
- [42] J. R. Kim, J.-P. Muller, J. G. Morley, and K. L. Mitchell, "Automated registration of MDIM with MOLA tracks," in *Proc. 32nd Annu. Lunar Planetary Sci. Conf. (LPSC)*, 2001.
- [43] M. I. Patel and V. K. Thakar, "Speed improvement in image registration using maximum likelihood based mutual information," in *Proc. IEEE Int. Conf. Adv. Comput. Commun. Syst.*, Jan. 2015, pp. 1–3.

- [44] G. Wolberg and S. Zokai, "Image registration for perspective deformation recovery," *Proc. SPIE*, vol. 4050, 2000. [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/4050/1/Image-registration-for-perspective-deformation-recovery/10.1117/12.395570.short>
- [45] I. Zavorin and J. L. Moigne, "Use of multiresolution wavelet feature pyramids for automatic registration of multisensor imagery," *IEEE Trans. Image Process.*, vol. 14, no. 6, pp. 770–782, Jun. 2005.
- [46] A. A. Zakharov, A. Y. Tuzhilkin, and A. L. Zhiznyakov, "Finding correspondences between images using descriptors and graphs," *Procedia Eng.*, vol. 129, pp. 391–396, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1877705815040151>
- [47] G. Brigot, E. Colin-Koeniguer, A. Plyer, and F. Janez, "Adaptation and evaluation of an optical flow method applied to coregistration of forest remote sensing images," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 9, no. 7, pp. 2923–2939, Jul. 2016.
- [48] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," in *Proc. Eur. Conf. Comput. Vis.*, 2006, pp. 404–417.
- [49] K. Li and S. Zhou, "A fast SIFT feature matching algorithm for image registration," in *Proc. Int. Conf. Multimedia Signal Process. (CMSP)*, 2011, pp. 89–93.
- [50] C. Silpa-Anan and R. Hartley, "Optimised KD-trees for fast image descriptor matching," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2008, pp. 1–8.
- [51] J. Zaletelj, "Reliable subpixel ground control point estimation algorithm using vector roads," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 9, pp. 1325–1329, Sep. 2016.
- [52] Q. Li, G. Wang, J. Liu, and S. Chen, "Robust scale-invariant feature matching for remote sensing image registration," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 2, pp. 287–291, Apr. 2009.
- [53] J. M. Murphy, O. N. Leija, and J. Le Moigne, "Agile multi-scale decompositions for automatic image registration," *Proc. SPIE*, vol. 9840, p. 984011, May 2016. [Online]. Available: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/9840/1/Agile-multi-scale-decompositions-for-automatic-image-registration/10.1117/12.2222182.short>
- [54] M. I. Patel and V. K. Thakar, "Application of radon transform for fast image registration," in *Proc. IEEE Int. Conf. Adv. Comput. Commun. Syst.*, Jan. 2015, pp. 1–4.
- [55] R. Gupta and R. I. Hartley, "Linear pushbroom cameras," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 9, pp. 963–975, Sep. 1997.
- [56] R. Hartley and T. Saxena, "The cubic rational polynomial camera model," in *Proc. ARPA IU Workshop*, 1997, pp. 649–653.
- [57] K. Gwinner *et al.*, "The High Resolution Stereo Camera (HRSC) of Mars Express and its approach to science analysis and mapping for Mars and its satellites," *Planetary Space Sci.*, vol. 126, pp. 93–138, Jul. 2016.
- [58] J. Bostelmann and C. Heipke, "Results from multi-strip bundle adjustment with HRSC imagery," in *Proc. ISPRS Working Group IV/8 Planetary Mapping Spatial Databases Meet.*, 2015.
- [59] P. C. McGuire *et al.*, "True- and false-color HRSC+OMEGA image mosaics of Mars," in *Proc. 47th Lunar Planetary Sci. Conf. (LPSC)*, 2016.
- [60] G. G. Michael *et al.*, "Systematic processing of Mars Express HRSC panchromatic and colour image mosaics: Image equalisation using an external brightness reference," *Planetary Space Sci.*, vol. 121, pp. 18–26, Feb. 2016.
- [61] N. Thomas *et al.*, "The colour and stereo surface imaging system for ESA's trace gas orbiter," in *Proc. 47th Lunar Planetary Sci. Conf. (LPSC)*, 2016.



Panagiotis Sidiropoulos (M'12) received the B.Sc. degree in electrical and computer engineering and the M.Sc. degree in informatics from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2003 and 2007, respectively, and the Ph.D. degree in machine learning from the Centre for Vision, Speech, and Signal Processing, University of Surrey, Surrey, U.K., in 2012.

He is currently a Post-Doctoral Research Associate with the Mullard Space Science Laboratory, University College London, London, U.K. His research interests include the use of remote sensing, machine learning, and multimedia processing techniques to planetary science.



Jan-Peter Muller received the B.Sc. degree (Hons.) in physics from the University of Sheffield, Sheffield, U.K., in 1976, the M.Sc. and D.I.C. degrees in atmospheric physics and dynamics from Imperial College London, London, U.K., in 1977, and the Ph.D. degree in planetary meteorology from University College London (UCL), London, in 1982.

He is currently a Professor of image understanding and remote sensing with UCL, where he has been on the faculty since 1984. He has been a member of the Moderate Resolution Imaging Spectrometer and Multi-Angle Imaging SpectroRadiometer Science Teams since 1990. He is a Deputy Coordinator of the EU-FP7 QA4ECV project, generating a unique 35-year record of earth's land surface albedo from European and U.S. sensors, and a Coordinator of the EU-FP7 iMars project, generating a 40-year long record of co-registered, orthorectified orbital images of Mars to map change detection using data mining and crowd sourcing methods. His research interests include the practical applications of machine vision to remote sensing of the earth and planets, with a special focus on the operational applications of stereo photogrammetry and other 3-D imaging technologies for applications in climate modeling, weather forecasting, and planetary exploration.

Dr. Muller was the Chair of the U.K. Aurora Advisory Committee, and the Chair of the Committee on Earth Observation Satellites WG on Cal/Val "Terrain Mapping Sub-Group," as well as the International Society for Photogrammetry and Remote Sensing Working Group on "Global DEM interoperability."