Extraction of Common Regions of Interest from Orbital and Rover (Ground) acquired Imagery

Evangelos Boukas*, Antonios Gasteratos*, Gianfranco Visentin**

*School of Engineering, Democritus University of Thrace, Vas. Sophias 12, GR-67100 Xanthi, Greece. e-mail: evanbouk@pme.duth.gr, agaster@pme.duth.gr

**Automation and Robotics Section, ESA
e-mail: Gianfranco.Visentin@esa.int

Abstract

Accurate global localization of space exploratory rovers is of high importance to the future extraterrestrial missions. In particular, the Mars Sample Return mission (MSR) needs precise localization techniques for the Sample Fetching Rover (SFR) to collect a previously drilled cache of soil. Moreover, exact positioning is sine qua non for a successful rendezvous of the SFR with a Mars Ascent Vehicle (MAV) that will launch the load into Martian orbit. The utilization of prior information to refine the localization has been the main approach in mobile robotics during the past decades. In extraterrestrial scenery, orbital images constitute an excellent source of such prior information. With long term aim to provide accurate global rover localization, as a first step, the proposed paper assesses the ability to automatically detect common regions of interest from both rover (ground) and orbital images. The approach comprises the separate extraction of salient regions, employing distinct techniques for the orbital and ground images.

1 Introduction

1.1 Motivation

Owing to the space exploratory roves being assigned with ever more complex tasks, the increase in their autonomy is considered to be the next big step towards advanced and cooperative extraterrestrial missions. The autonomy of space exploratory rovers is highly connected to their ability to effectively localize themselves. In particular, the Mars Sojourner rover (Pathfinder) beared autonomy of 2-3 m per solar day (sol) employing only wheel encoders and sun sensors. The MER rovers (Spirit and Opportunity) have increased their autonomy by incorporating visual odometry along with Inertial Measurement Unit (IMU) measurements. The most recent rover on Mars, namely the MSL rover Curiosity, has travelled more than $5km$ (500+sols), however the visual odometry has been used mostly for slip check, rather than for autonomous driving. On the other hand, ExoMars rover will utilize visual odometry for autonomous navigation as a nominal mode, thanks to a dedicated coprocessor [1].

An example of such an advanced mission is the forthcoming Mars Sample Return (MSR) campaign, which is a joint effort of the European Space Agency (ESA) and the National Aeronautics and Space Administration (NASA). The MSR has received international attention and its main purpose is to return $500g$ of Mars soil sample to Earth. The Martian soil will be, thereafter, analysed in Earth laboratories that include state of the art analysis instrumentation. The hardware and software of the rover that is required to fetch the previously collected cache of Martian soil has been targeted by ESA as another contribution to the MSR campaign, namely the Sample Fetching Rover (SFR) [2].

1.2 Related Work

Bearing in mind that an one-way signal to Mars requires approximately up to 20 *min*, it is apparent that the non autonomous navigation of rovers poses a noticeable overhead to the exploration of the red planet. Therefore, recent and future space exploratory rovers are designed for long traverses [3, 4]. This is the reason why the rovers are equipped with several cameras. An example of camera setup design for such a specific application is reported in [5], where a camera system for localization and mapping of space exploratory rovers is proposed. The algorithms required for the long range autonomous navigation are 3D reconstruction, mapping [6], localization and path planning.

Moreover, notwithstanding the advanced and noteworthy mechanical and electrical design of space exploratory rovers, their capabilities in terms of computational and power needs are modest, especially compared to contemporary robots. The algorithms for the localization, as well as for any other task, should be as lightweight as possible. Toward this end, the authors in [7] have implemented a Visual Odometry (VO) algorithm that is capable of achieving state of the art results, while being of low computational cost. In terms of power consumption, the authors in [8] have implemented the localization and

mapping algorithms on FPGA devices for ESA's future space exploratory rovers.

In the context of orbital based localization, the authors in [9] employ descent imagery to create a Digital Elevation Map (DEM) of the area around the landing site. Nextly, they solve an image network utilizing bundle adjustment based on manual selected tie points. The authors in [10] propose a localization framework based on skyline matching. The skyline from the rovers imagery is firstly extracted and then matched to possible skylines that have been extracted from a DEM. The matching reveals the position of the rover. Li et al. [11] noted that with the aim to produce high quality localization estimates, orbital and rover image networks should be connected through tie points. The work presented in [12] employed a Bayesian recursive algorithm, namely a modified particle filter able to retrieve the location of the rover on a global DEM. The authors in [13] proposed the employment of a LIght Detection And Ranging (LIDAR) device in a system that is able to globally localize a rover laying within the imaging region of a georeferenced DEM. Nevertheless, the inclusion of a such devices on Space Missions is not feasible, at least based on current Space Agencies' predictions. For a more detailed review of the state of the art on the localization of planetary exploration rovers with orbital imaging can be found in [14].

Considering the aforementioned related works it is obvious that one of the most important features of a global rover localization system is the acquisition of tie points among the orbital and rover imagery. One of the most complete works in the detection of salient regions for scientific analysis is the Autonomous Exploration for Gathering Increased Science (AEGIS) system [15]. The AEGIS focuses on the automated targeting of regions of interest on rover imagery. It includes an algorithm, namely the *ROCKSTER*, for the identification of such regions. This algorithm employs edge detection, flood fill and contour calculation to cope with edge fragmentations. However the method described in this paper targets the regions that are observable both on ground and aerial imagery.

1.3 ESA Network/Partnerning Initiative

As is the case with the aforementioned MSR mission, it is clear that one of the most important abilities of future rovers, such as the Sample SFR, is the accurate global localization. Towards this end, the ESA funded Network/Partnering Initiative entitled: "Methods to Refine the Self-Localization of Planetary Rovers Using Orbital Imaging" aims to create a localization framework that will employ orbital imagery. The rover is considered to be "lost in space", yet it is supposed to lay within an area portrayed in an orbital image. Thus, the proposed method firstly assesses the extraction of regions of interest on the orbital images. In particular, it employees the mixture of appropriate computer vision, morphology analysis and machine learning techniques to extract regions, which are likely to be detected from the rover as well. Then, taking advantage of the redundancy of the rover imagery, a probabilistic approach is employed to extract regions that are potentially detectable on the orbital images. From this point, landmarks are appropriate to be fed into a system able to solve the "lost in space" problem. As the first step to this system, this paper presents the methods employed in order to extract commonly observed regions of Interest (ROIs) from both orbital and rover imagery. The rest of the paper is organized as follows:

- The utilized dataset is presented in Sec. 2.
- The detection of Regions of Interest on orbital imagery is presented in Sec. 3
- The detection of Regions of Interest on rover imagery is presented in Sec. 4

2 Dataset

Since our system is ought to be utilized in extraterrestrial scenery, the data upon which the model is designed and tested are of most importance. The dataset being utilized in the context of this research is the one acquired along the framework of the SEEKER activity, which was carried out under the European Space Agency's StarTiger initiative [16]. This dataset was collected in the Chilean desert, Atacama after thorough analysis and has been proven to be the most Mars like site on earth. The dataset contains information acquired both by a rover and an orbital imagery analog.

2.1 Orbital Imagery (Analog)

The orbital images have been produced employing aerial imagery mounted on a Unmanned Aerial Vehicle (UAV). The UAV was utilized in order to produce orthoimages with spatial resolution greater than $0.1m$. However, both the orthoimages and DEMs are downsampled to most accurately reflect the real extraterestrial data that are available nowadays. The orthorectified and georeferenced orbital image that was considered in our setup and evaluation, along with the respective DEM is depicted in Fig. 1. The orbital images contain large rocks, both wide spread and gathered. The rocks that appear on the orbital images have dimensions at least greater the resolution of the imagery. This information is taken into consideration on the detection of rocks on the rover imagery. Another type of landmark that is clearly observable on the orbital images are the dry salt beds: *salars*. The salars are prominent, reflective and assemble the outcrops that are targeted and analyzed by Mars rovers. These landmarks are as well usable for a global localization system and, therefore, are included in our model.

(a) Orthorectified & Georeferenced Orbital Image.

(b) Digital Elevation Model (DEM).

Figure 1. Complete Orbital Image and Digital Elevation Model (DEM) of the SEEKER Atacama Dataset.

2.2 Rover Imagery

The rover employed in the dataset acquisition has been built over the robovolc, which is a six wheel mobile robot capable to reach velocity up to $1m/s$ with maximum operational time up to 7 hours. The SEEKER team has customized the robot in order to include, among others, the Bumblebee XB3 stereo camera , an Inertial Measurement Unit (IMU) and encoders for wheel odometry. The stereo cameras acquired images at a resolution of $512px \times 384px$. Some indicative instances of the data acquired by the rover imagery are depicted in Fig. 2. The total route of the rover is more than $5km$ long and can be segmented in three parts: The first segment is a $850m$ route characterized by direct sunlight that leads the path planning and localization to become quite challenging. The second part is a $1.2km$ traverse that is distinguished by the existence of large rocks, soft sand and noteworthy slopes. The final segment follows a $3km$ course with direct sunlight.
However, the advanced difficulties of the dataset should be noted, as well:

- The great instantaneous alteration in the roll axis, between consecutive frames, pose an extra burden towards accurate local motion estimation.
- The direct sunlight augments the difficulties of image processing algorithms, namely the Disparity estimation and the blob detection.

Moreover, the rover was equipped with a Differential Global Positioning System (DGPS). Unfortunately, due to some specific DGPS system failures the measurements were strangled and de-synchronized with rover image acquisition. However, via a filtering method and extensive processing we were able to result with a sparser than intended but, nevertheless, highly accurate ground truth (Figure 1).

3 Detection on Orbital Imagery

Concerning the detection on orbital images, the method employs the mixture of several techniques such as a Hessian analysis for blob identification and an entropy based salient detector. The final step comprises a classifier that distinguishes between the different types of salient regions (rocks, outcrops, etc.).

3.1 Multiscale Blob Detection

3.1.1 Two-Sigma filter

The orbital images contain speckle noise due to numerous factors, such as the mosaicing procedures. Therefore, the first step of the multiscale blob detection comprises the spatial filtering in order to remove any false pixel values that were introduced. The filter employed in the noise removal is the two-sigma algorithm [17] which is know to suppress speckles while retaining edges. The two-sigma filter is based on the assumption that the information follows a gaussian distribution with a mean value μ and variance σ^2 and therefore the 95.4% of the values lies within the range $[\mu - 2\sigma, \mu + 2\sigma]$. If the number of neighboor values lying within the range $[\mu - 2\sigma, \mu + 2\sigma]$ is greater than a variable number K, the center cell value is assigned to μ. In any other case the central case is assigned with the average of the neighboring cells.

3.1.2 Hessian analysis

The Hessian analysis has been used for the detection of prominent structures, including blob-like and ellipse-like shapes [18]. The computation of the Hessian matrix

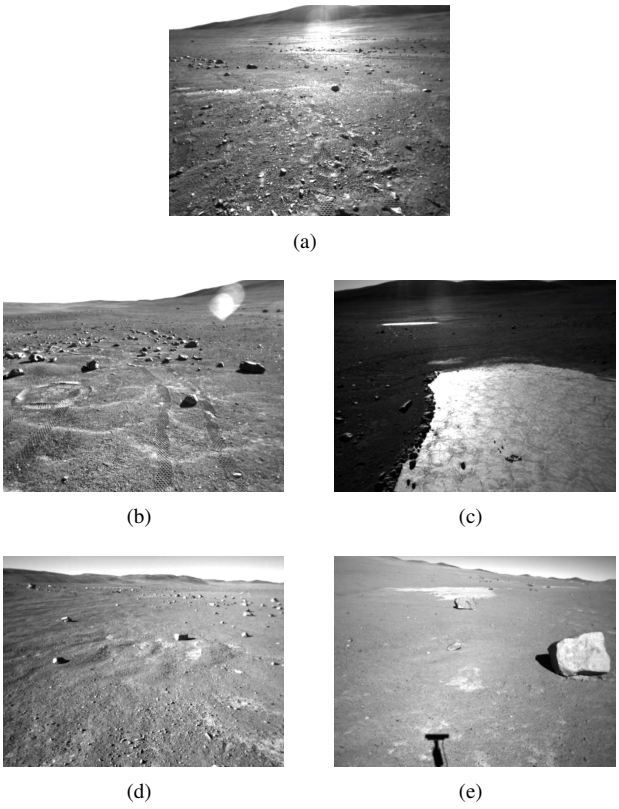

Figure 2. Indicative instances of the rover's route. Special characteristics of the dataset: (a) Direct sunlight; (b) Dense Rocks; (c) Salar regions; (d) Sparse rocks; and, (e) Large rocks.

is performed employing the derivatives of Gaussian. The protruding and well defined shapes in the image, usually, are expressed as extreme eigenvalues of the Hessian matrix [18]. Therefore, we perform an analysis on the distribution of the eigenvalues where only the image points that lay outside the 2σ around the mean of the distribution are considered as blob candidates.

The σ of the Gaussian can be considered as the main parameter in the context of the scale space theory [19]. The higher values of σ lead to the detection of larger regions. However, the scenes taken into consideration in this work contain various types of blobs and, therefore, the detection on multiple scales is required. At this point we should note that the method introduces some false detections which are, nevertheless, handled by means of classification as described in section 3.2.

3.1.3 Local Entropy

Local entropy, which is based on Shannon entropy [20], is a metric that represents quantitatively the randomness of the image intensity levels in an image (Eq. 1). The calculation employs the probabilities of each intensity to appear in an image neighborhood, similarly to the whole image entropy. Therefore the probabilities inside the neighborhood are calculated by firstly computing the histogram of intensities and then by normalizing the histogram to sum up to 1. The local entropy filter LE is presented in the following equation:

$$LE = -\sum_{i=0}^{255} p(i) \cdot \ log_2(p(i)), \tag{1}$$

where $i \in [0, \ldots, 255]$ is the intensity and $p(i)$ is the probability of each intensity.

The blobs extracted with the aforementioned methods are able to detect different and complementary regions. Therefore, a logical OR operation is performed on the detected regions to get the final blob image. The resulting detection includes areas that contain holes and, therefore, we apply a binary image filing operation [21]. The next step comprises a crucial parameter for the detection of commonly observable regions from both orbital and ground imagery, i.e. the size of the blobs to be considered. Depending on the orbital image resolution one should select the appropriate size. Since the best resolution of state of the art orbital imagery on Mars (HiRISE) is able to capture images at a resolution of $0.25m/px$, we define the minimum size of blobs to be detected $0.5m \times 0.5m$. Then the labeling of the the blobs is performed employing a connected component algorithm [22]. An example of the ROI detection is presented in figure 3. While the detection is able to extract most of the regions of interest, some false detections (mere sand) occur. The next section discusses the methods employed in the system that enable, apart from the distinguishing between the types of ROI, i.e. rocks and outcrops, the discarding of these false detections.

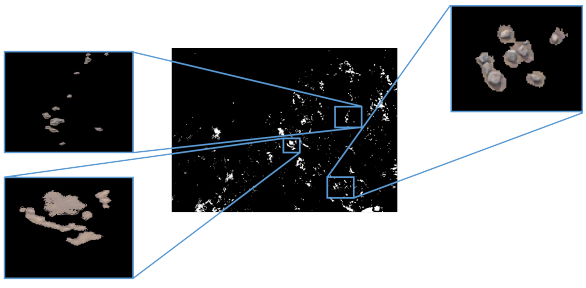

Figure 3. An indicative example of the ROI detection on a sample scene. It can be seen that the detection method is able to detect rocks as well as salars (outcrop analogous).

3.2 Classification for Sample Selection and Annotation

This section describes the classification of previously detected ROIs on three classes: Rocks, Outcrops analogs (mostly salars) and Ground (mostly sand). Through the classification procedure the annotation of the blobs is performed and any falsely detected regions, i.e. ground, are removed.

3.2.1 Feature Space

Based on the detected blobs, the feature space for the classification can be defined. The features employed in our method are related to color, contrast and texture. Firstly, the actual intensity values of the blobs are considered. The colorspaces that are taken into consideration are: Grayscale, RGB and HSV. Additionally, the local entropy has been employed to measure the randomness of the values in the ROI (sec. 3.1.3). Moreover, in order to infuse the system with the ability to identify the variations in texture, the Gabor features have been investigated. The $2D$ Gabor filter is the modulation of a $2D$ Gaussian distribution with a sinusoidal with a specific orientation. The Gabor features are computed by the convolution of the input image with a number of Gabor filters with various: standard deviations of the $2D$ Gaussian (σ), orientations of the sinusoidal and phase offsets (θ,ϕ).

3.2.2 Feature vector generation

A feature vector of a ROI is defined as a N-dimensional vector containing information about the blob. The structure containing the feature vectors of all samples, i.e. ROIs, when concatenated in a matrix, is usually referred to as *data matrix* and is the basis upon which the classification procedure is performed. The size of the data matrix is $M \times N$, with M being the number of ROIs. Due to different ROI sizes, in order to keep a fixed size of the data matrix, all blobs are resized to a specific size, i.e $20px \times 20px$ independently the blob being larger or smaller than this value.

The next step towards the feature generation is the dimensionality reduction (DR) one, i.e. the projection of the N_{raw} dimensional sample to an N dimensional space, where $N < N_{raw}$, with as little loss as possible in information. The need for dimensionality reduction is twofold. On the first hand, there is the problem of computational cost. One of the most commonly used methods for dimensionality reduction is the Principal Component Analysis (PCA) or KarhunenLove transformation which is an unsupervised method to produce a linear mapping of a high dimensional feature vector (and hence the data matrix) to a low dimension one [23]. An example of the dimensionality reduction in the current system is the reduction of 32800 dimensions (for all the types of features) to 119 with a projection to the principal components that preserve 99% percent of the information.

3.2.3 k-NN classifier

The k nearest neighbors (k-NN) algorithm is a NN generalization where, instead of one, k neighbors are employed [24]. The k-NN classifier has the following form: Given a set of N-dimensional labelled data and an unlabelled vector X of a new sample, the algorithm finds the k nearest neighbors based on a metric. The label of the majority of the nearest neighbors is assigned to the new sample. The only variable in the classifier is the k, which is usually an odd number (for voting reasons) and may not be a multiple of the number of classes. Considering the fact that this system employs three classes the candidate values for k are: {1,5,7,11,13}. The selection of feature types and parameter k was performed utilizing a *leave one out* validation approach. The selected feature vector is the one containing the mixture of {HSV,GABOR,ENTROPY} and a $k = 1$. Although the over all classification rate is not excellent, in the context of this work the classification of Rocks and Outcrops are adequate. The fact that really low false positive exist, makes the classification sufficient.

Based on the aforementioned evaluation of the classification the system is able to detect and label regions of interest from aerial imagery efficiently. Moreover the presented method is able to perform sample selection, i.e to discard any falsely detected ROIs, in order to retain the output free of noise. The detection on orbital imagery was performed over most of the parts of the Atacama dataset. Some indicative examples of detection and annotation are presented in figure 4.

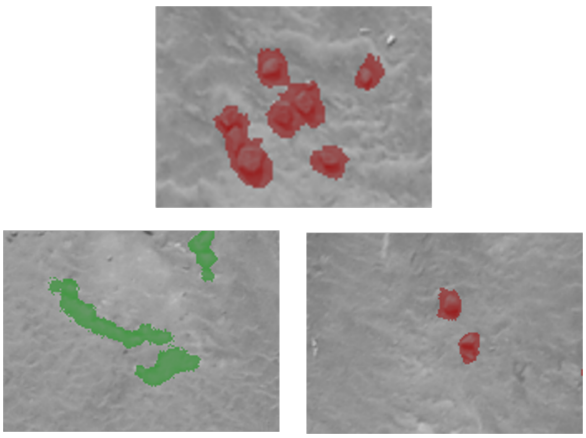

Figure 4. Indicative examples of the ROI detection and annotation on the Atacama dataset. The red color indicates rocks and the green color indicates outcrops

4 Detection on Rover Imagery

In the context of rover localization the most important data are those stemming from its own sensory, such as stereo cameras, inertial measurement unit (IMU), sun sensor, etc. In this work the detected ROIs from rover images, in contrast to previous works, have the main constraint and requirement to be observable by orbital imagery as well.

4.1 Single Instance Detection

4.1.1 Hessian analysis & Entropy

The first step on the detection on the rover's images is similar to the ones presented in sections 3.1.2 and 3.1.3 for the orbital imaging with small alterations due to the change of perspective and level of detail. The hessian analysis which is based on the extreme values of the eigenvalues of the image. The main difference between this implementation and the one presented in the previous section is the range of the variable σ, which is the main parameter in the scale space theory and, contains hereby higher values. The entropy has also been employed in this part of the system with small modifications, i.e the size of the neighborhood has been extended in order to detect larger blobs.

4.1.2 Saliency

In theory an intelligent operator driving a vehicle (even on the ground of Mars) would use visual landmarks to specify its relative position and route. This is the inspiration that led to the enhancement of the system with the ability to detect visually salient regions. The saliency map models are focused on replicating the human attentional behavior on a scene, i.e. the region that a human would fixate in a specific scene. There are several approaches that have been proposed, the most important of which are [25]. The method employed on this work is the "Graph-Based visual saliency"[26].

4.1.3 3D Reconstruction for Rock Extraction

The 3D reconstruction of a scene in front of the rover utilizing a stereo camera is the process of assigning each pixel in the stereo image to a 3D point in the world. The algorithm employs the computation of a disparity map, which represents the difference of the viewpoint on the left and right image of the stereo rig and is presented in [27]. The simplified equation of the 3D reconstruction (neglecting lens distortions) can be expressed as follows:

$$[x, y, z] = \left[\frac{x_c \cdot z}{f}, \frac{y_c \cdot z}{f}, \frac{b \cdot f}{disp(x_c, y_c)}\right], \qquad (2)$$

where $[x_c, y_c]$ are the coordinates of a point in the image plane and $[x, y, z]$ are the 3D coordinates of the same point after the reconstruction. In order to extract the rocks that protrude from the ground the V-disparity has been utilized. The V-disparity is a line-wise histogram of the disparity of a stereo image [28]. The V-disparity of a stereo pair that contains mostly ground points, has the property to present the pixels of the image that depict the ground, on a line. Localizing the points in the V-disparity that are "above" the slope, one can detect the pixels in the image that correspond to rocks. An example of the detection of rocks utilising the V-disparity is shown in figure 5.

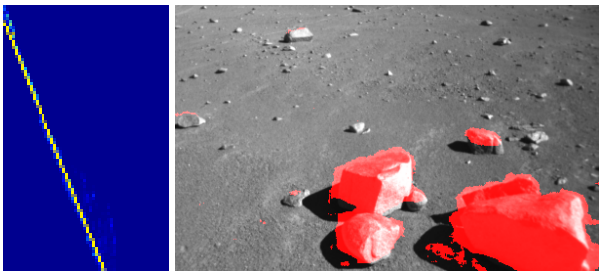

(a) The V-disparity. (b) The detected rocks over-plotted on the original image.

Figure 5. The detection of ROIs in a rover scene employing Entropy and Hessian analysis.

4.2 Classification

Following the detection of the regions of interest on a still rover image the next step comprises their classification in one of the three types: rock, outcrop, ground. The feature space of the the blobs cannot be the same as the one presented in section 3.2.1 due to rover cameras. The possible feature vectors are formed utilizing the following types of information: Gray, Gabor, Entropy. Principal Component Analysis is performed on the feature vector for dimensionality reduction and the resulting dimensions of the specific feature vector are 98. The classifier utilized in the system is the k-NN with the variable k set to 5. The evaluation was performed utilizing a *leave one out* cross correlation approach. The final feature vector utilized is the: {Gray, Gabor, Entropy}

4.3 Sequential Detection on Rover's Route

Although the detection and annotation of the ROIs has been presented adequately in this paper, there is still a very important factor that has not been discussed yet. That is the fact that the rover is not still, but is continuously moving. Therefore, the detection on the rover should not be performed on a single instance but rather on a sequence of images. The basis of the detection on the sequence of images is the motion estimation (sec. 4.3.1).

4.3.1 Visual Odometry

The calculation of a rover's route based on information stemming from its cameras is referred to as visual odometry. The authors in [7] proposed a computational efficient, yet effective algorithm for the computation of a rovers motion estimation, based on a novel outlier detection filter. This method has been employed as well in this system with the addition of an algorithm for the optimal fitting of consecutive point clouds, namely the Iterative Closest Points (ICP) [29]. Assuming a point cloud ${}^{t}\mathrm{P}$ that was created through 3D reconstruction, as explained in section 4.1.3, the same set of point cloud is observable on the time $t+1$, i.e. ${}^{t+1}\mathrm{P}$. Given the transformation between the two consecutive frames the point cloud ${}^{t+1}\mathrm{P}$ can be projected relatively to the robot's initial frame. However, due to erroneous calculation of the transformation, the point clouds are never perfectly aligned and, thus, the rover estimation may be subject to drifts. The ICP algorithm can provide a near-optimal solution to this problem by calculating the transformation required to accurately match the two misaligned point clouds.

In the context of the global rover localization the detection of ground stemming ROIs on an fixed frame, in the same way the orbital images are referenced to a fixed geographical point. By incrementally computing incremental motion estimations the system can acquire the rover's pose (position and orientation) relatively to its initial frame.

5 Conclusions

As space exploratory missions advance the need of global rover localization arises. This need is inextricably connected with the detection of commonly observed regions of interest on both orbital and ground imagery. The paper in hand presented the design, implementation and validation of a system that serves the aforementioned necessity. The system takes into consideration appearance features as well as geometrical characteristics of the rover's surrounding. The extraction of regions on the orbital images is performed based on a mixture of entropy and hessian analysis. The detected objects are assigned with a position relative to the rover's initial frame, towards the global localization end. Our analysis indicates that the classification on both types of images is performing adequately while, on the same time, it ensures that no false detection is inserted into our model. The proposed methodology enables the detection of commonly observable regions from rover and orbital imagery.

5.1 Acknowledgment

The proposed work is funded by the European Space Agency via the Network/Partnering Initiative under the contract No. 4000109064/13/NL/PA.

References

[1] P. Baglioni and L. Joudrier. Exomars rover mission overview. In *Robotics and Automation Workshop: on Planetary Rovers (ICRA Workshop), 2013 IEEE International Conference on*. IEEE, 2013.

[2] Andrea Merlo, Jonan Larranaga, and Peter Falkner. Sample fetching rover (sfr) for msr. In *Proceedings of the 12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2013)*, 2013.

[3] M Van Winnendael, P Baglioni, and J Vago. Development of the esa exomars rover. In *Proceedings of the 8th International Syymposiumon Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 2005)*, pages 5–8, 2005.

[4] John P Grotzinger, Joy Crisp, Ashwin R Vasavada, Robert C Anderson, Charles J Baker, Robert Barry, David F Blake, Pamela Conrad, Kenneth S Edgett, Bobak Ferdowski, et al. Mars science laboratory mission and science investigation. *Space science reviews*, 170(1-4):5–56, 2012.

[5] Ioannis Kostavelis, Evangelos Boukas, Lazaros Nalpantidis, Antonios Gasteratos, and Marcos Aviles Rodrigalvarez. Spartan system: towards a low-cost and high-performance vision architecture for space exploratory rovers. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1994–2001. IEEE, 2011.

[6] L. Nalpantidis, G.C. Sirakoulis, and A. Gasteratos. Non-probabilistic cellular automata-enhanced stereo vision simultaneous localization and mapping. *Measurement Science and Technology*, 22, 2011.

[7] Ioannis Kostavelis, Evangelos Boukas, Lazaros Nalpantidis, and Antonios Gasteratos. Visual odometry for autonomous robot navigation through efficient outlier rejection. *IEEE International Conference on Imaging Systems and Techniques, 2013, In Press*, 2013.

[8] Ioannis Kostavelis, Lazaros Nalpantidis, Evangelos Boukas, Marcos Aviles Rodrigalvarez, Ioannis Stamoulias, George Lentaris, Dionysios Diamantopoulos, Kostas Siozios, Dimitrios Soudris, and Antonios Gasteratos. Spartan: Developing a vision system for future autonomous space exploration robots. *Journal of Field Robotics*, 31(1):107–140, 2014.

[9] Rongxing Li, Fei Ma, Fengliang Xu, Larry H Matthies, Clark F Olson, and Raymond E Arvidson. Localization of mars rovers using descent and

surface-based image data. *Journal of Geophysical Research Planets (1991–2012)*, 107(E11):FIDO–4, 2002.

[10] Fabio Cozman, Eric Krotkov, and Carlos Guestrin. Outdoor visual position estimation for planetary rovers. *Autonomous Robots*, 9(2):135–150, 2000.

[11] Ron Li, Kaichang Di, Juwon Hwangbo, and Yunhang Chen. Integration of orbital and ground images for enhanced topographic mapping in mars landed missions. In *Proceedings of the Annual NASA Science Technology Conference (NTSC), College Park, MD*, 2007.

[12] Bach Van Pham, Artur Maligo, and Simon Lacroix. Absolute map-based localization for a planetary rover. In *Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA)*, 2013.

[13] Patrick JF Carle, Paul T Furgale, and Timothy D Barfoot. Long-range rover localization by matching lidar scans to orbital elevation maps. *Journal of Field Robotics*, 27(3):344–370, 2010.

[14] Evangelos Boukas, Antonios Gasteratos, and Gianfranco Visentin. Localization of planetary exploration rovers with orbital imaging:
a survey of approaches. In *Robotics and Automation Workshop: on Modelling, Estimation, Perception and Control of All Terrain Mobile Robots (ICRA Workshop), 2014 IEEE International Conference on*. IEEE, 2014.

[15] Tara A Estlin, Benjamin J Bornstein, Daniel M Gaines, Robert C Anderson, David R Thompson, Michael Burl, Rebecca Castaño, and Michele Judd. Aegis automated science targeting for the mer opportunity rover. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):50, 2012.

[16] Mark Woods, Andrew Shaw, Estelle Tidey, Bach Van Pham, Unal Artan, Brian Maddison, and Gary Cross. Seeker-autonomous long range rover navigation for remote exploration. In *Proceedings of the 11th International Syymposiumon Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 2012)*, 2012.

[17] Jong-Sen Lee. Digital image smoothing and the sigma filter. *Computer Vision, Graphics, and Image Processing*, 24(2):255–269, 1983.

[18] Ruan Lakemond, Sridha Sridharan, and Clinton Fookes. Hessian-based affine adaptation of salient local image features. *Journal of Mathematical Imaging and Vision*, 44(2):150–167, 2012.

[19] Tony Lindeberg. *Scale-space theory in computer vision*. Springer, 1993.

[20] Claude Elwood Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001.

[21] Pierre Soille. *Morphological image analysis: principles and applications*. Springer-Verlag New York, Inc., 2003.

[22] Akmal Rakhmadi, Nur ZS Othman, Abdullah Bade, Mohd SM Rahim, and Ismail M Amin. Connected component labeling using components neighborsscan labeling approach. *Journal of Computer Science*, 6(10):1099, 2010.

[23] Peter N. Belhumeur, João P Hespanha, and David Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 19(7):711–720, 1997.

[24] Jerome H Friedman. Flexible metric nearest neighbor classification. *Unpublished manuscript available by anonymous FTP from playfair. stanford. edu (see pub/friedman/README)*, 1994.

[25] Laurent Itti, Christof Koch, Ernst Niebur, et al. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 20(11):1254–1259, 1998.

[26] Jonathan Harel, Christof Koch, Pietro Perona, et al. Graph-based visual saliency. *Advances in neural information processing systems*, 19:545, 2007.

[27] Lazaros Nalpantidis, Georgios Ch Sirakoulis, and Antonios Gasteratos. A dense stereo correspondence algorithm for hardware implementation with enhanced disparity selection. In *Artificial Intelligence: Theories, Models and Applications*, pages 365–370. Springer, 2008.

[28] Raphael Labayrade, Didier Aubert, and J-P Tarel. Real time obstacle detection in stereovision on non flat road geometry through" v-disparity" representation. In *Intelligent Vehicle Symposium, 2002. IEEE*, volume 2, pages 646–651. IEEE, 2002.

[29] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Robotics-DL tentative*, pages 586–606. International Society for Optics and Photonics, 1992.