MATCHING SPARSE NETWORKS OF SEMANTIC ROIS AMONG ROVER AND ORBITAL IMAGERY

Evangelos Boukas①②, Antonios Gasteratos①, Gianfranco Visentin②

①*School of Engineering, Democritus University of Thrace, Vas. Sophias 12, GR-67100 Xanthi, Greece. e-mail: evanbouk@pme.duth.gr, agaster@pme.duth.gr*

②*Automation and Robotics Section, ESA e-mail: Gianfranco.Visentin@esa.int*

ABSTRACT

The ability of rovers to localize themselves on a global reference is of high importance for the evolution of future robotic space exploration activities. This is why ESA has funded a research through the Network/Partnering Initiative (NPI) to conceptualize and develop methods concerning the global self-localization of rovers. The approach followed in our NPI activity is the one of the distinctive landmarks, or Regions of Interest (ROIs). In this scenario the rover is fed with the location and properties of all the ROIs in a predefined area, e.g. the uncertainty ellipse of its landing, as captured by means of satellite imagery. The rover detects ROIs utilizing its onboard sensors and then matches these constellations with the ones that had been fed to it initially. The paper in hand presents the methods utilized, in the context of global self localization under the aforementioned NPI activity, to match sparse networks of semantic ROIs among rover and orbital imagery. The matching method has been tested on real experimental data as well as on simulated ones so as to assess the performance optimally.

Key words: Global Localization; Rovers; Robot Vision; Navigation; ESA.

1. INTRODUCTION

The accurate localization of rovers is a perquisite to create truly autonomous space exploration. The first rover that had some level of autonomy (eventhough only 2-3 m per sol) employing two main sensors, i.e. wheel encoders and sun sensors [1], was the Pathfinder mission's rover "Mars Sojourner". The successors of Sojourner, Mars Exploration Rovers (MER) *Spirit* and *Opportunity* have higher autonomy including stereo camera visual odometry (VO) aided by Inertial Measurment Units (IMU) [2], yet at first stages they where using VO mostly for slip checking. The latest NASA rover *Curiosity* employs VO at intervals of approximately 1m in trajectory, along with obstacle avoidance, to perform autonomous driving [3]. ESA's ExoMars rover has been designed in a way that it will use nominally visual odometry on a high rate employing a dedicated coprocessor [4].

The space exploratory rovers have evolved enough to be able to navigate safely (VO and obstacle avoidance) on planetary terains and, thusly, the next step is the assignment of more complex activities. When discussing advanced mission scenarios, an obvious example that comes to mind is always the Mars Sample Return (MSR) mission, which is foreseen to be an ESA/NASA cooperation [5]. The increase in rovers' autonomy is considered to be the next big bet towards advanced and cooperative extraterrestrial missions, such as the MSR. It is obvious that in order for the future space exploratory rovers to be capable of participating into complex missions and to coordinate with the rest of the apparatus, a very crucial feature should be their ability to localize themselves within an absolute global frame, e.g. the Mars body-fixed rotating frame. In order to globally localize a rover on Martian terrain, based on the fact that no Global Navigation Satellite System (GNSS) is available, the utilization of known landmarks is mandatory. The localization of a rover on a georeferenced orbital image is equivalent to global localization [6] and, hence, should considered sufficient. Toward this end, ESA has funded an Networking/Partnering Initiative entitled: "Methods to Refine the Self-Localization of Planetary Rovers Using Orbital Imaging", the preliminary results of which we are presenting here.

2. STATE OF THE ART

The localization of space exploratory rovers has been an active area of research for over 10 years [7]. However, most of the research is based on the derivation of their position and trajectory relatively to their initial landing position, usually based on visual odometry [8, 9, 10]. The authors in [11] have implemented visual odometry algorithms on FPGA devices in order to overcome contemporary power and processing restrictions of space exploration rovers. In order to be able to perform large scale

and more ambitious tasks, rovers must be able to localize themselves on a global scale. The authors in [12] made use of SIFT features and a 3D morphological filter to calculate the position of the rover within the orbital image. A LIDAR device has been employed to globally localize a rover within the imaging region of a georeferenced digital elevation model (DEM) [13, 14, 15]. The approach utilizes prominent features from local and global DEMs. The correspondences between the rover and ground features are computed by means of DARCES method [16]. Lourakis and Hourdakis [17] present a method to utilize boulders extracted from robot stereo cameras and orbital imagery to improve the visual odometry of a rover. A review of the related work concerning global localization is presented in [18]. The approach that is investigated under the ESA NPI activity entitled "Methods to Refine the Self-Localization of Planetary Rovers Using Orbital Imaging" comprises of the following sub-modules as presented in [19]:

- Detection and classification of Regions of Interest (ROI) on orbital images, covering the area of the rovers traverse, and formation of a Global Network (GN) of annotated ROIs.
- Relative localization based on rover sensory (cameras, IMU).
- Detection and classification of Regions of Interest (ROI) along rovers route and formation of a Local Network of labelled ROIs.
- Matching (or fitting) of the LN on the GN.

The extraction of commonly observable ROIs has been presented in [20], while the paper in hand focuses on the assessment of our method for the matching of LN and GN.

3. ROVER RELATIVE LOCALIZATION

The relative localization module of our system employs a stereo visual odometry algorithm and an Inertial measurement Unit (IMU) sensor fusion. We managed to emulate the contemporary rovers' capabilities by executing the VO algorithm only on consecutive frames, and not a buffer of previous ones. These restrictions drove as to use neither a total optimizer such as Bundle Adjustment (BA) nor even a local (windowed) BA. The presented relative localization system achieves high accuracy in a large - more than $1km$- route. The main components of the rover localization algorithm that were implemented are:

- Feature Extraction
- 3D reconstruction
- Visual Odometry
  - Outlier Rejection
  - Motion Estimation
- VO & IMU Fusion.

4. DETECTION & CLASSIFICATION OF ROIS ON ROVER AND ORBITAL IMAGERY

In this research we focus mainly on two types of ROI, namely rocks and large flat rock outcrops. These basic types of boulders have been investigated as well by previous works in the field of automatic targeting on Mars terains [21, 22]. In the following subsections, the detection will be described, followed by the classification.

4.1. Detection Metrics

Two main metrics are employed in the detection of "candidate" rocks or outcrops: the Hessian analysis and the Entropy response. The first metric is the Hessian of the image.

The hessian of an image has been proven to provide a a good reference for detection of blob-like and ellipse-like ROIs [23] It can be calculated by the second order partial derivatives of an image along its directions (Eq. 1).

$$H(x,y) = \begin{Bmatrix} \frac{\partial^2 I}{\partial x \partial x} & \frac{\partial^2 I}{\partial x \partial y} \\ \frac{\partial^2 I}{\partial y \partial x} & \frac{\partial^2 I}{\partial y \partial y} \end{Bmatrix} \quad (1)$$

We know from the literature that the prominent blobs result in extreme eigenvalues of the Hessian matrix [23]. A good definition of extreme values is the difference from the mean of the eigenvalues higher than 2σ. This method of detecting blobs is obviously introducing errors which are, nevertheless handled by means of classification. The second metric utilized is the Local entropy, which defines the "randomness" of an image. Since, the Entropy does not differentiate between low and high differences in intensity we employ the weighted entropy (Eq. 2):

$$LE^w = -\sum_{i=0}^{255} w(i) \cdot p(i) \cdot \ log_2(p(i)), \quad (2)$$

where $w(i)$=1-$G(\mu,\sigma_e)$, μ is the mean of the neighborhood and ${\sigma_e}^2 = 8$.

4.2. Candidate ROI extraction

The extraction of possible ROIs is performed with both appearance (texture) and geometric thresholds. The appearance metric thresholds are adaptive in the sense that they are calculated as a percentage of the distribution, while the geometric constrains are fixed. The thresholds for the extraction of rocks and outcrops on orbital and rover data is as follows:

Orbital

- Rock:
  - ∘Geometric Constrains: More than $0.6m^2$.
  - ∘Texture Constrains: ($> 2 \cdot \sigma$) on the distribution of entropy and ($> 2 \cdot \sigma$) on the distribution the eigenvalues of the Hessian
- Outcrop:
  - ∘Geometric Constrains: More than $8m^2$.
  - ∘Texture Constrains: ($< -2 \cdot \sigma$) on the distribution the eigenvalues of the Hessian.

Rover

- Rock:
  - ∘Geometric Constrains: More than $0.6m^2$, higher than $0.25m$.
  - ∘Texture Constrains: Fusion of ($> 2 \cdot \sigma$) on the distribution the eigenvalues of the Hessian and ($< -2 \cdot \sigma$) on the distribution the eigenvalues of the Hessian.
- Outcrop:
  - ∘Geometric Constrains: More than $5.5m^2$.
  - ∘Texture Constrains: Logical OR of the ($< -2 \cdot \sigma$) on the distribution the eigenvalues of the Hessian and ($< -2 \cdot \sigma$) on the distribution of Entropy.

Since the rover is traversing we have implemented a method to track the ROIs as they appear and disappear in the rovers FOV. When a new ROI is detected then an ellipse containing this object is created. Then as the rover detects a new ROI, it compares it with the previously saved one, if the ROIs coincide, and then the ellipse is updated. and the process goes on.

4.3. ROI Classification

The classifier utilized in our system is the k-NN classifier. The k-NN classifier used a previously trained dataset to compare with a new sample and classify it appropriately. The approach used for the classification problem is the one of multiple binary classifiers. That means that in oru test we have 4 classifiers: “Rock Orbital”, “Outcrop Orbital”, “Rock Rover”, “Outcrop Rover”. The classification scheme for our system comprises 3 main parts:

- Feature extraction
- Dimensionality reduction
- Classification

Firstly, for each candidate ROI the features that are taken into consideration are the following: [“RGB”, “HSV”, “GRAYSCALE”,“GABOR”, “ENTROPY”]. Then a dimensionality reduction algorithm is performed. We employed the Principal Component Analysis (PCA) (Karhunen−Loéve method), an unsupervised method, able to map linearly a high dimensional feature vector to a low one [24]. Finally the classifiers decide on which class, every new sample, is closer to.

5. ROI NETWORK FORMATION

As we mentioned there exist two networks in our system: 1) The Global Network (GN), which includes the ROIs as seen from the orbital images and, 2) the Local Network (LN), that contains the ROIs that have been explored by the rover. The global Network is created offline and is an one time procedure, while the LN is somewhat more complicated as it evolves as the rover moves. The GN contains the cartographic coordinates of the ROIs, while the Local Network contains the location of ROIs based on its relative localization.

6. ROI MATCHING

As the rover explores the area it localizes itself on a relative frame and creates a LN of ROIs, as explained previously. The matching procedure is similar to the one presented in [15] with substantial differences, such as the utilization of labeled landmarks attained from stereo cameras instead of LIDARs, the utilization of 2D instead of 3D points and the network nature of our matching method. For simplicity reasons, the orbital-GN correspondence is performed based on a 2D version of the DARCES algorithm [16]. Since the orbital images are orthorectified, an accurate planar projection of the ground surface is already available allowing us to consider 2D DARCES. As a first step, the DARCES algorithm creates several hypotheses of correspondence among three static nodes of the LN (usually referred to as the control points) and several GN nodes. The triplets of points form triangles. A hypothesis is created if and only if the edges of the triangles (candidates for correspondence) differ less than a predefined threshold d_D. Initially, all points within the GN are considered as candidates to match with the first control point (usually referred to as the primary point). Then, the algorithm seeks for correspondences to the the second point (secondary point) within the range of a radius. For the 2D case we consider, the locus of the correspondence points is, thusly, a circle, whereas in [16] it is a sphere. Next the algorithm reckons the rest of the triangles’ edges, thus forming a list of hypotheses. This process is repeated for several control points in the LN.

In this paper we propose improvements on the original DARCES algorithm by introducing extra constrains in the detection of hypotheses, with a view to simultaneously improve its computational efficiency and accuracy.

Firstly, the orientation (clockwise or anticlockwise) of the triangles is taken into consideration, so that only triangles with the same orientation can be matches. Last, a semantic constraint is suggested in the calculation of the hypotheses, due to the fact that the nodes are labeled (Sec. 4.3). The modified DARCES algorithm may even return no correspondences, providing that none of the above criteria is met. This has been proven particularly important in dealing with false detections in either orbital or rover imagery. We should note here that the method described is an opportunistic one and, should the rover not approach major ROIs in its route, then it can not be globally localized. As far as robustness and noise tolerance are considered, we have not followed the "winner takes it all" approach, but we have rather implemented a voting scheme in which the whole (minor and major) global and local ground networks are contributing. The two corresponding pairs of points provide the transformation ${}^{l}T_{g}$ which aligns the local map to the global one. Then whole relative localization trajectory is aligned with this transformation leading to a globally localized rover position and route.

7. EXPERIMENTAL RESULTS

7.1. Seeker Activity Dataset in Atacama Desert

We have carefully selected the dataset upon which our algorithm is going to be tested. The dataset of the Seeker activity [25] has been crated in the Atacama desert, Chile, which has been proven to be the place in Earth with the most similar characteristics to Mars environment [26]. The dataset contains information acquired both by a rover and an orbital imagery analog. Figure 1 depicts the detection and classification of ROIs on orbital and rover images of the Atacama Seeker dataset. Following the extraction of orbital and rover ROIs, the Global and Local networks, respectively. Then the matching procedure is performed to fit the LN in the GN. The final alignment of local and global networks and, hence, the solution to the global localization problem is depicted in fig. 2. The total route of the rover is more than $1000m$ and the localization error after global alignment is $13.34m$.

7.2. Simulation Dataset Creation

In order to evaluate the repeatability of our matching procedure, multiple simulated Local Networks have been considered. The simulated *LNs* were created based on the real Global Network. According to our analysis of the Atacama dataset, the rover's route of 1 km coincides with 24 ROIs, of which the rover was able to detect and classify 21, i.e. 87.5%. In this typical day of exploration the rover has travelled a trajectory of 1km or 600m of Euclidean distance relative to its starting point, i.e. 60% of Euclidean distance over actual route travelled.

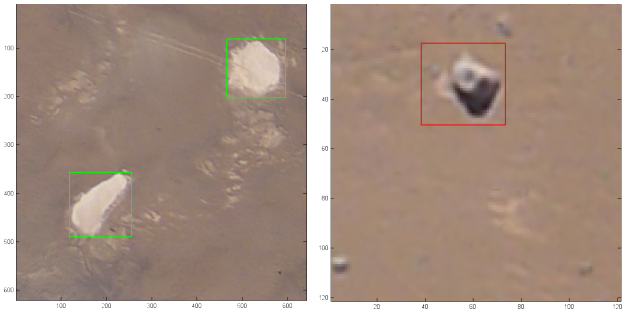

(a) Example Detection and Classification for two indicative major ROIs on the orbital image. Note, how the classification offers efficient outlier rejection

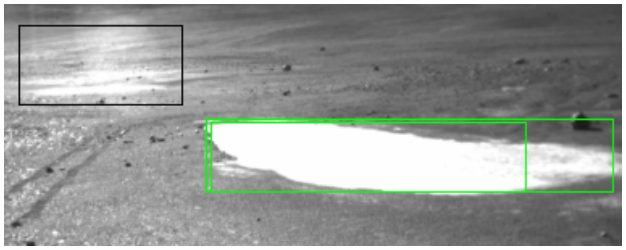

(b) Example Detection and Classification of ROIs on rover images

Figure 1. The orbital and rover ROI extraction.

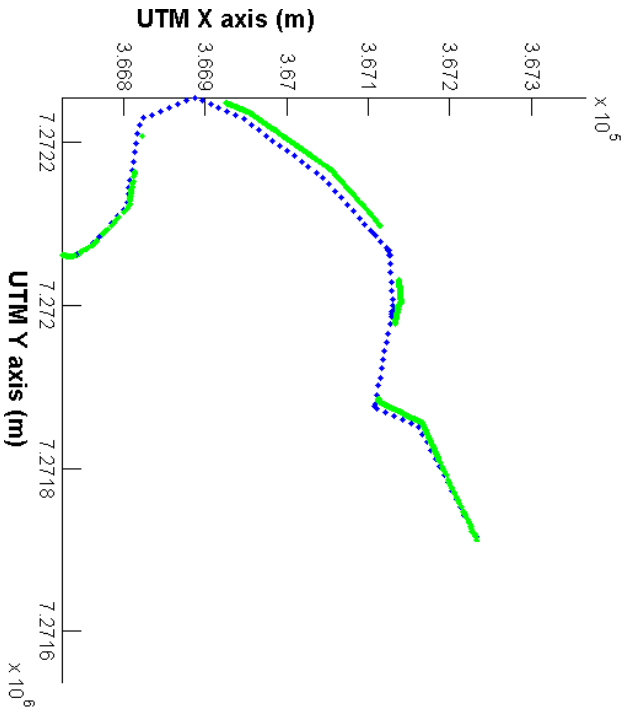

Figure 2. Fitting of the local (LN) to the global network (GN) to provide global Localization.

Bearing in mind the afforementioned observations we created the simulated *LNs*, which correspond to simulated trajectories, as follows:

- Create the starting and ending points of simulated trajectories:

  For each point of the *GN* find the points which have a distance D_{pair}: $600m \leq D_{pair} \leq 1km$

  Create pairs of starting and ending points $\{S_p^N, E_p^N\}$, where N is the number of pairs.

- Create the simulated trajectories: Select all the possible trajectories connecting the S_p to E_p with a minimum of $24 - 20\% = 19$ and a maximum of $24 + 20\% = 29$ nodes, which have a total length of less than 1km.

- Restrict the detectability of ROIs: On every trajectory remove 20% of the nodes to simulate the ability of the rover to detect and classify ROIs.

- Simulate the localization error: on each simulated LN move each ROI randomly in a range of $[-1\%,1\%]$ of the total distance from the starting point, so as to simulate the maximum error of the visual odometry.

7.3. Simulation Dataset Evaluation

The *GN* of the Atacama dataset includes 174 ROIs that result in more than 58000 simulated *LNs*. The matching algorithm was executed against the *GN*. The *LN* with the best score of matching is counted as a fixed match. The accuracy of the matching algorithm is higher than 99.3%. We have to point that eventhough these results are very encouraging, there is no evidence that this algorithm is a panacea to the global localization. Our algorithm is an opportunistic one, in the sense that it will work only if there exist enough ROIs in the scene. In fact our results here prove that, should enough ROIs exist, our method is able to localize a rover on a global frame.

8. CONCLUSIONS

In this paper we have thoroughly described and assessed the method utilized in our global localization system to match the landmarks recognized by a rover to the ones detected in orbital images. We have tested our algorithm both on real and on simulated data. We have proven that our algorithm is able to provide global localization, as long as there exist some distinctive ROIs. However there exist some improvements that could be implemented in our system. One of them is the ability of the system to incrementally compute the global localization, in a probabilistic scheme, rather than on the end of the route.

ACKNOWLEDGMENTS

This work was funded by the European Space Agency via the Network/Partnering Initiative under the contract No. $4000109064/13/NL/PA$.

REFERENCES

[1] Andrew H Mishkin, Jack C Morrison, Tam T Nguyen, Henry W Stone, Brian K Cooper, and Brian H Wilcox. Experiences with operations and autonomy of the mars pathfinder microrover. In *Aerospace Conference, 1998 IEEE*, volume 2, pages 337–351. IEEE, 1998.

[2] Khaled S Ali, C Anthony Vanelli, Jeffrey J Biesiadecki, Mark W Maimone, Yang Cheng, A Miguel San Martin, and James W Alexander. Attitude and position estimation on the mars exploration rovers. In *Systems, Man and Cybernetics, 2005 IEEE International Conference on*, volume 1, pages 20–27. IEEE, 2005.

[3] Mark Maimone. Curiouser and curiouser: Surface robotic technology driving mars rover curiositys exploration of gale crater. In *Robotics and Automation Workshop: on Planetary Rovers (ICRA Workshop), 2013 IEEE International Conference on*. IEEE, 2013.

[4] 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.

[5] Richard Mattingly and Lisa May. Mars sample return as a campaign. In *Aerospace Conference, 2011 IEEE*, pages 1–13. IEEE, 2011.

[6] RL Kirk, E Howington-Kraus, MR Rosiek, JA Anderson, BA Archinal, KJ Becker, DA Cook, DM Galuszka, PE Geissler, TM Hare, et al. Ultrahigh resolution topographic mapping of mars with mro hirise stereo images: Meter-scale slopes of candidate phoenix landing sites. *Journal of Geophysical Research: Planets (1991–2012)*, 113(E3), 2008.

[7] Mark Maimone, Yang Cheng, and Larry Matthies. Two years of visual odometry on the mars exploration rovers. *Journal of Field Robotics, Special Issue on Space Robotics*, 24(3):169–186, 2007.

[8] D. Scaramuzza and F. Fraundorfer. Visual odometry: Part I: The first 30 years and fundamentals. *IEEE Robotics and Automation Magazine*, 18(4): 80–92, 2011.

[9] Friedrich Fraundorfer and Davide Scaramuzza. Visual odometry: Part II: Matching, robustness, optimization, and applications. *Robotics & Automation Magazine, IEEE*, 19(2):78–90, 2012.

[10] A. Johnson, S. Goldberg, C. Yang, and L. Matthies. Robust and efficient stereo feature tracking for visual odometry. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 39–46, 2008.

[11] 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.

[12] Kaichang Di, Zhaoqin Liu, and Zongyu Yue. Mars rover localization based on feature matching between ground and orbital imagery. *Photogrammetric engineering and remote sensing*, 77(8):781–791, 2011.

[13] Paul Furgale, Pat Carle, and Timothy D Barfoot. A comparison of global localization algorithms for planetary exploration. In *Intelligent Robots and Systems (IROS), 2010 IEEERSJ International Conference on*, pages 4964–4969. IEEE, 2010.

[14] 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.

[15] Patrick JF Carle and Timothy D Barfoot. Global rover localization by matching lidar and orbital 3d maps. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 881–886. IEEE, 2010.

[16] Chu-Song Chen, Yi-Ping Hung, and Jen-Bo Cheng. Ransac-based darces: A new approach to fast automatic registration of partially overlapping range images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 21(11):1229–1234, 1999.

[17] Manolis Lourakis and Emmanouil Hourdakis. Planetary rover absolute localization by combining visual odometry with orbital image measurements. In *Proc. of the 13th Symposium on Advanced Space Technologies in Automation and Robotics (ASTRA'15)*. European Space Agency, 2015.

[18] Evangelos Boukas, Antonios Gasteratos, and Gianfranco Visentin. Localization of planetary exploration rovers with orbital imaging: a survey of approaches. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) workshop on "Workshop on Modelling, Estimation, Perception and Control of All Terrain Mobile Robots"*, Hong Kong, China, 31 May - 7 June 2014.

[19] Evangelos Boukas, Antonios Gasteratos, and Gianfranco Visentin. Towards orbital based global rover localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, Washington, 26-30 May 2015.

[20] Evangelos Boukas, Antonios Gasteratos, and Gianfranco Visentin. Extraction of common regions of interest from orbital and rover (ground) acquired imagery. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*, Montreal, Canada, 2014.

[21] Mark Woods, Andy Shaw, Dave Barnes, Dave Price, Derek Long, and Derek Pullan. Autonomous science for an exomars roverlike mission. *Journal of Field Robotics*, 26(4):358–390, 2009. ISSN 1556-4967. doi: 10.1002/rob.20289. URL `http://dx.doi.org/10.1002/rob.20289`.

[22] 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.

[23] 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.

[24] 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.

[25] 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. *International Symposiumon Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2012.

[26] David Wettergreen, Nathalie Cabrol, Vijayakumar Baskaran, Francisco Calderón, Stuart Heys, Dominic Jonak, RA Luders, David Pane, Trey Smith, James Teza, et al. Second experiments in the robotic investigation of life in the atacama desert of chile. In *Proc. 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space i-SAIRAS*, 2005.