Q1
Q2

Field Trial Results of Planetary Rover Visual Motion Estimation in Mars Analogue Terrain

Joseph Nsasi Bakambu, Chris Langley, Giri Pushpanathan, W. James MacLean, and Raja Mukherji
MDA Corporation, 9445 Airport Road, Brampton, Ontario, Canada L6S 4J3
e-mail: joseph.bakambu@mdacorporation.com, chris.langley@mdacorporation.com, giri.pushpanathan@mdacorporation.com, raja.mukherji@mdacorporation.com, james.maclean@utoronto.ca
Erick Dupuis
Space Exploration, Canadian Space Agency, 6767 Route de l'Aéroport, Saint-Hubert, Quebec, Canada J3Y 8Y9
e-mail: erick.dupuis@asc-csa.gc.ca

Received 7 January 2011; accepted 4 November 2011

This paper presents the Mojave Desert field test results of planetary rover visual motion estimation (VME) developed under the "Autonomous, Intelligent, and Robust Guidance, Navigation, and Control for Planetary Rovers (AIR-GNC)" project. Three VME schemes are compared in realistic conditions. The main innovations of this project include the use of different features from stereo-pair images as visual landmarks and the use of vision-based feedback to close the path-tracking loop. The multiweek field campaign, conducted on relevant Mars analogue terrains, under dramatically changing lighting and weather conditions, shows good localization accuracy on the average. Moreover, the MDA-developed inertial measurement unit (IMU)-corrected odometry was reliable and had good accuracy at all test locations, including loose sand dunes. These results are based on data collected during 7.3 km of traverse, including both fully autonomous and joystick-driven runs.

1. INTRODUCTION

One of the continuing challenges for future unmanned exploration rover missions on the Moon and Mars is the ability to accurately estimate the rover's position and orientation. In the absence of an equivalent to the Global Positioning System on Earth, designers of future systems must exploit a combination of vehicle odometry, inertial sensing, and visual information to obtain this localization information. Future missions such as ESA's ExoMars and NASA's Mars Science Laboratory rover require the rover to autonomously traverse from hundreds of meters to one kilometer daily at speeds of up to 100 m/h (Volpe, 2006). Equally challenging is the need to localize the position of the rover to an accuracy of between 1% and 4% of the distance traveled. Accurate localization is arguably the most fundamental competence required for long-range autonomous navigation. For this reason, MDA Space Missions and the Canadian Space Agency (CSA) have embarked on a development path to further their capability for visual motion estimation (VME) of planetary rovers.

VME algorithms have recently seen considerable interest from the planetary exploration rover community as a solution for accurate localization. On the Mars Exploration rovers (MERs), visual odometry was not considered part of the main system for localization, but was shown to perform relatively well. In Maimone, Chang, and Matthies (2007), JPL reported that "Visual Odometry software has enabled precision drives over distances as long as 8 m on slopes greater than 20 degrees, and has made it possible to safely traverse the loose sandy plains of Meridiani." The algorithm works by tracking Harris corner features (Harris & Stephens, 1988) in a stereo image pair from one frame to the next. Thus, the problem is one of determining "the change in position and attitude for two pairs of stereo images by propagating uncertainty in a 3D to 3D pose estimation formulation using maximum likelihood estimation." The evaluation tests conducted at the JPL Marsyard and Johnson Valley, California showed that the absolute position errors were less than 2.5% over the 24-m Marsyard course, and less than 1.5% over the 29-m Johnson Valley course. The rotation error was less than 5.0 deg in each case.

LAAS/CNRS (Laboratoire d'Architecture et d'Analyse des Systèmes/Centre National de la Recherche Scientifique) VME is based on the frame-to-frame pixel-tracking method. Landmarks are extracted from images by finding points of interest identified by image intensity gradients. The test results using the Lama rover (Lacroix et al., 2000) showed an error of 4% on a 25 m traverse. After improvement of the algorithm in Lacroix et al. (2002), an overall error of 2% on a 70 m traverse was achieved.

A survey of the literature (Biesiadecki et al., 2005; Corke et al., 2004; Maimone et al., 2007) shows that the state of the art in localization has yet to meet the ExoMars localization accuracy requirement of 1% of the traveled distance, without the use of computationally intense

View this article online at wileyonlinelibrary.com • DOI: 10.1002/rob.21409

methods (for example, bundle adjustment[1]) that either necessitate much higher power flight avionics or execution as a postprocessing step. The most promising results to date are published in Konolige, Agrawal, and Sola (2007) and Souvannavong, Lemaréchal, Rastel, and Maurette (2010); both of the approaches use sparse bundle adjustment but are tested using a limited number of outdoor data sets. Sibley, Mei, Reid, and Newman (2010) derived a relative bundle adjustment that, instead of optimizing in a single Euclidean space, works in a metric space defined by a manifold. The results on over 850,000 images (covering 142 km) indicate the accuracy and scalability of the approach. Previous work at MDA has shown that a full simultaneous localization and mapping (SLAM) approach yields good accuracy, but has a high computational burden, restricting the vehicle to lower speeds (Bakambu et al., 2008). Conversely, if no map is maintained, a frame-to-frame pixel-tracking technique has been shown to operate with acceptable speed, but at the cost of decreased accuracy. MDA proposed a multiframe VME algorithm that attempts to obtain a suitable balance between accuracy and speed. The results in hardware indicate that traverses of greater than 200 m are possible to an accuracy of between 1% and 4% under planetary conditions. Our approach uses 3D odometry and a stereo pair to identify visual landmarks using the scale invariant feature transform (SIFT) (Lowe, 1999). Although the result is commensurate with the ExoMars localization accuracy requirements, further development is needed to improve accuracy and achieve robustness under a variety of terrain conditions, and decreased computational cost for greater practicality for a flight mission.

The latest MDA VME approach is presented in this paper. This extended localization system was tested in a three-week field campaign in the Mojave Desert using the robot shown in Figure 1. The new features being tested, which yield improvements from the previous work (Bakambu et al., 2008), include the following:

- Use of different features from stereo-pair images as visual landmarks. SIFT features are invariant to image translation, scaling, and rotation and partially invariant to illumination changes and affine projection, making them suitable landmarks for use in motion estimation (Se, Barfoot, & Jasiobedzki, 2005). However, in testing of the VME system to date it has been found that in some cases SIFT features alone are insufficient for tracking. For example, on sandy terrain with smooth rocks, the SIFT detector often produces only small-scale features, which are often unstable. Other features extracted from the stereo-pair images in the current work are maximally stable extremal regions (MSERs) (Donoser & Bischof, 2006; Kristensen & MacLean, 2007; Matas, Chum, Urban, & Pajdla, 2002) and Harris-Laplace (Mikolajczyk et al., 2005) features.
- Extension of the localization system to use inclinometry to improve roll/pitch angle estimation. Inclinometry maintains an observable estimate of the gravity vector, instead of relying on roll and pitch rate integration.
- Extension of the localization system to use a long-range and wide-field-of-view active 3D sensor to extract long-range fixed landmarks for enforcing VME observability, and thus improving the accuracy of the approach.
- Use of the improved VME pose estimate as feedback for closed-loop motion control. Our novel approach uses frequency lifting (Davis & Vinter, 1985) to allow VME to be used while the rover is moving, rather than as an a posteriori estimate once the traverse is completed.

Figure 1. Rover test bed in the Mojave Desert.

The VME technologies described above were integrated as part of a complete long-range GN&C solution for planetary rovers, including 3D terrain modeling and assessment, motion planning and tracking, hazard detection and avoidance algorithms, and high-level supervisory control based on CORTEX, a tool developed by the Canadian Space Agency (CSA) for design and real-time execution of state machines (Dupuis et al., 2005).

The rest of this paper is organized as follows. Section 2 details a typical field test scenario and describes the test environments through pictures and maps reconstructed from sensor data. The challenges and the issues observed during the field test campaign are also presented. The experimental results and lessons learned are presented and discussed in Section 3. Section 4 concludes the paper.

2. SYSTEM OVERVIEW

2.1. Operational Scenario

To better understand the sequence of operations during the test campaign, this subsection describes a typical

[1]Triggs, McLauchlan, Hartley, and Fitzgibbon (2000): "Bundle adjustment is the problem of refining a visual reconstruction to produce jointly optimal 3D structure and viewing parameter (camera pose and/or calibration) estimates."

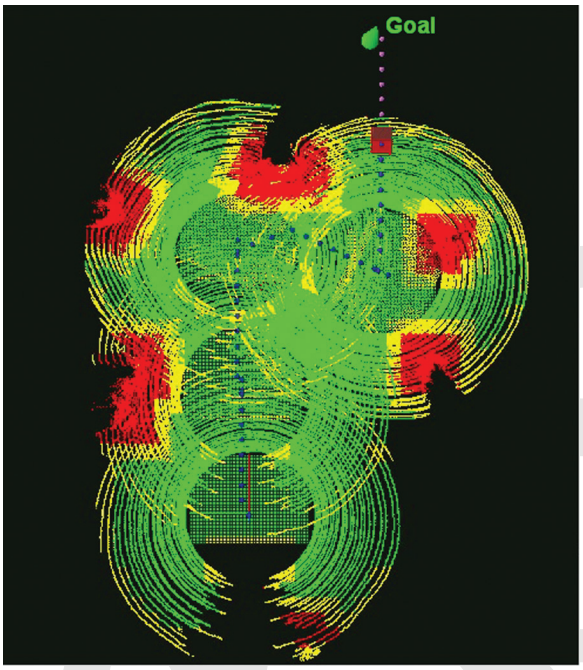

Figure 2. Typical field test scenario, showing 3D terrain data from the sensor, color-coded according to traversability. Dots indicate the planned path.

long-range exploration scenario for a planetary rover. This scenario can be classified as autonomous navigation in an unknown terrain

A typical autonomous traverse (see Figure 2) consists of the following operations:

- A remotely located operator or scientist selects one or more goal locations that are outside of the rover's onboard sensor ranging envelope and/or obscured by an obstacle when viewed from the start location.
- The rover takes a high-resolution, medium-range (10 to 15 m) scan of its surrounding environment using its integrated sensor suite.
- The rover constructs an internal representation of the surrounding terrain, conducts the terrain traversability assessment, and plans a hazard-free path toward the goal. This path is safe inside the sensed area, and aims directly at the goal outside the sensed area.
- The operator station displays the terrain map, including traversability map overlays, and the planned path.
- The rover tracks the planned path up to the boundary of the medium-range map and stops. Telemetry and environment data are collected during the motion and at the stop location.
- Mapping, terrain assessment, and path planning are automatically repeated until the final destination is reached. If the goal is not reachable, the rover will move as close as possible to the defined goal (as illustrated in Section 0, Figure 19) and request a new goal or a new task.

All of the above steps are coordinated and monitored by a high-level supervisory control and data acquisition module developed using the CSA's CORTEX autonomy framework (Dupuis et al., 2005), described in the next subsection.

2.2. High-Level Supervisory Control

Exploration of large unknown planetary environments will rely on rovers that can autonomously cover distances of kilometers and maintain precise information about their location with respect to local terrain features. This requires the capability of perceiving, modeling, and assessing the environment, planning and tracking collision-free paths, and detecting and avoiding hazards autonomously. Obviously, planetary exploration robots will require a high level of autonomy to perform tasks more efficiently.

Over the last few years, the Canadian Space Agency (CSA) has designed, implemented, and tested different autonomy techniques on typical autonomous robotics scenarios: finite state machines (FSM), hierarchical task networks (HTN), and goal decomposition hierarchies (GDH). Details on enumerated autonomy techniques are provided in (Dupuis et al., 2005). Based on the experience gained on testing these autonomy techniques, CSA designed and implemented the CORTEX Autonomy Toolbox, which merges the advantages of these techniques. CORTEX implements hierarchical FSM (HFSM), which allows a high-level FSM to invoke a lower-level FSM.

To take full advantage of the CORTEX autonomy toolbox, HFSM formalism is used for the AIR-GNC system. Note, however, that to actually design the state machines, a formal method is required that guarantees that the system behavior obeys a desired specification and that all desired behaviors are executed in a nonblocking and nonconflicting fashion (Ramadge & Wonham, 1987). The methodology employed mimics the theory of state space methods in control theory. It is known in the literature as the theory of discrete event systems, originally developed by W. M. Wonham in 1987 (Ramadge & Wonham, 1989).

2.3. Terrain Modeling, Traversability Assessment, and Path Planning

Two different path planning algorithms were developed, implemented, and tested in this project. The MDA-developed path planner, which is based on a grid cost map computed from a terrain traversability assessment and D* search (Stentz, 1995), is used for short-range path planning (see Figure 2). The CSA-developed hybrid path planner, which uses a triangulated mesh, is used to plan medium-range paths. The hybrid path (see Figure 3) is obtained by first computing an A* path using a triangulated mesh,

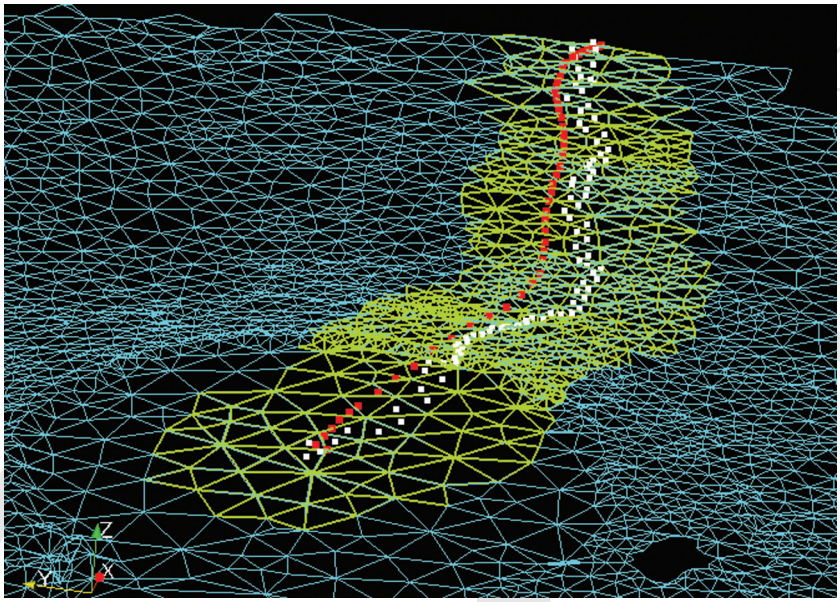

Figure 3. Triangulated mesh (cyan) with extracted traversable corridor (green) along the planned A* path (white dots) and hybrid path (red dots).

extracting a traversable corridor along the A* path, and then applying a fluid theory-based path planner to find a smooth and short path within the traversable corridor. More details on CSA's hybrid path planning can be found in (Gingras et al., 2010). The terrain modeling function is based on (Gingras et al., 2010).

2.4. Localization

The rover test bed is equipped with wheel odometry, an AIMS FOG 2 navigation-grade inertial measurement unit (IMU), and a Point Grey XB3 commercial stereo camera. The autonomous traverse can be executed using localization estimates from either the enhanced IMU-corrected odometry or the VME, as described in Section 1, and in (Bakambu et al., 2008). It is worth noting that prior work in this field has focused on VME as an open-loop observer for providing accurate localization a posteriori when a destination is reached. In our novel approach, the low-rate and delayed VME measurements and the high-rate IMU–odometry measurements are fused to yield a high-rate and smooth position and attitude measurement signal. This technique is called *frequency lifting* (Davis and Vinter, 1985). The frequency-lifting output is then used as the feedback signal for closed-loop path tracking.

3. FIELD TRIAL CAMPAIGN

The goal of the AIR-GNC project was to advance the capabilities of autonomous navigation (terrain modeling and assessment, path planning, path tracking, and localization) for future planetary missions. To demonstrate these capabilities, it was necessary to test them in a relevant planetary environment, which includes rocks, cracks, loose sand, and slope hazards. The dry lake beds of the Mojave Desert in the Southern United States have been used as Mars analogues in the past by NASA's Jet Propulsion Laboratory (Huntsberger et al., 2002). The terrain offers a variety of soil conditions, from dry, fine-grained sand to hard-packed clay to loose gravel. Slopes and rock fields can also be found near the edges of the lake beds. The area is also relatively free of vegetation. The test bed with the newly developed software modules was fully exercised because of the presence of slope and rock hazards, the likelihood of wheel slippage, and the presence of representative natural image features. The main focus of this paper is on presenting the performance of the various VME schemes under realistic conditions.

The advantage of conducting a multiweek test campaign is the large body of data that can be stored and analyzed. This allows a more meaningful statistical analysis to be conducted, and hence, a truer representation of the localization performance accuracy across a wide variety of terrains. During this field campaign, 104 separate runs were logged, spanning 7.3 km of traverse in six distinct terrain classes. In the literature, very few research groups have conducted multiweek test campaigns in representative environments. Wettergreen et al. (2005) and Wettergreen et al. (2008) deployed the rover "Zoe" in the Marslike Atacama Desert of Chile. During three seasons of field testing campaigns, this robot has traveled over 250 km autonomously. Furgale and Barfoot (2010) deployed a six-wheeled rover

Figure 4. Map showing the Desert Studies Center, Baker, and the relative locations of the test sites. *Source:* Google.

platform in planetary analogue terrain on Devon Island in the Canadian Arctic. This rover has traveled more than 32 km, 99.6% of the distance without human intervention.

3.1. Field Test Terrain Landscapes and Reconstructed Maps

The home base for the field trials was the Desert Studies Center (DSC), located near Baker, California. The DSC is a field station of the Biology Department of the California State University at Fullerton, and hosts desert biology and geology researchers and students from academia and industry. Considering its isolated location, the DSC is an excellent facility for use as a home base for field trials.

The field tests were conducted in and around Silver Dry Lake, Silurian Dry Lake, and Dumont Little Dunes in the Mojave Desert (see Figure 4). The test environs offer many interesting terrains (in terms of soil characteristics, visual features, and rock, bush, and slope hazards), all within a small geographical area.

3.1.1. Silver Dry Lake Test Locations

Three locations were selected for testing in the Silver Dry Lake bed: the Playa, Boulder, and Plateau locations. The Playa location is characterized by relatively flat, open, and hazard-free terrain. The soil is hard clay, which has cracked as it dries. This produces a fractal-like pattern of fissures on the surface. The fissures also make the surface bumpy, which can cause the rover to shake as it moves (see Figure 5).

Figure 5. Typical view of the Playa location at Silver Dry Lake.

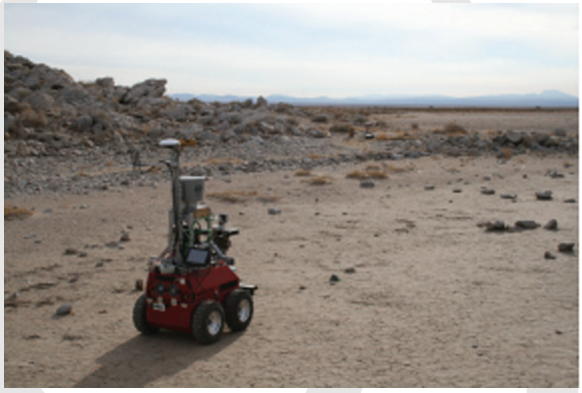

Figure 6. Typical view of the Boulder location at Silver Dry Lake.

The Boulder field is characterized by relatively flat loose sand and playa, with large rock hazards randomly scattered throughout (see Figure 6). The hazard detection function of the terrain assessment system was fully exercised in this location.

The Plateau area is characterized by loose, small-grained gravel with few rocks, but many medium-sized bushes (see Figure 7 and Figure 8). This proved to be a visually challenging location, because there were few rock features of significant size in the image, and the wind-induced motion of the bushes caused their appearance to change rapidly, which affects the pixel-tracking motion estimate.

3.1.2. Silurian Dry Lake Test Locations

Two locations were selected for testing in the Silurian Dry Lake bed: the Playa and the Shore. The open playa of Silurian Lake is virtually the same as the playa of Silver Lake, with the exception that the fissures in the clay are slightly larger. The Shore location is simply the shore of the Silurian Dry Lake bed, where the playa turns into loose gravel with some bushes (see Figure 9). The advantage of this test area

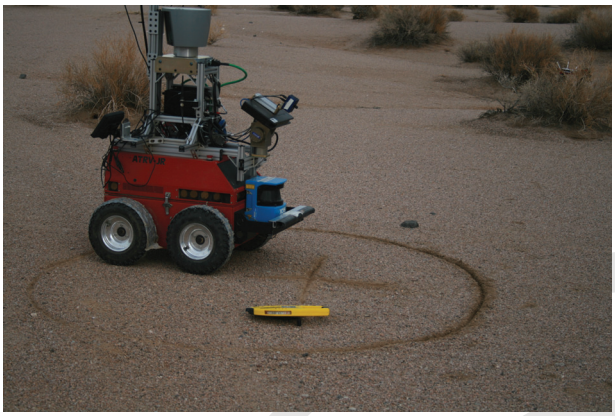

Figure 7. Typical view of the Plateau location at Silver Dry Lake.

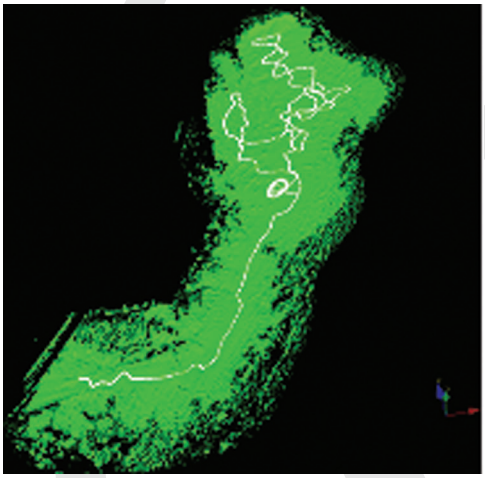

Figure 8. Reconstructed map of the Plateau and the traveled path (200 × 255 m, 960 m traveled distance).

was the ability to test the transition from playa to loose gravel and back.

3.1.3. Other Test Locations

The Mudflat area was found near Route 127, between Silurian Dry Lake and Dumont Little Dunes. It contains a lot of loose sand and gravel, with many ravines and gullies. There are only a few bushes. Most of the hazards are due to slope rather than rocks (see Figure 10 and Figure 11).

Dumont Little Dunes provided a test location with loose sand dunes, with occasional sparse bushes, as shown in Figure 12 and Figure 13. This terrain provided a great deal of wheel slippage. Visual features were mostly due to vehicle tracks rather than natural terrain.

Figure 9. Typical view of the Shore location at Silurian Dry Lake.

Figure 10. Typical view of the Mudflat location near Route 127, between Silurian Dry Lake and Dumont Little Dunes.

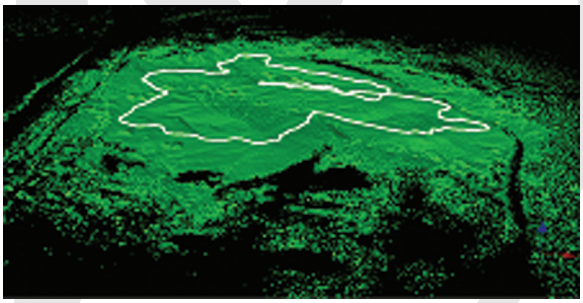

Figure 11. Reconstructed Mudflat map and traveled path overlays (160 × 140 m, 445 m traveled distance).

3.2. The Challenges and Issues Observed

This subsection addresses some of the general issues observed that affected the field campaign. The performance of the localization system itself will be addressed in Section 4.

Figure 12. Typical view of the Dumont Little Dunes location.

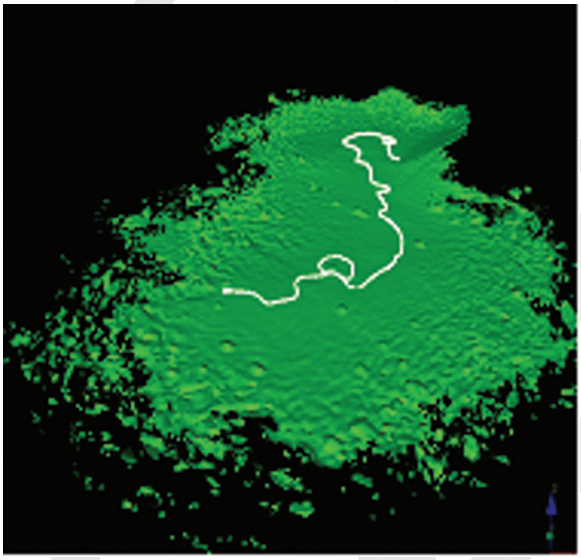

Figure 13. Reconstructed Dunes map and traveled path overlays (250 × 340 m, 360 m traveled distance).

General issues during the field campaign included the following:

- *Weather:* January is the rainy season in the Mojave Desert. Unfortunately, in 2010 it rained more often and more heavily than usual. As a result, a few full days of testing were completely lost, and about 3 days had to be cut short because of rain. Rainfall could be heavy at times, flooding the dry lake beds. With any more than a light sprinkle, the playa becomes very slick and loose, quickly causing risk to the vehicles. Further, the ground needs time to dry out after a heavy rainfall, meaning that more than 1 day can be lost at a time.
- *Glare:* Specular reflections from the sunlight on the moistened, high-reflectivity sand frequently caused glare in the VME stereo camera images, which had a negative impact on the performance of VME localization. For a more complete discussion, see Section 4.2.
- *Rover Terrain Ability:* The terrain ability of the rover (shown in Figure 1) is very limited. Having no suspension and limited ground clearance means that very small rocks can be hazardous. Although the existing sensor suite certainly has the resolution to detect hazards for a planetary-representative chassis, with the current chassis there are still some undetected rocks that, although not catastrophic, can cause the rover to stick and skid, or undergo abrupt pitch and roll motions.
- *Shadows in the Images:* No effort was made to control the direction of travel with respect to the sun, and there are many instances in which the shadow of the rover appears in the stereo camera images. This ensures that the motion filtering in the feature-matching algorithms is fully exercised.

4. EXPERIMENTAL RESULTS

4.1. Performance Metrics and Statistics

Path planning, autonomy, and localization are key capabilities for planetary rovers, and thus, the field trial looked to quantify the performance of the AIR-GNC system in these areas. Success of the terrain assessment and planning function is defined as finding a safe path to the goal where one exists, and correctly identifying when a safe path does not exist. Success of the full autonomous system is defined as reaching the goal (or the closest traversable point to the goal) when a safe path exists. The distance between the end point of the localization estimate and the end point of the ground truth, divided by the total path length estimated from the ground truth, is used to express the localization accuracy as a percentage of the distance travelled.

The ground truth position was measured using a NavCom 3020M real-time kinematic differential global positioning system (RTK DGPS). The manufacturer's specification states a precision of 1 cm in RTK mode (RT-3020 GPS Products User Guide, 2008). GPS latitude, longitude, and altitude data were converted to the local frame, where positive X is East, positive Y is North, and positive Z is up (i.e., coincident with the local gravity direction).

Localization data were logged with respect to the initial rover body frame at the start of the traverse. Thus, the ground truth must be expressed in the initial rover frame before accuracy analyses can begin. The position offset is easy to correct by simply subtracting the initial GPS position from all subsequent measurements. Pitch and roll are also compensated for, because the VME initializes itself using the inclinometry from the IMU's accelerometers, which give very clean measurements of the gravity direction because the rover is stationary. However, in the absence of an absolute heading sensor on the rover, the alignment of the ground truth becomes somewhat more challenging. Here

Table I. Traverse distances and mean localization errors, broken down by test area.

		Traverse distance (m)			Mean localization error (%)			
Location	# runs	Min dist.	Mean dist.	Max dist.	IMU-odom.	H-L	MSER	SIFT
Boulder	5	12.8	32.1	48.2	2.9	8.7	3.2	4.2
Dunes	4	12.7	246.6	599.5	3.5	7.7	2.5	4.8
Mudflat	6	20.9	41.2	84.3	6.3	7.3	7.0	6.5
Plateau	6	34.7	174.7	867.8	1.7	4.6	3.5	6.0
Playa	5	13.9	179.1	316.6	3.3	9.8	3.0	25.6
Shore	2	16.4	68.4	120.4	2.1	4.2	2.3	2.5
All locations	28	12.7	124.1	867.8	3.5	7.3	3.9	8.8

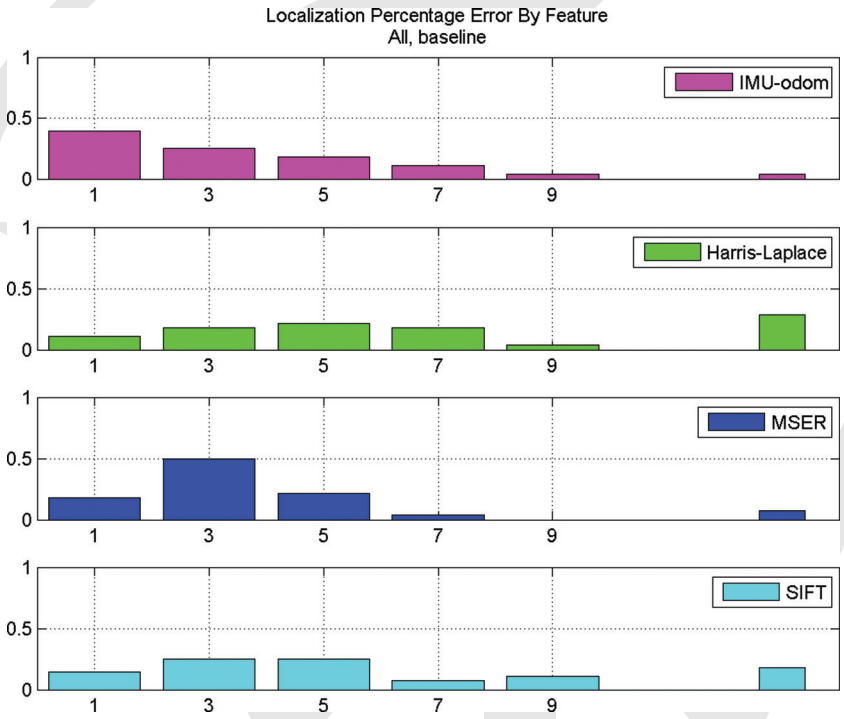

Figure 14. Localization percentage error by feature, cumulated over all test sites.

the localization estimate and ground truth were aligned using the first 5 m of the traverse.

In this project, with so many different options for the VME algorithm, it was more desirable to align the ground truth with the IMU-corrected odometry. Note that no matter which signal is used for the alignment, the accuracy results will be biased (even if very slightly) in favor of that signal over the others. In the absence of absolute attitude ground truth, this problem cannot be avoided.

4.2. Results and Discussion

Of the 104 total runs recorded, 28 were focused on testing the VME system. VME localization accuracy results are differentiated based on the visual feature used for localization, and as well as the location of the test (Table I). Figure 14 shows the localization percentage error by feature based on the data collected in all field test locations. For clarity, the histograms are normalized by the number of runs. The width of each bin of the histogram is 2% (of distance traveled), and for clarity, the axis is cut off at 10% error. All runs with errors greater than 10% (the likely outliers) are lumped into a single histogram bin, which is presented to the right of the regular histogram.

As can be seen in Figure 14, on the average MSER outperformed both Harris–Laplace and SIFT. This suggests that in these terrains the image regions had greater stability than the local corner features. On the average, enhanced

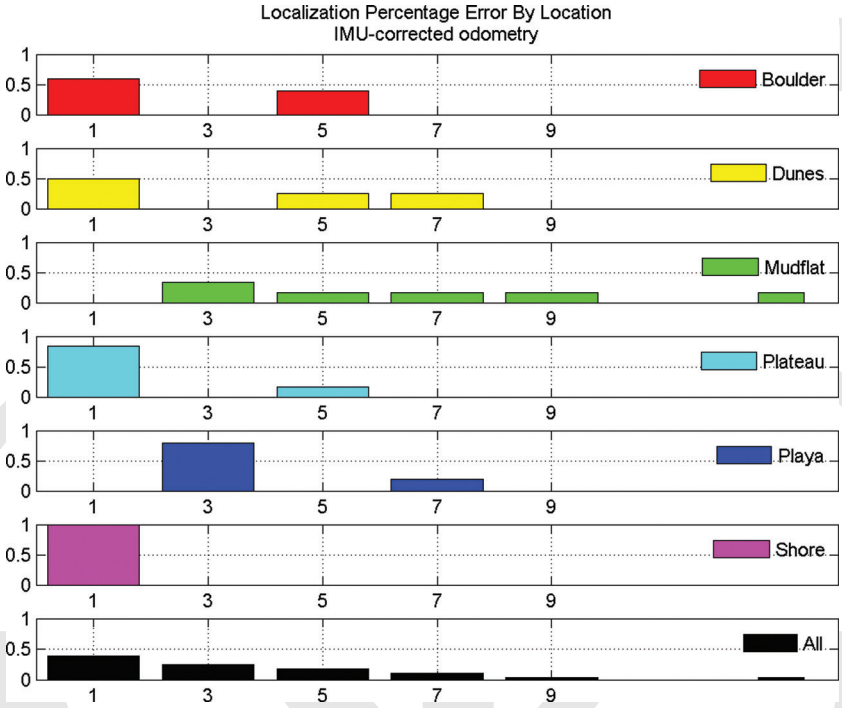

Figure 15. Enhanced IMU-corrected odometry localization percentage error across all test locations.

IMU-corrected odometry had slightly better accuracy than MSER. Figure 15 and Figure 16 show the normalized localization results of respectively the enhanced IMU-corrected odometry and MSER, differentiated across the various test locations. From these figures and the numerical results in Table I, it can be seen that MSER outperformed enhanced IMU-corrected odometry in the Playa area and the Dunes area, the latter certainly because of the large amount of wheel slippage.

The Mudflat location was particularly difficult for every localization method, including enhanced IMU-corrected odometry. The terrain was quite loose, which caused considerable slippage. VME was primarily affected by the high reflectivity of the terrain (for example, the glare that was present in many runs). Moreover, the stereo cameras used in this project for VME have automatic brightness and contrast adjustment capability. One of the lessons learned from this field trial is that automatic exposure adjustment may actually hinder the performance of the stereo camera vision system. In many runs there is a significant and sometimes intense glare from the sunlight, which washes out large portions of the image and increases the contrast. Example runs in which the glare is very prominent include those on the Playa and in the Mudflat location before the sun goes behind the mountains (see Figure 17). Further, this glare is specular, so it changes with the relative angle to the sun. Because of this, the glare patches themselves do not make very reliable features, as they change with viewpoint. This effect was not noticed during the previous field campaign in 2008, which took place during the dry season (Bakambu et al., 2008), and is not considered representative of Martian or Lunar surface conditions.

The autonomous system performed with a high success rate. Of the 104 total runs recorded, 93 were executed autonomously. In 12 cases, the rover did not reach its goal location because of a GNC-related issue:

- In four runs, the rover stopped at an intermediate waypoint with a hazard within the blind zone of the terrain assessment sensors. When the next intermediate traverse was executed, the rover moved into the hazard.
- In four runs, the terrain assessment algorithm was unable to distinguish between sparseness of data points due to field of view and range limitations of the sensor due to hazardous terrain (e.g., a gulley). The algorithm acted conservatively and classified the cell as hazardous, and no path to the goal was found. This issue has since been corrected.
- In four runs, the accumulated localization error caused the rover to approach a hazard.

In total, the success rate of the planner was 91%, and the success rate of the fully autonomous system (i.e., reaching the selected goal) was 87%.

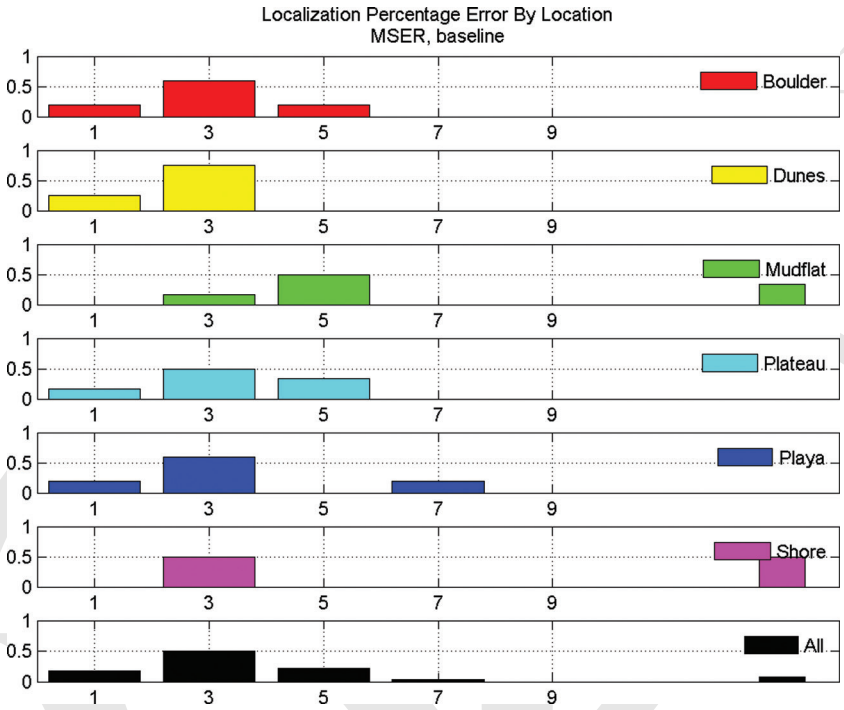

Figure 16. MSER localization percentage error across all test locations.

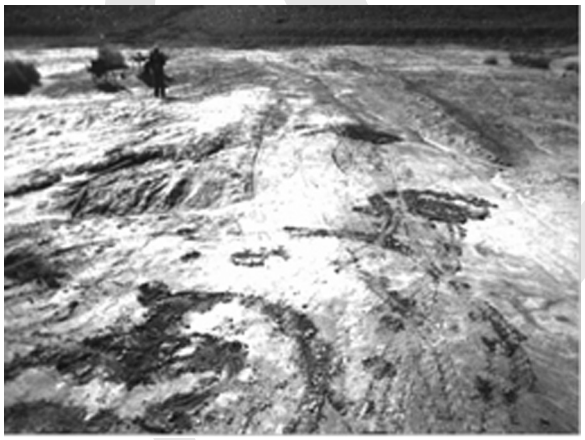

Figure 17. Mudflat terrain with high contrast due to glare off of the moist terrain.

Figure 18. Image of the terrain in Mudflat location: The goal location is shown in yellow behind the bush.

Several challenging scenarios were constructed to test the decision-making ability of CORTEX-based high-level supervisory control. One example (not shown) placed the rover within a ring of hazards, with the only escape route within an occluded region of the sensor; the rover was able to move and replan until it could detect the escape route, and then proceed to the goal. In another scenario, shown in Figure 18, the goal (denoted by the yellow object) was selected in an untraversable location behind a bushy hill. To reach the goal, the robot had to follow the ditch and move around the hill. Figure 19 shows the path planned and executed by the robot. As can be seen, the robot did exactly what was expected.

4.3. Summary of the Results

Table II summarizes the localization accuracy of stereo camera feature detectors shown in Figure 14.

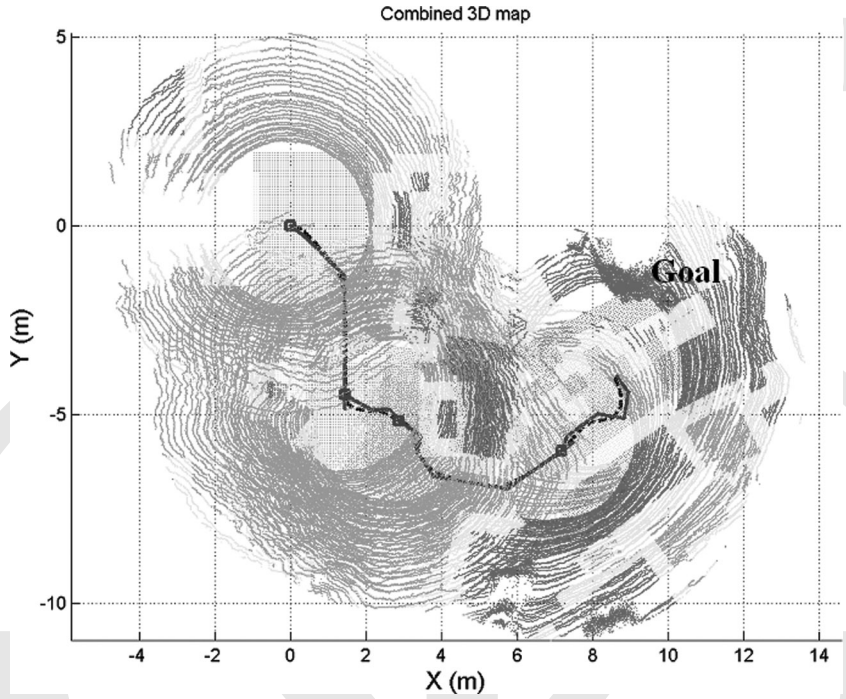

Figure 19. Combined 3D cost map and path overlays. Note that the goal location is untraversable, so the rover stops at the nearest traversable location.

Table II. Localization accuracy summary for stereo camera feature detectors.

Feature	Min % error	Mean % error	Std. dev. (%)
Harris–Laplace	0.8	7.3	5.2
MSER	0.4	3.9	2.8
SIFT	1.2	8.8	11.9
IMU–odom.	0.4	3.5	2.2

MSER had better accuracy and a lower standard deviation than either Harris–Laplace or SIFT. This suggests that in these types of terrain, region-based features are more useful than the local (corner) features. This is particularly noticeable in the Playa areas. MDA enhanced IMU-corrected odometry did very well in most locations—on the average, slightly better than MSER. MSER outperformed enhanced IMU-corrected odometry in the Playa and Dunes areas. Harris–Laplace performed nearly as well as enhanced IMU-corrected odometry in the Shore area (the plot is not shown). The Mudflat location was particularly difficult for all methods of localization. As mentioned above, enhanced IMU-corrected odometry was subject to slippage, and camera images were subject to glare.

Based on the above results, an elegant localization solution would be to rely on enhanced IMU-corrected odometry and continuous monitoring for wheel slippage. If the slippage starts, the localization will switch from enhanced IMU-corrected odometry to VME and stay in this mode until the slippage ends. Another solution is to robustly extract and match long-range landmarks online to enforce the observability of the estimator. The concept behind the approach is to localize the landmarks to a fixed frame (for example, the initial starting pose of the rover), rather than as part of the SLAM state. Observability of the estimator is maintained as long as the landmarks are within the field of view and operating range of the sensor. Offline results (see Table III) using data collected during the field trials have shown very promising results when compared against the IMU–odometry and MSER-based estimates.

5. CONCLUSIONS AND FUTURE WORK

This paper has presented the performance of the VME in the field trials of the AIR-GNC project conducted by MDA in collaboration with and funded by the CSA.

The operational scenario of AIR-GNC includes terrain scanning, modeling and traversability assessment, and path planning and tracking. CORTEX-based supervisory control was effective for monitoring and executing traverses, which resulted in a noticeable increase in the autonomy of the test bed. VME-based feedback was successfully

Table III. Results of offline long-range landmark-based observable VME.

Location	Distance (m)	IMU-odom. (%)	MSER VME (%)	3D landmark VME (%)
Mudflat, run 1	85	2.0	3.0	0.6
Mudflat, run 2	21	3.0	9.4	0.9
Plateau	34	1.5	5.3	1.3

used to close the path-tracking loop. A total of 3.2 km of autonomous traverses was executed.

Enhanced IMU-corrected odometry performed very well; on the average it performed slightly better than MSER-based VME. MSER outperformed Harris–Laplace and SIFT as a feature-extraction front end to VME.

Future work will include wheel slippage monitoring and estimation using visual evidence. Improved odometry would be a wide-ranging benefit. This could potentially involve better state estimation using a dynamic model of the rover. Robust long-range 3D feature extraction will be critical if long-range observation is to be used in practice.

ACKNOWLEDGMENTS

This work was funded by the Canadian Space Agency. The authors would like to thank Régent L'Archevêque and David Gingras from the CSA for their valuable contributions in path planning and CORTEX-based supervisory control. The authors also thank Manickam Umasuthan from MDA for his contributions to the 3D feature-extraction implementation, and Tung Nguyen and Kelvin Tao, students at the University of Toronto, for integration of sensors and initial testing of the system. Final thanks go to Rob Fulton, Jason Wallace, and the staff of the Desert Studies Center, who were extremely accommodating during our field trials.

REFERENCES

Bakambu, J., et al. (2008). Visual motion estimation: localization performance evaluation tool for planetary rovers. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Los Angeles.

Biesiadecki, J., et al. (2005). Mars exploration rover surface operations: Driving opportunity at Meridiani Planum. In IEEE Conference on Systems, Man and Cybernetics, The Big Island, HI.

Corke, P., et al. (2004). Omnidirectional visual odometry for a planetary rover. In Proceedings of Intelligent Robots and Systems.

Davis, M. H. A., & Vinter, R. B. (1985). Stochastic modelling and control. London: Chapman and Hall.

Donoser, M., & Bischof, H. (2006) Efficient maximally stable extremal region tracking. In IEEE Conference on Computer Vision and Pattern Recognition.

Dupuis, E., et al. (2005). Autonomous robotics toolbox. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Munich, Germany.

Furgale, P., & Barfoot, T. (2010). Visual teach and repeat for long-range rover autonomy. Journal of Field Robotics 27, 534–560.

Gingras, D., Dupuis, E., Payre, G., & De Lafontaine, J. (2010). In Path planning based on fluid mechanics for mobile robots using unstructured terrain models. IEEE International Conference on Robotics and Automation, Anchorage, AK.

Gingras, D., Lamarche, T., Dupuis, E., & Bedwani, J. L. (2010). Rough terrain reconstruction for rover motion planning. Paper presented at the Seventh Canadian Conference on Computer and Robotics Vision (CRV), Ottawa, Canada.

Harris, C., & Stephens, M. (1988) A combined corner and edge detector. In Proceedings Fourth Alvey Vision Conference, Manchester, UK (pp. 147–151).

Huntsberger, T., et al. (2002). Rover autonomy for long range navigation and science data acquisition on planetary surfaces. In IEEE International Conference on Robotics and Automation.

Konolige, K., Agrawal, M., & Sola, J. (2007). Large scale visual odometry for rough terrain. In Proceedings of the International Symposium on Research in Robotics (ISRR).

Kristensen, F., & MacLean, W. J. (2007). Real-time extraction of maximally stable extremal regions on an FPGA. In IEEE International Symposium on Circuits and Systems.

Lacroix, S., et al. (2000). Autonomous rover navigation on unknown terrains. In 7th International Symposium on Experimental Robotics, Honolulu, HI.

Lacroix, S., et al. (2002). Autonomous rover navigation on unknown terrains. International Journal of Robotics Research, 21(10–11), 917–942.

Lowe, D. G. (1999). Object recognition from local scale-invariant features. In Proceedings of the International Conference on Computer Vision, Corfu.

Maimone, M., Chang, Y., & Matthies, L. (2007). Two years of visual odometry on the Mars exploration rovers. Journal of Field Robotics, 3(24), 169–286.

Matas, J., Chum, O., Urban, M., & Pajdla, T. (2002) Robust wide baseline stereo from maximally stable extremal regions. Paper presented at the British Machine Vision Conference.

Mikolajczyk, K., et al. (2005). A comparison of affine region detectors. International Journal of Computer Vision, 65(1–2): 43–72.

Ramadge, P. J., & Wonham, W. M. (1987). Supervisory control of a class of discrete event processes. SIAM Journal of Control and Optimization, 25(1), 206–230.

Ramadge, P. J., & Wonham, W. M. (1989). The control of discrete event systems. In Special issue on discrete event dynamic systems, Proceedings of the IEEE, 77(1), 81–98.

RT-3020 GPS Products User Guide (2008). NavCom Technologies, Inc. **Q20**

Se, S., Barfoot, T., & Jasiobedzki, P. (2005). Visual motion estimation and terrain modeling for planetary rovers. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Munich, Germany. **Q21**

Sibley, G., Mei, C., Reid, I. D., & Newman, P. N. (2010). Vast-scale outdoor navigation using adaptive relative bundle adjustment. International Journal of Robotic Research, 29(8). **Q22**

Souvannavong, F., Lemaréchal, C., Rastel, L., & Maurette, M. (2010). Vision-based motion estimation for the ExoMars rover. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Sapporo, Japan. **Q23**

Stentz, A. (1995). The focused D* algorithm for real-time replanning. Paper presented at the 14th International Joint Conference on Artificial Intelligence, Montreal, Canada. **Q24**

Triggs, B., Mclauchlan, P., Hartley, R., & Fitzgibbon, A. (2000) Bundle adjustment—A modern synthesis. In Vision algorithms: Theory and practice, Lecture notes in computer science (pp. 298-375). Springer-Verlag. **Q25**

Volpe, R. (2006). Rover functional autonomy development for the Mars mobile science laboratory. In Proceedings of the IEEE Aerospace Conference, Big Sky, MT (Vol. 2, pp. 643–652). **Q26**

Wettergreen, D., et al. (2005). Second experiments in the robotic investigation of life in the Atacama desert of Chile. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), 2005. **Q27**

Wettergreen, D., et al. (2008). Long-distance autonomous survey and mapping in the robotic investigation of life in the Atacama Desert. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), 2008. **Q28**

Queries

Q1: Figures 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 18, 19 are in color in the source files. Figures appear in color online at no cost to you. If you would like any figure to appear in color in the print issue, please advise and we will send you a formal quote for the cost. Otherwise they will appear in black and white in the print issue and in color online.

Q2: If desired, one proof page may appear in color in the print issue at no charge. Please indicate the page you would like printed in color, if any. Any additional figures can be in color for a charge of $850 per page.

Q3: AU: To identify which specific works are being referred to, cite as Gingras, Lamarche, et al., 2010 and as Gingras, Dupuis, et al., 2010, respectively.

Q4: AU: Give names of all authors, month of conference.

Q5: AU: Give names of all authors, month and city of conference.

Q6: AU: Give names of all authors. If journal, give volume, page range; if conference proceedings, give month and city of conference.

Q7: AU: Give month and city of conference.

Q8: AU: Give names of all authors, month of conference.

Q9: AU: Give month of conference.

Q10: AU: Give month of conference.

Q11: AU: Give month of conference.

Q12: AU: Give names of all authors, month of conference.

Q13: AU: Give month of conference.

Q14: AU: Give month and city of conference.

Q15: AU: Give names of all authors, month of conference.

Q16: AU: Give names of all authors.

Q17: AU: Give month of conference.

Q18: AU: Give month and city of conference.

Q19: AU: Give names of all authors.

Q20: AU: Give city where published.

Q21: AU: Give month of conference.

Q22: AU: Give page range.

Q23: AU: Give month of conference.

Q24: AU: Give month of conference.

Q25: AU: Mclauchlan, or McLauchlan? Give city where published.

Q26: AU: Give month of conference.

Q27: AU: Give names of all authors, month of conference.

Q28: AU: Give names of all authors, month of conference.