MURFI 2016 - FROM CARS TO MARS: APPLYING AUTONOMOUS VEHICLE NAVIGATION METHODS TO A SPACE ROVER MISSION

Brian Yeomans[1], Horia Porav[1], Matthew Gadd[1], Dan Barnes[1], Julie Dequaire[1], Tom Wilcox[1], Stephen Kyberd[1], Simon Venn[1], and Paul Newman[1]

[1]*Oxford Robotics Institute, University of Oxford, United Kingdom*
{brian,horia,mattgadd,dbarnes,julie,tomw,skyberd,svenn,pnewman}@robots.ox.ac.uk

ABSTRACT

The Mars Utah Rover Field Investigation (MURFI) 2016 mission was a Mars Rover field analogue mission run by the UK Space Agency (UKSA) in collaboration with the Canadian Space Agency (CSA). MURFI consisted of a "field team" at the Mars analogue site near Hanksville (Utah, USA), and an "Operations Team" based in the Mission Operations Centre (MOC) at UKSA's Harwell Campus (Oxfordshire, UK).

The Rover platform for the mission comprised the Oxford Robotics Institute's (ORI) ARC Q14 Space Rover (Q14), and the mission provided a unique opportunity for ORI to test the performance of several of their advanced navigation and autonomous driving algorithms in realistic planetary exploration conditions over a period of several weeks (Figure 1). MURFI's core objectives included a realistic imitation of the first 10 sols of a possible future Mars Rover mission.

Following completion of the core mission, the ORI Rover team engaged in a series of ambitious trials and data gathering scenarios based on ORI's broad suite of navigation and autonomy algorithms which have been developed primarily for on-road terrestrial applications. The objectives were to (i) assess the performance of these systems in this radically different environment, (ii) determine what systems modifications (if any) might be needed to operate effectively in a typical planetary surface scenario, and (iii) to identify which new capabilities these techniques might bring to the field of planetary surface navigation and exploration. The results of these trials were very encouraging, showing good baseline performance of the techniques deployed even though these had not been modified in any way to reflect the change of environment, and indicating several avenues for further exploration and development of ORI's techniques which could generate substantial benefits for the Space community.

During the core mission, the Rover drives for each sol were conducted autonomously using a single mast-mounted stereo camera sensor. Waypoint files for the planned drives were designed by the MOC, transmitted to the field site, and uploaded to Q14 for autonomous execution using ORI's OxfordVO visual odometry (VO) application - the same application which forms the kernel of the VO software for the forthcoming ExoMars mission. The accuracy in following the planned drives during this phase was highly satisfactory, with the Rover typically finishing the drive within a few centimetres of the planned location.

During the post-MURFI trials programme, the team implemented ORI's "Dub4" suite to provide "teach-and-repeat" (T&R) functionality. Dub4 uses a single stereo camera to create vision-based maps in highly unstructured environments which are then used to localise and navigate autonomously. This facilitates safe, rapid retracing of a previously driven path, enabling interesting science sites subsequently identified through data analysis to be rapidly revisited by the Rover. Accuracy in following the previously driven path was good despite large areas of the driven environment being relatively featureless. The use of an affordable monocular camera as an effective localisation sensor using the feature-rich desert floor was investigated, with encouraging results. Data was collected so that dense reconstructions of the terrain around the rover can be generated in a future phase of our work. This reconstruction capability has the potential to create an extremely powerful visualisation tool for generating rich 3-D mesh representations to be utilised by mission scientists to more effectively focus the science effort. Future work will also concentrate on performance enhancements by adapting existing ORI techniques to the specifics of the planetary surface environment, and developing enhanced machine learning autonomy approaches along the path towards the implementation of a true "robot geologist".

Key words: robots; rover; localisation; mapping; autonomy; teach-and-repeat.

1. INTRODUCTION

The MURFI mission was a rover field mission intended to demonstrate the capability of the UKSA, in collaboration with international partners including the CSA, to

Figure 1. The Q14 Rover and PANCAM emulator. The PanCam is mounted at the top of the mast and can be rotated through 360 *degrees and angled up/down using a mast mounted pan/tilt unit. The XB3 is the gold unit half way up the mast, the LIDAR is just below, and and the Grasshopper is mounted on the front of the rover platform bed.*

deploy an "ExoMars-like" rover analogue mission and carry out a geologically-focused traverse of Mars-like terrain. As part of the mission, ORI provided the Rover and rover operations support team to the field location in the Utah desert near Hanksville, as well as support to operations in the MOC located at Harwell, UK. The mission had a number of scientific objectives and activities which are described in more detail in [1]. In this paper, we focus on rover operations carried out during both the MURFI mission proper, and the subsequent post-mission phase. This post-mission phase in particular provided an excellent opportunity to deploy ORI's suite of navigation and autonomous driving techniques in a realistic planetary analogue context.

2. THE ROVER

Figure 1 shows an image of the rover including the PanCam and navigation instruments.

2.1. Payload platform

The rover platform comprised a Q14 robot from Advanced Robotics Concepts (ARC). With active 4-wheel steering and drive, and passive dynamic suspension system, the rover provides a reasonable payload capacity and good mobility over a range of terrains within a relatively low mass package. The rover mass without payload is approximately 30kg and it can carry up to 40kg of payload. For the MURFI mission activities, the 4-wheel steering capability enabled the path planning exercise to be simplified through construction of the paths as a series of linear drives linked by point turns. 4-wheel steering also means that wheel-slip is much reduced compared with simpler differential steering platforms, reducing the impact of the rover on the terrain and minimising track deposition.

2.2. Navigation sensors

As seen in Figure 1, the primary navigation sensor comprised a Point Grey Bumblebee XB3 stereo camera mounted mid-way up a central mast fitted to the rover. This was the sole sensor used for navigation during the MURFI mission proper. In addition, the rover was fitted with a Velodyne VLP-16 3-D LIDAR located just below the XB3 camera and a Point Grey Grasshopper wide-angle monocular camera facing forwards and looking past the forward wheel set towards the terrain in front of the rover. The LIDAR was used to gather data for the forthcoming dense reconstruction exercises, and the Grasshopper camera for single-camera navigation. The platform was also fitted with a Lord Microstrain 3-DM-GX4-45 inertial sensor, which was primarily utilised for automatic logging of the platform orientation during imaging sessions. All data from every sensor was routinely logged for subsequent analysis and evaluation.

2.3. Instrumentation

The overall instrumentation package was similar to that planned for the ExoMars mission and comprised:

- The Aberystwyth University PanCam Emulator (AUPE) [2] to simulate the ExoMars PanCam instrument [3], and the High Resolution Camera (HRC) emulator,
- A Digital SLR camera with macro lens, mounted to simulate the ExoMars Close-up Imager (CLUPI [4]) range of motion and field-of-view,
- An ASD Inc. FieldSpec4 field reflectance spectrometer to simulate the Infrared Spectrometer for Mars Instrument (ISEM, [5]), and
- A Raman Spectrometer, the use of which on the final drill-samples acquired would signify "mission success".

3. SURFACE NAVIGATION

3.1. MOC communication protocol

During the MURFI mission, the rover operated in limited autonomy mode. The planned drive for the upcoming sol cycle was dictated by the MOC team at Harwell, based primarily on PanCam imagery taken during the drive on the previous sol and uploaded from the rover to the MOC at the end of the drive. Each drive was planned in coordinates relative to the start point of the drive, corresponding to the end point of the previous sol's drive. The rover position and orientation was marked physically at the end of each sol, the rover removed, and replaced on the following sol in the exact same position ready for the drive. The planned drive was described simply as the (x, y) coordinates in metres relative to the start position of a series of waypoints linked by point turns. The orientation at each point turn was specified, including the goal orientation. A text descriptor file was generated by the MOC, sent to the field team, and uploaded to the rover.

3.2. Executing sol drives

When driving, the rover operated autonomously. To ensure the rover actually drove the planned track, the rover utilised the XB3 stereo camera linked to OxfordVO [6], which generated frame-by-frame estimates of Q14's ego-motion. This application has been selected as the VO component for the forthcoming ExoMars mission, and is described in more detail in [7].

Using a simple differential controller which calculated the difference between the rover's actual and planned pose, corrections were generated to ensure the rover adhered to the planned path. The system performance was found to be excellent, with the actual position at the end of each sol's drive measured to be within just a few centimetres of the planned position, and no problems were experienced in meeting the MURFI mission rover drive requirements.

4. TEACH-AND-REPEAT WITH DUB4

"Wherever, Whenever, Whatever the Weather" (Dub4) is ORI's state-of-the-art visual navigation system. In this section we describe an approach where the rover is first manually driven along a path ("teach") and then subsequently drives that same path autonomously ("repeat"). To do this, Dub4 consumes live images from a stereo camera and compares them to a database of previously recorded locations to determine the precise position of the rover with respect to the "taught" path [8][9][10].

4.1. Mapping

The database is created during a first survey of the rover's route (an experience) by storing both visual snapshots of the places that it sees and the relative coordinates (rotation, translation) between these snapshots. As the rover undertakes more drives, many more experiences of either the same route or of different routes can be added and fused together seamlessly in order to improve the localisation performance of the rover, and this accumulated database of information can then be used to localise the rover under a diversity of viewpoint, illumination, and weather conditions. The rover finds its position along the route by first determining which visual snapshot (keyframe) matches best with what is currently observed, and then computing a metric 6DoF pose relative to that snapshot. The rover can then compute its relative position against any other snapshot recorded in the database by examining the chain of stored poses between snapshots [8][9][10].

First, the rover's trajectory is estimated using VO [11][12]: the current stereo image pair F_t is used to create and store a bank of 3-D landmarks (SURF, BRIEF, ORB), and a 6-Degree-of-Freedom (6DoF) transformation is computed between F_t and the previous stereo pair, F_{t-1}. The 3-D landmarks and 6DoF poses between them are stored in an "experience" graph $\mathcal{G}$, with nodes composed of 3-D landmarks and the 6DoF transforms saved as edges between nodes (see Figure 2), under the assumption that 6DoF transforms between nodes that are close will be metrically accurate, while nodes that are far away from each other will maintain only a topological ordering [13][14][8]. At localisation time, landmarks extracted from F_t are compared against the stored 3-D landmarks, and a 6DoF transform is computed between F_t and the experience frame F_{exp_i} that best matches the current frame.

4.2. Localisation

During localisation, the rover's pose is encoded as a reference node N_i and a 6DoF transform T_{N_i,R_t} between the rover and this reference node. Localisation is a two-step process: Firstly, large-scale localisation (so called place-recognition) is done using FABMAP [15], which detects loop-closures using image similarity and yields a set of candidate nodes N_i that are most similar to the live frame, without actually computing a pose relative to any of those nodes. Secondly, local-scale metric localisation is performed on these candidate nodes, by solving for a 6DoF transform between each candidate node and the live frame. Finally, Dub4 will output a 6DoF transform with respect to the node that has the most inliers(correct landmark correspondences between the live frame and the keyframe stored on the node).

4.3. Path memory

However, the number of nodes that need to be checked by FABMAP will grow with each added experience, which might become problematic, especially when running on resource-constrained hardware such as a rover. Given that localisation must run in real-time or near-real-time, a naïve approach would only be able to check a reduced number of experience nodes before the next live frame is fed into the system. To increase the probability of localisation success under these constraints, Dub4 uses a ranking policy [8] which sorts the candidate experience nodes by their distance to the current position estimate obtained using VO. Additionally, Dub4 ranks candidate nodes based on which experience is currently used for localisation, as it is likely that the same experience will be used to localise in future frames. This probabilistic approach is further explained in [8] and [16].

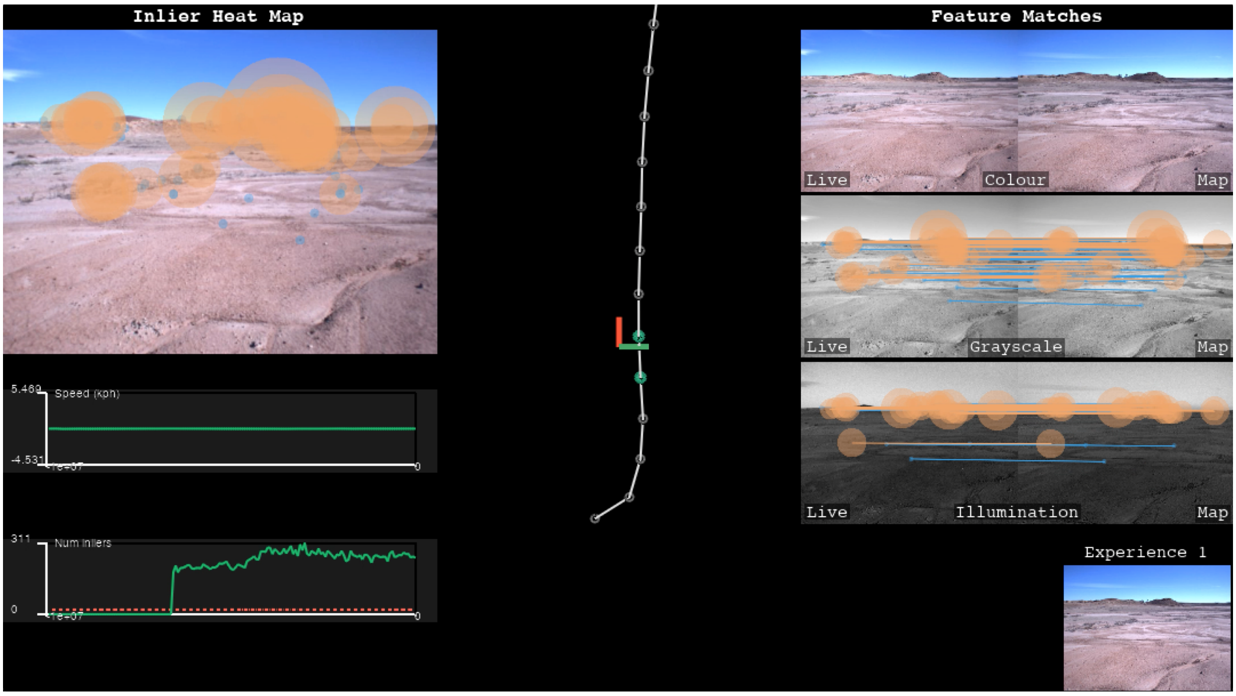

Figure 2. DUB4 running on the Q14 rover. The lower left side shows the number of inliers (correct image correspondences between the current frame and an experience keyframe). The right side of the figure shows the 3D landmark correspondences between the current live frame and the experience keyframe, for grayscale and illumination-invariant transforms of the original RGB images. The center of the figure shows the graph-like representation of a Dub4 experience, where nodes represent keyframes, and edges represent 6DoF transforms between keyframes.

4.4. Illumination invariance

There are also instances where changes in illumination result in localisation failure, especially in the case of moving shadows. To bolster the robustness to such changes, Dub4 uses an illumination invariant transform of the input color images, similar to the method described in [17].

4.5. Path-following

We have tested the applicability of Dub4 to Mars-like environments by simulating a full autonomy process. This is a three-step exercise:

- The rover is manually driven around a predetermined path several times, while recording full video logs from the stereo camera mounted on its mast. The rover operators attempt to follow exactly the same route for each subsequent drive by using the tracks left by the rover as a guide,
- Afterwards, an experience map is created using the video log of the first manual drive,
- Finally, localisation and autonomy are simulated, by using the additional video logs as simulated stereo video input to Dub4.

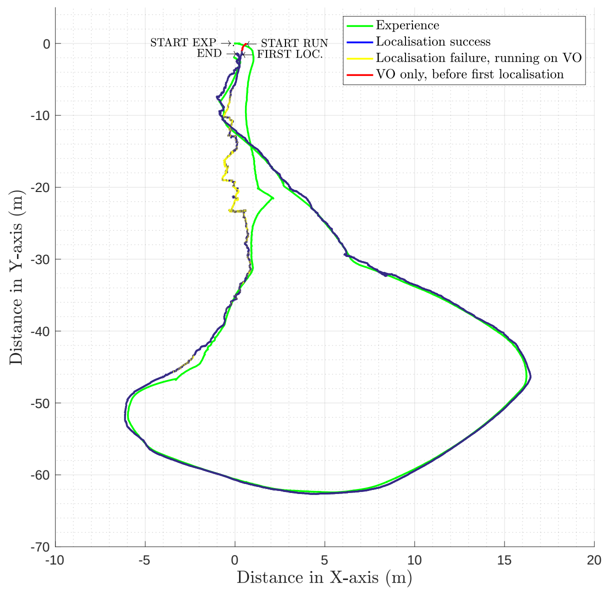

Figure 3. Path driven by the rover during the first simulated autonomous run, with the experience map in green, and the first simulated autonomous run in yellow, blue and red.

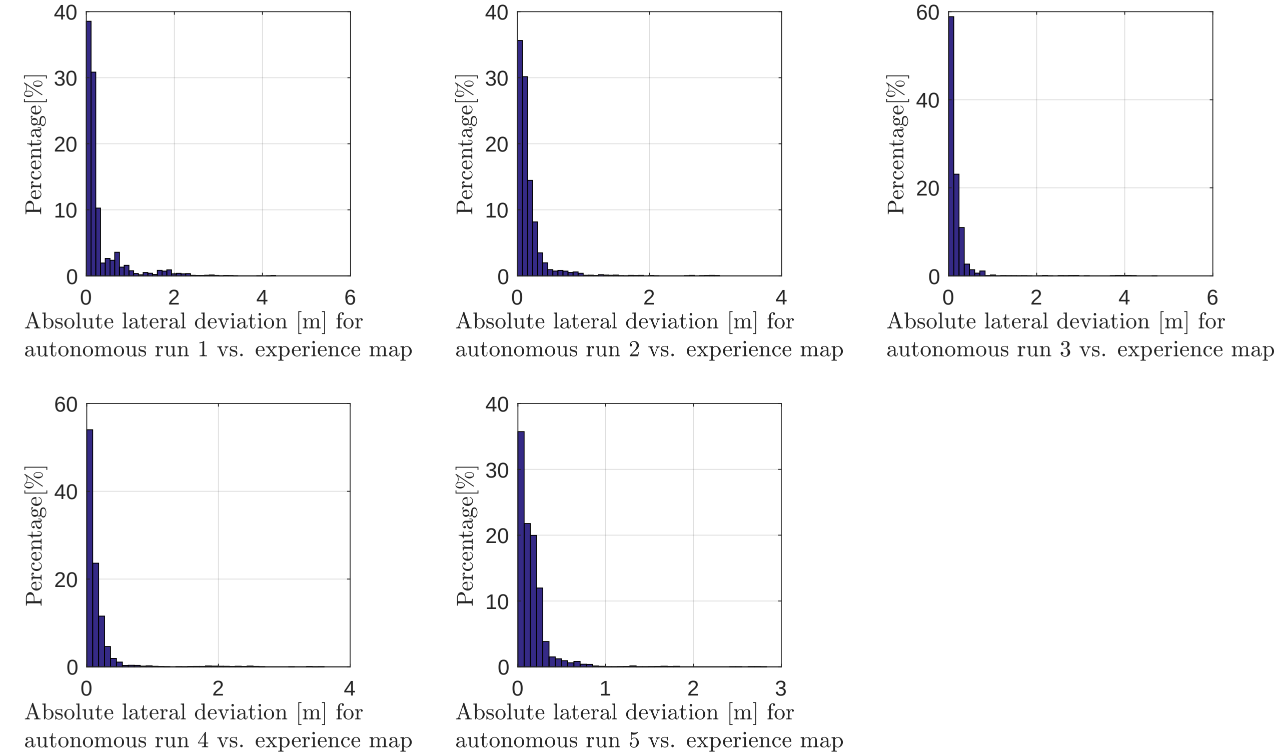

Figure 4. Lateral deviation histograms of the simulated autonomous runs from the reference experience map (not cumulative). The horizontal axis represents the lateral deviation in metres. These results show that Dub4 is capable of closely following the experience map, with most deviations within 0.5m.

4.6. Results

We benchmark Dub4's suitability to Mars-like environments by computing the probability of running without localisation for a certain distance. Figure 5 shows the probability of running without localisation (in open-loop mode) for the simulated autonomous runs. This simple metric is both readily humanly-understandable (and should be read as: "the robot should travel no more than X meters on dead-reckoning alone") and a good way of comparing the performance of different localisers.

The path driven by the rover is shown in Figure 3. The green loop represents the experience map, which serves as the T&R reference. The second loop (yellow, blue, and red) represents a simulated autonomous run, and shows sections of the route where Dub4 successfully accomplished localisation against the experience map (blue), portions where localisation against the experience failed, but the rover's position was successfully estimated using VO(yellow), and a short portion at the beginning of the simulated autonomous run where Dub4 was unable to estimate the rover's position (red).

The lateral deviation of the simulated autonomous runs with respect to the experience map created from the first run is shown in Figure 4. This visualisation is important when determining the performance and accuracy of the T&R system. The low deviations seen here would allow the mission planners to design more intricate autonomous traversals while maintaining rover safety.

Overall, Dub4 performed very well in a Mars-like environment, showing a very low overall probability of losing localisation. Future work could involve improving the invariance to changes in illumination and view-point, and developing feature detectors that are better suited for environments with few distinct landmarks and textures. It is important to note that Dub4 depends on stereoscopic vision, demanding a strong calibration with increased demand on computational resources to process sensor imagery and a heavier sensor payload when compared with monocular vision systems.

5. SINGLE-CAMERA NAVIGATION

Dub4 and other experience-based navigation (EBN) frameworks [16, 18, 19] are highly relevant to the rover navigation problem, giving the rover access to a network of reusable paths within which it can efficiently localise, and upon which it can repeatedly expand as it engages in forays into the world. Unlike state-of-the-art T&R systems using stereo vision, an example of which we described in Section 4, we developed in [20] a light-weight, low-cost, and reduced-bandwidth alternative which exploits the planarity of imagery captured by a single downward-facing camera. Here, odometric constraints are available by tracking the perceived texture of the ground in front of the robot, and computing a simple frame-to-frame homography.

5.1. Topometric graphs for navigation

Generalising the graph-based localisation problem, we require the map nodes, M_i, to store data which is representative of discrete and distinguishable places in the

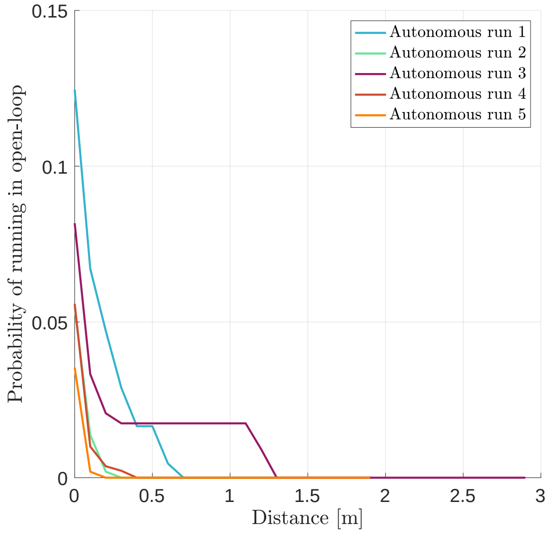

Figure 5. Probability of running without localisation (in open-loop mode) for the simulated autonomous runs. The horizontal axis represents the distance travelled in open-loop. The vertical axis represents the probability of travelling a certain distance in open-loop mode.

world (e.g. image, laser scan, etc). An edge between nodes i and j must characterise a relationship between places ${}^{i}\phi_{j}$. To search, we must be able to compose the edge constraints ${}^{i}\phi_{k} = {}^{i}\phi_{j} \oplus {}^{j}\phi_{k} \ldots$, along a path $P = \{{}^{i}\phi_{j}, {}^{j}\phi_{k}, \ldots\}$ in the database. Finally, to decide if two places are close together, we need some symmetrical metric $|{}^{i}\phi_{j}| = |{}^{j}\phi_{i}| \geq 0$.

5.2. Frame-to-frame edge constraints

As the robot moves (from time $k-1$ to k), it will observe a point of interest (in 3-D space) at slightly different image coordinates $\mathbf{x}_{k-1}$ and $\mathbf{x}_{k}$, due to its motion. If this point is on or near an approximately planar surface, the image coordinates are related by a (3×3) matrix (called a planar-induced homography) such that $\mathbf{x}_{k} = \mathbf{H}_{k,k-1}\mathbf{x}_{k-1}$. For every pair of frames that the rover observes as it moves, we can estimate this homography by iteratively discarding points not on the plane in a RANSAC process [21]. Examples of these interesting image pixels are shown in Figure 6, including many outliers (e.g. points on the horizon) that need to be filtered by the estimation process.

5.3. Characterising places

Originally, we detected and matched interesting points using speeded-up robust features (SURF) [22], due to its robustness to viewpoint changes, and broad support in the

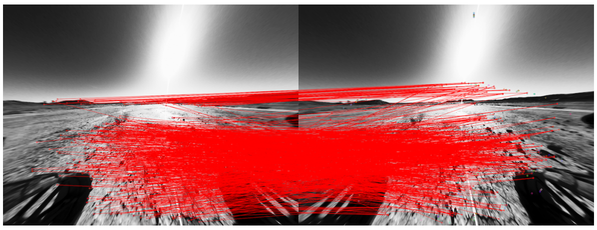

Figure 6. Sample imagery showing interesting features (detected and described by SURF) on two subsequent images captured by the rover as it moves, and matches from frame-to-frame

software community. However, this framework leaves the designer free to choose place characterisation and connectivity. Indeed, we have successfully integrated GPU acceleration [23] and a fast binary descriptor [24].

5.4. Localisation

Having populated the database with an initial mapping run, localisation is then achieved by optimising the metric for "closeness" over the set of nodes in a modest (in terms of total edges traversed) graph neighbourhood $M = \{m \mid \rho - \sigma \leq m \leq \rho + \sigma\}$ around the current pose, ρ, which can be obtained by an arbitrary graph search algorithm (e.g. breadth-first, depth-first). Candidates for this metric may include a distance in homography-space (see [25]), or the overlap of the two images when the homography is used to warp one of them.

5.5. Recovering motion

While localisation in this example is achieved purely in the space of homographies, it is possible to recover the rotation and (up-to-scale) translation, as the homography decomposes to $\mathbf{H}_{k,k-1} = \mathbf{R}_{k-1,k} + \frac{1}{d}\mathbf{t}_{k-1,k}\mathbf{n}^{\pi T}$. For this, the intrinsic camera parameters (distortion, etc) and a singular-value decomposition (SVD) are required (see [26]). Furthermore, if the camera height is well calibrated, the absolute scale of the motion is available.

5.6. Experiments

We tested this system by repeatedly driving Q14 around a highly textured and somewhat flat area near the MURFI base. As shown in Figure 7(a), the greyscale imagery was captured by a Point Grey Grasshopper monocular camera at resolution 1036×1084. The camera's intrinsics were used to rectify images (Figure 7(b)) and was computed using the OCamCalib toolbox [27]. The route was piloted at regular intervals (8 traverses) over a period of 24 hours, totalling approximately 800m. As such, the data present us with variation in environmental appearance, predominantly due to the waxing and waning of available light as

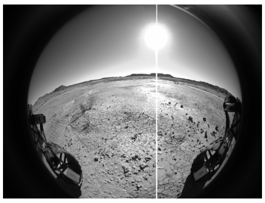

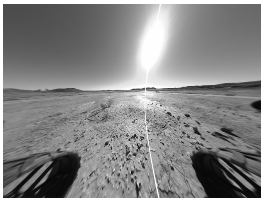

(a) Greyscale imagery captured by the fish-eye camera.

(b) Undistorted by applying the camera intrinsics.

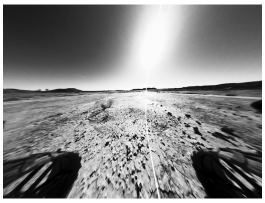

(c) Histogram equalisation can remove specularities caused by harsh glare.

Figure 7. Sample imagery from the downward-facing Grasshopper mounted on Q14, as well as examples of simple preprocessing steps taken to ensure good estimation and some robustness to changes in lighting.

the sun changed position. To this end, we ensured some robustness to the variation in lighting by normalising images using histogram equalisation [28], shown in Figure 7(c).

5.7. Results

In Figure 8 we show the ground-plane trace of the robot's motion across the desert floor. These traces, showing the decomposed motion of Section 5.5, are subject to a lot of drift, but we have shown in previous work (see [20]) that the scale factor (corresponding to the camera height) is consistent. Indeed, Figure 8 is a good illustration of the power of relative topometric maps, as we maintain regular localisation in the space of homographies (see Section 5.2). Here also for the first time, we show localisation between multiple experiences, as the robot grows its map and enriches it with topological (localisation) links.

6. CONCLUSIONS AND FURTHER WORK

As expected, Oxford VO performed very well during the MURFI mission proper. The T&R exercises showed that Dub4 can be successfully used in a Mars-like environment, with successful localisation on most portions of the "taught" trajectory and resilience to changing lighting conditions leading to low deviations ($\leq$ 0.5m) from the planned mission. This would in turn allow the mission planners to design complex and intricate autonomous exploration missions while maintaining the safety of the rover. The monocular camera based localisation testing also generated good results, indicating that this reduced-

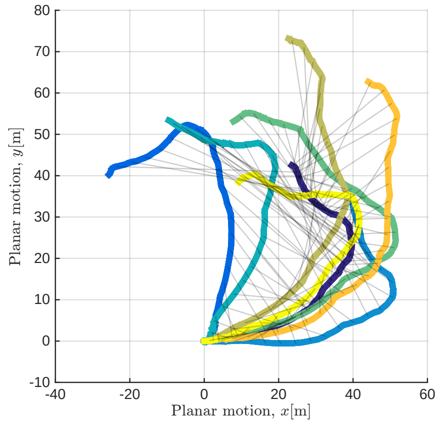

Figure 8. The gray lines connecting these colourful traces indicate the localisation of one drive (experience) to another. The frequency of these connecting lines increases with the accuracy of the metric pose estimate relative to the corresponding experience. This shows that despite accumulating error in the global position of the robot, it can still reliably localise (enabling accurate control), aided by a constantly evolving representation of the world.

bandwidth approach shows considerable promise. Overall we think it likely that only minor modifications to the underlying processes of each technique will be needed to ensure they are highly effective in planetary surface environments, and future work will implement these modifications. This will also address the potential to generate dense reconstructions of the terrain surrounding the rover using depth map data from the stereo camera fused with VO information. In addition, this phase of the work will investigate whether also fusing LIDAR-sourced data improves the quality of the reconstruction. This area has shown great potential in a terrestrial context to generate feature rich 3-D meshes of the scene. In a Space context, the ability to generate life-like representations of the scene could be extremely useful to mission scientists and others as an aid to better understanding the features of the area near the rover and helping to better focus the science effort. Looking further ahead, application of advanced machine learning techniques and advanced autonomy could bring the possibility of a true "robot geologist" closer to reality.

7. ACKNOWLEDGEMENTS

Horia Porav is supported by the Oxford-Google DeepMind Graduate Scholarship. Matthew Gadd is supported by the FirstRand Laurie Dippenaar, Oppenheimer Memorial Trust, and Ian Palmer scholarships. Paul Newman is supported by EPSRC Programme Grant EP/M019918/1.

REFERENCES

[1] Balme et al. Mars Utah Rover Field Investigation 2016 (MURFI 2016): Overview of Mission, Aims and Progress. *Lunar and Planetary Science XLVIII*, 2017.

[2] Harris et al. Remote detection of past habitability at Mars-analogue hydrothermal alteration terrains using an ExoMars Panoramic Camera emulator. *Icarus*, 252:284–300, 2015.

[3] Coates et al. The PanCam instrument for the ExoMars rover. *Astrobiology*, in press, 2017.

[4] Josset et al. The Close-Up Imager (CLUPI) on board the ESA ExoMars Rover: Objectives, description, operations, and science validation activities. *Astrobiology*, in press, 2017.

[5] Koroblev et al. Infrared spectrometer for ExoMars (ISEM), a mast-mounted instrument for the rover. *Astrobiology*, in press, 2017.

[6] Winston Churchill. *Experience Based Navigation: Theory, Practice and Implementation*. PhD thesis, University of Oxford, Oxford, United Kingdom, 2012.

[7] Shaw et al. Robust visual odometry for space exploration. *Advanced Space Technologies in Robotics and Automation*, 2013.

[8] Linegar et al. Work Smart, Not Hard: Recalling Relevant Experiences for Vast-Scale but Time-Constrained Localisation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.

[9] Linegar et al. Made to measure: Bespoke landmarks for 24-hour, all-weather localisation with a camera. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016.

[10] Nelson et al. Building, Curating, and Querying Large-scale Data Repositories for Field Robotics Applications. In *International Conference on Field and Service Robotics (FSR)*, Toronto, ON, Canada, June 2015.

[11] Maimone et al. Two years of visual odometry on the mars exploration rovers. *Journal of Field Robotics*, 24(3):169–186, 2007.

[12] Nister et al. Visual odometry for ground vehicle applications. *Journal of Field Robotics*, 23(1):3–20, 2006.

[13] Sibley et al. Planes, trains and automobiles: Autonomy for the modern robot. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 285–292. IEEE, 2010.

[14] Strasdat et al. Double window optimisation for constant time visual slam. In *2011 International Conference on Computer Vision*, pages 2352–2359, Nov 2011.

[15] Mark Cummins and Paul Newman. Appearance-only slam at large scale with fab-map 2.0. *The International Journal of Robotics Research*, 30(9):1100–1123, 2011.

[16] Winston Churchill and Paul Newman. Experience-based navigation for long-term localisation. *The International Journal of Robotics Research*, 32(14):1645–1661, 2013.

[17] McManus et al. Shady dealings: Robust, long-term visual localisation using illumination invariance. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 901–906. IEEE, 2014.

[18] Maddern et al. Leveraging Experience for Large-Scale LIDAR Localisation in Changing Cities. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.

[19] Michael J Milford and Gordon F Wyeth. Seqslam: Visual route-based navigation for sunny summer days and stormy winter nights. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1643–1649. IEEE, 2012.

[20] Matthew Gadd and Paul Newman. A framework for infrastructure-free warehouse navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.

[21] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

[22] Bay et al. SURF: Speeded up robust features. *Computer vision–ECCV 2006*, pages 404–417, 2006.

[23] Pulli et al. Real-time computer vision with opencv. *Communications of the ACM*, 55(6):61–69, 2012.

[24] Rublee et al. Orb: An efficient alternative to sift or surf. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2564–2571. IEEE, 2011.

[25] Heise et al. Pm-huber: Patchmatch with huber regularization for stereo matching. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 2360–2367. IEEE, 2013.

[26] Yi Ma. *An invitation to 3-d vision: from images to geometric models*, volume 26. Springer, 2004.

[27] Scaramuzza et al. A toolbox for easily calibrating omnidirectional cameras. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 5695–5701. IEEE, 2006.

[28] Finlayson et al. Illuminant and device invariant colour using histogram equalisation. *Pattern recognition*, 38(2):179–190, 2005.