Image Based Localisation and Autonomous Image Assessment for a Martian Aerobot

Mark Woods(1), Andy Shaw(1), Malcolm Evans(1), Roger Ward(1) Dave Barnes(2), Phil Summers(2), Gerhard Paar(3), Arnold Bauer(3), Derek Pullan(4)
(1)SciSys, Clothier Road, Bristol, BS4 5SS, UK, mark.woods@scisys.co.uk
(2)Dept. of Computer Science, Aberystwyth University, SY23 3DB, UK , dpb@aber.ac.uk
(3) Joanneum Research, Graz, Austria, A-8010, gerhard.paar@joanneum.at
(4) Dept. of Physics and Astronomy, University of Leicester, LE1 7RH, UK, dpu@star.le.ac.uk

Abstract

There is currently a strong demand for a range of intelligent vehicles to assist in solar system exploration and to pave the way for possible human missions. Balloon based planetary Aerobots offer a unique platform for carrying out remote science and can be used for a variety of applications such as high resolution imaging and rover guidance. However short to medium term missions of this type will be constrained in terms of power, communications, data storage and processing capability. To be of use, they must be able to localise and manage image data in an autonomous and intelligent manner, including prioritisation of images based on pre-defined science goals. This paper outlines a proposed scheme for a demonstrator testbed which will help provide and develop such a capability.

1. Introduction

For Mars, Venus and Titan, flying robots would provide a unique exploration perspective, complementing the views offered by ground based platforms or orbital surveillance elements. Such systems could supply planetary scientists with detailed views at regional scales and support a range of science tasks such as atmospheric meteorology, terrain mapping, instrument deployment, data relay and vehicle navigation.

A significant feature of Lighter Than Air (LTA) balloon systems is that a majority of the technology building blocks already exist compared to Heavier Than Air (HTA) systems. Simple LTA balloons have already been flown on Venus as early as 1985 as part of the Franco-Russo VEGA mission [1] which lasted for two days. As such they represent a relatively simple, low-cost way of achieving complex scientific goals.

In the near future it is likely that the first Martian aerobot mission take the form of an uncontrolled, free flying high resolution imaging mission. As the aerobot system will randomly drift in the Martian wind across its often rugged terrain, it will only be able to communicate with an orbiter as chance permits. This gives rise to challenging autonomy requirements in the area of image storage and transmission, and localization.

As an aerobot system will have limited resources of memory and power, the main problem will be economic storage and use of images acquired. All unnecessary imagery will need to be suppressed. Given the communication and memory constraints it will be unlikely that the aerobot will be able to store and transmit all of the images it can acquire. It will therefore require sufficient intelligence to autonomously prioritise images according to pre-defined set of scientific goals.

One other major requirement for any on-board scientific package will be positional data so that any results can be put in to context. Although other means of navigation could be provided, using already available imagery would be the most economic. Thus there is a requirement to investigate image based localisation to the required accuracy to support image management and uplink-scheduling.

2. Objectives

The objective of our work was to design and prototype an Imaging and Localisation Package (ILP) for a free-flying Martian balloon on a high-resolution imaging mission. The work was carried out for ESA/ESTEC in order to investigate and demonstrate the viability of such a mission. The original design was for an ILP package that would enable optimal acquisition of images to allow the reconstruction of accurate 3D topographical models of the surface of the explored planet, while also providing an accurate location of the balloon with respect to the Martian environment. To reduce the system weight and power requirements the proposed ILP system used a downward facing camera and computer vision techniques to:

- acquire and store images at various resolutions
- construct and update a 3D model (DEM) of the surface topography
- constantly estimate the position (latitude, longitude and altitude) of the aerobot as well as its motion with respect to the surface and
- decide on the basis of the communications budget, the morphology of the surface and through an intelligent assessment of the information content of the images at which resolution/compression they needed to be transmitted to Earth

In order to test the ILP software functionality the study produced a bi-modal demonstrator framework consisting of a balloon simulator and an actual prototype balloon with wireless communication capability.

The demonstrator system (DemoShell) could be used in either an all software mode where a balloon simulator was used to provide synthetic images under a variety of conditions to the ILP or in a hardware mode where the images originated from the balloon mounted camera. In both cases a PC-based GUI application was used to configure and control the demonstrator as well as allowing an interactive display of the ILP output. As a result the development was split into the following areas:

- DemoShell: Configuration and interface GUI
- On-board ILP system software
- Martian balloon simulator
- Hardware balloon system

3. ILP System Overview

Figure 1 outlines the basic components and processing flow of our ILP system. Each of the main components is explained further in the following sections.

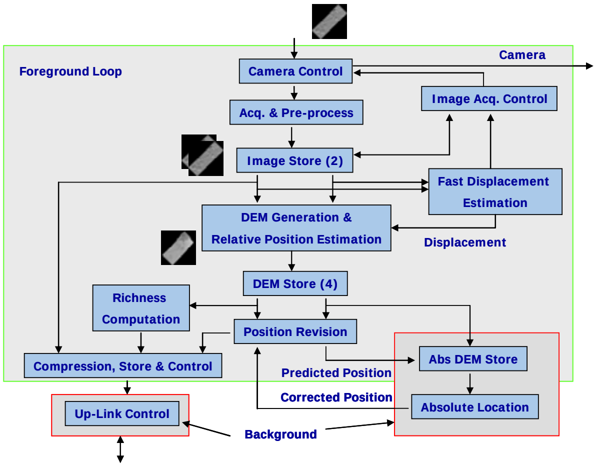

Figure 1: ILP System Overview

3.1 Image Acquisition

A key feature of the proposed operational scenario is that there are hard real-time processing requirements to consider which are a function of the Aerobot's speed and height as Figure 2 indicates.

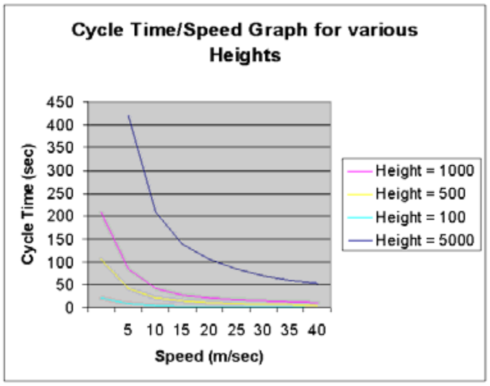

Figure 2, The impact on the ILP cycle times for various altitudes and velocities.

The image acquisition process was therefore designed to reduce the computational overhead of the Digital Elevation Model (DEM) generation process. An initial pre-processing algorithm determines the aerobots displacement from successive images and uses this measurement to control the timing of the image acquisition. The aim of the control scheme is to ensure almost ideal overlap (70%) between successive images in order to reduce the time taken to generate the DEM and prevent redundant processing. DEMs require the image pairs to be similar but well contrasted. Therefore, with changing planetary conditions, camera exposure parameters may need to be controlled automatically.

3.2 DEM Generation

A DEM is a standard data structure for the digital representation of a planetary surface. For each x/y coordinate within the represented area a height can be directly derived from the DEM.

In this particular aerobot application the DEMs were used in the ILP to support the global localisation by comparison with a global DEM and image richness assessment. By displaying received DEMs on the ground segment GUI, scientists can complement knowledge obtained from image mosaics on local geology.

In the aerobot case, DEMs were built using images taken at different positions using the aerobot's single camera system along its trajectory. For the photogrammetric determination of the distance to a certain scene point (without loss of generality in the coordinate system of the temporally first image) the relative orientation between the two camera positions (pointing and displacement vector) must be known, the respective scene point must be visible in both images. Consequently as the images involved need to overlap only the region within the overlap area can be used to generate a DEM. Any missing information leads to ambiguities in the resulting DEM data, such as unknown scaling or rotation and displacement with respect to the global coordinate system. Scaling problems can be resolved by one independent distance measurement to a point identified in both involved images (e.g. using an altimeter to provide distance to a point whose direction is known in the camera coordinate system), or the knowledge about the displacement between the stereo images. The real-time constraints identified in Figure 2 meant that the DEM generation had to be carried out as efficiently as possible. A number of techniques were developed to speed up this process and reduce computational load including the provision of ideal image overlap through controlled acquisition thus removing the need to estimate displacement. Advance knowledge of displacement itself also shortens the DEM generation process.

Integration into the global coordinate system was then possible though the identification of DEM landmarks that are also present in the global DEM. DEM generation is also highly connected to relative localisation since tracking of landmarks in principle is the same process as the stereo matching needed for photogrammetric DEM reconstruction.

3.3 Global and Relative Localisation

Autonomous localisation is another key requirement for a free-flying aerobot mission, as it is required to provide a reference for the scientific image analysis and to enable a prediction of forthcoming uplink windows, which in turn are required to manage on-board memory and image compression. However this was not trivial as Mars does not currently possess a Global Positioning System (GPS) and systems such as rate gyros, accelerometers, compasses and sun/star sensors are subject to error and drift and are therefore inadequate by themselves.

Although it is possible to derive positional information when in contact with an orbiter whose position is known, however these contacts are infrequent and cannot be used as the sole means of localising the aerobot.

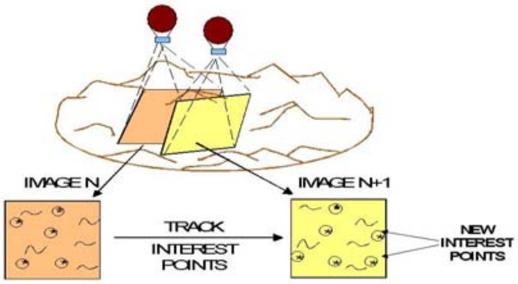

Figure 3, Relative localisation using landmark tracking.

An alternative scheme was adopted for this work which used a combination of computer vision based global and relative localisation in conjunction with orbiter contacts to provide continuous aerobot localisation. The relative method used landmark tracking shown in Figure 3, which was a by-product of the DEM generation process to provide a relative location fix. Initial testing on a small-scale Martian analogue terrain showed repeatability within 1% of the travelled distance for the relative localisation. Although through the traverse these relative measurements must be corrected with an absolute reference. To make up for infrequent orbiter contacts a novel landmark comparison scheme [2] was used to provide an absolute position.

The landmark comparison method consists of two individual elements; the first is the analysis of the terrain for geological features, providing a density profile and the second, a gradient profile of the surface. The feature analysis is obtained by calculating the radius of curvature at each elevation position in the DEM and through analysis of the curvature encompassing that point, classification in obtained. Figure 4 shows a false height and feature extracted (the features are represented through different colours) local DEM, the region selected for the feature density profile (area within the circle). The feature density and the gradient profile within this area were used during the search of the global DEM.

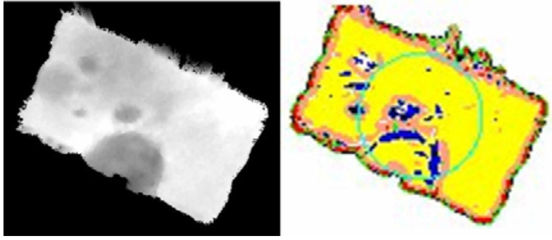

Figure 4, Local DEM false height coloured and the feature extracted DEM, region selected for analysis is within the circle.

Figure 5 shows the global Martian test terrain feature extracted with the same geological parameters as those used in the local analysis. An actual location for the local DEM in Figure 4 within the global DEM can be generated by using the feature and gradient profile region (form the local DEM) and matching it to a region as these have to be the same in the global DEM.

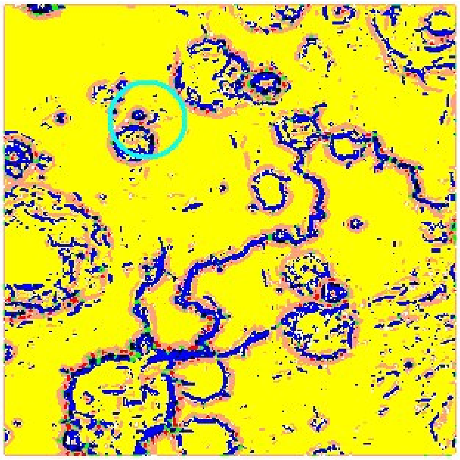

Figure 5, Feature extracted global DEM of the Martain test terrain; the circle shows the matched location for the local DEM from Figure 4.

Through experimentation it was found that the largest influencing factor for obtain a match was the size of the local DEM. A local DEM that covered an area of approximately 30x30, global resolution (or more) produced 98% correct matching and areas that covered approximately 15x15 global resolution, produced over 70% correct positions. These results were also highly dependant on the number of features present.

There are a number of possible sources of reference topographical information of the Martian surface data. For this study both Mars Orbiter Laser Altimeter (MOLA) and the Mars Orbiter Camera (MOC) from Mars Global Surveyor (MGS) were used. Absolute localisation can therefore be achieved by extracting naturally occurring features (peaks, ridges, channels etc.) from the topographical maps. By categorising the surface by its features, then by matching these features in a high resolution DEM generated by the aerobot, with the same features in the low resolution global map (e.g. MOLA) a global position estimate is obtained.

The obvious limitation of the scheme and one that has been shown during experimentation is that it is better suited to areas which are reasonably feature rich and diverse such as the Southern Highlands. However this constraint is fully compatible with the scientific goals which seek to explore such geologically rich areas.

3.3 Image Richness Estimation

An Image Richness Index (IRI) was used to determine the priority that should be given to the acquired images to ensure that those of high interest were not lost, if the amount of image data that could be transmitted to the orbiter became restrictive.

Ideally, this requires a sophisticated model of the planetary science analysis capability currently associated with human teams. Clearly this is a complex problem and the research area itself is relatively immature involving a wide range of technology areas including computer vision and various strands of AI. A further issue in this work is that the science assessment must be carried out in real-time i.e. before the next image is taken as the aerobot transits and area. Time available for assessment depends on the speed and altitude of the aerobot.

In our approach the science members of the team defined a set of features of interest at a number of levels. At the primary, ‘first-pass’; level where significant geomorphological features such as impact crater’s, channel’s, volcano’s, dunes etc. At the secondary level, evidence of features such as cross-bedding which could be used to distinguish between volcanic or aqueous channel creation are considered important. Our work focused on the detection of these primary features in the first instance to generate a gross assessment of an individual area. Based on the science team input we have assumed, that altitude has a ‘fractal effect’ on the representation of features in images. For practical purposes therefore, our detection algorithms consider feature structure to be effectively invariant to this parameter. Consequently, macro or contextual relationships between nested features are not considered in this phase of detection.

Of particular interest in determining the presence of the primary features are the DEM’s generated on-board by the ILP. Raw DEM's provide rich, 3-D structural information which are more suited to primary feature detection than 2D orthoimages. Our core assessment is based on the determination of key morphological variables (MV's) such as gradient and various curvatures which can be derived using first and second derivatives of a DEM. These variables have are widely used by the geomorphology community to assist in the characterisation of land surfaces. Our MV estimation method is based on the algorithms presented in [3]. Once derived, slope, horizontal and vertical curvature are used to classify individual pixels as being planar or ridge etc. based on a modified version of an approach presented in [4]. A global pixel assessment is then carried out and used to provide the overall rating for a particular DEM. The presence of a large number of ridge or plane pixels are be used to provide a coarse indication of the geomorphological richness of the area.

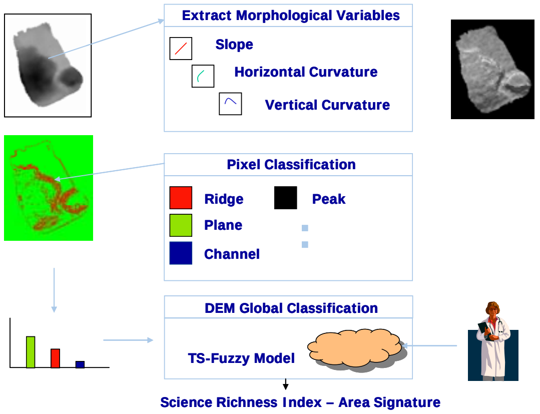

Figure 6: Basic architecture of our image richness assessment system

The nominal ILP implementation uses the coarse assessment as a basis for a richness measure. We have implemented a Takagai-Sugeno Fuzzy Inference System to capture subjective, non-linear human views of the feature vector-richness mappings. This implementation is attractive as the model is relatively computationally efficient, and is well-suited to initialisation through explicit natural language rules or off-line training based on observed input/output pairings.

We have also implemented a simple time variable analysis model which configures the level of assessment/classification carried out based on the amount of predicted processing time. This allows the aerobot to autonomously configure the amount of analysis to be carried out in order to meet its real-time constraints.

3.4 Image Storage, Management and Transmission

It was assumed that the amount of mass memory available would not be sufficient to hold all of the images produced before they could be transmitted. Some means of controlling compression and deletion of images on-board was therefore required. The image richness index was used to determine the priority given to each image, and its compression ratio to preserve the image's most interesting areas.

Compression provided a means to greatly reduce and control the actual amount of storage required. Progressive compression techniques (ECBOT wavelet JPEG 2000) were used to compress the images and DEM data. This technique proved resilient against drop out and also allowed the entropy coding to be tailored to meet any data length against image richness profile required by controlling the scaling against signal to noise, resolution, visual quality, or of regions of lesser interest etc. Since the window of communication would vary as much as the vagaries of the wind, continuous update of the capabilities of the up-link were calculated, allowing the adjustment of the compression ratio by layered deletion of the compressed image to roughly meet the link capability, and the memory storage requirements.

When a link was established the IRI was used to order images for transmission to ensure that the most important images were transmitted first. This represents an important paradigm shift for a surveillance mission and represents a significant increase in mission autonomy.

4. Demonstrator Overview

The ILP demonstrator system developed required two distinct modes. An all software mode (see Figure 7) that used an environment and balloon simulator and a hardware mode that encompasses a real PC controlled balloon, with a camera, altimeter and wireless connection. Both modes used the same interface to a demonstrator shell GUI and associated ILP software.

Using the simulation based mode as one of the core demonstrator components allowed a number of advantages over just a hardware-only solution:

- Repeatability of experiments.
- Simulation of environments that could no be re-constructed within the laboratory.

Simulators are able to re-construct environments and terrain to a very accurate level whilst still allowing the user to maintain control over the specific parameters. Terrain, weather, atmosphere and hardware devices, such as cameras, can be simulated in a realistic way and noise can be modelled to allow for random fluctuations in the environment or manufacturing tolerances for instruments. A simulator is well suited for modelling the flight of a balloon over a terrain and experiments cannot always be carried out due to numerous factors, so the use of simulation overcame any of these potential problems.

The trajectory of the balloon was initially supplied to the simulator, then it could either run in a "free floating" or "controlled" trajectory mode. A simulator also affords the ability to evaluate algorithms for localisation and control that may be difficult to examine on real hardware in an Earth-based environment. A major factor in the deployment of planetary balloons is the effect the environment and atmospheric conditions have upon the vehicle.

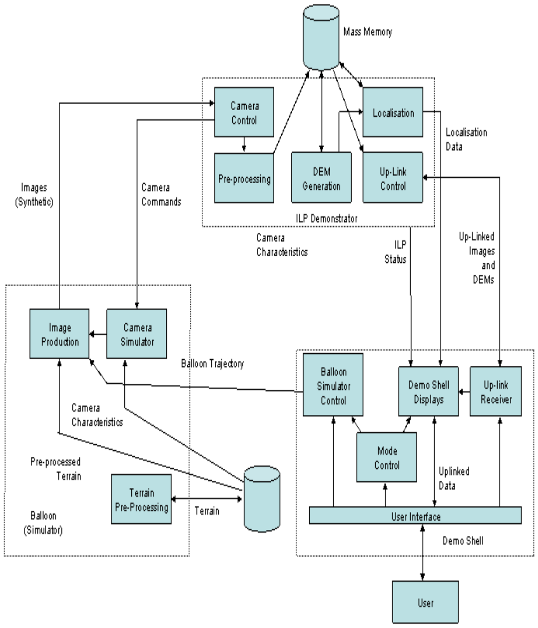

Figure 7, Demonstrator - Software mode. The Hardware mode replaces the simulator module with the actual gondola and camera system

Despite the availability of a simulator there was a clear need to test the ILP software in a more realistic environment. To this end a model balloon complete with propulsion, downward looking camera, processing and wireless control and communication was developed. A realistic test facility was necessary which was provided by the ESA/ESTEC PTB with its mock Martian terrain and high ceiling. To complement the PTB facility a small Martian surface was set up to allow for algorithm testing and development.

The demonstrator shell communicates with both the ILP and balloon simulator in order to receive up-linked DEMs and images and to download camera models and balloon trajectory information. The shell also allows the selection of the demonstration mode, hardware or software, provides an interface to allow selection and upload of camera specifications, antenna and terrain models and enables the user to define the communication window function. The user is also able to start, stop and pause the simulation at will.

Another functionality of the demonstrator shell was the display of the demonstrator software status. This consisted of the data provided to the demonstrator shell by the balloon simulator and the data produced by the ILP itself. A graphical interface was used to display the data in a clear and concise form. Presentation of this ILP data incorporated the display of image mosaics, balloon trajectory markers, communication windows and transfers and altitude. This requires the demonstrator shell to reconstruct the sequence and determine the placing of the individual images within these mosaics. It must also display the DEM data received by the demonstrator shell. The DEM data was presented in a format suitable for easy visualisation using the DEMView software addition to the demonstrator shell. Through the functionality of the DEMView it was possible to perform measurements on the DEMs such as the co-ordinates of individual points and the distance between the selected points.

5. Results

The development and testing of the various algorithms mainly took place within the simulated environment due to the ease of setting up and running the experiments and simulating actual Martian visual and flight conditions. Once DEM's were produced an image richness assessment applied. A global pixel assessment was carried out and used to provide the overall rating for a particular DEM. The presence of a large number of ridge or plane pixels was used to provide a coarse indication of the geomorphological richness of the region. Figure 8 shows results from the Image Richness analysis, the ortho-image for context, associated DEM, and a subsequent extraction of the 'raw' boundaries of a crater in the image.

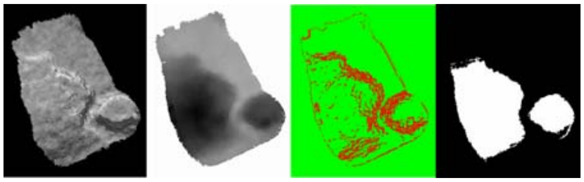

Figure 8, Local DEM Image Richness, (from left to right) Context ortho-image, Associated DEM, Raw ridge extraction and feature segmentation.

Figure 9 shows the results from the absolute matching algorithms over the test Martian data. Here it is possible to see the oval shaped trajectory the aerobot took and the individual local DEM's the ILP software created (placed around the global data) and the calculated aerobot position on the surface.

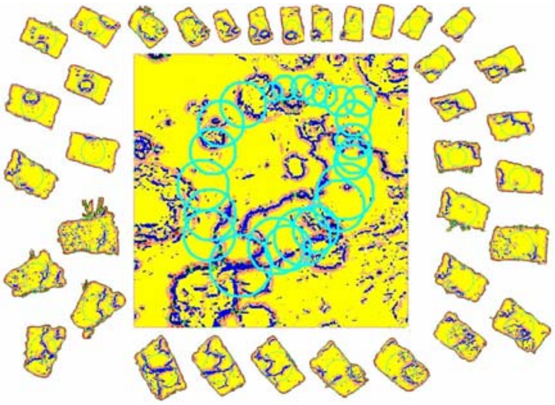

Figure 9, Results from the absolute matching algorithms over the test environment.

During the development of the ILP system the simulation environments were used frequently, because of the repeatability they offered. Figure 10, shows a scene shot from the balloon system simulator (BSim). Through the simulator it was possible to test the accuracy of the reconstructed DEMs because the scenery topography was accurately known. It was also possible to test the ILP system at various altitudes and velocities, with different camera specifications, as well as with various surface textures.

After numerous experiments using the simulated environments, work started on the real hardware aerobot. Rather than having a small RC balloon system transmitting all the data back to a base PC for analysis it was decided that all the processing should take place on board, and only the important data would be sent back

Figure 11 shows the developed aerobot system at the European Space Agency's Planetary Utilisation Test Bed facility. In this image it is possible to see the mock Martian terrain with Martian rock distributions and one of their test rovers that they use on the surface

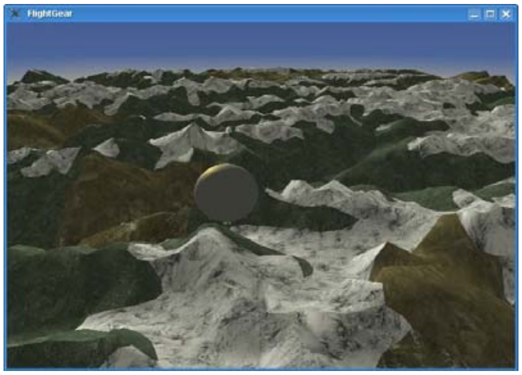

Figure 10, BSim simulated aerobot and environment used for system testing (Note Himalayas test scenery pictured).

Figure 11, ILP aerobot undergoing acceptance trials at the ESTEC PUTB facility.

The PUTB facility is approximately 8x8m in size with several terrain types, which allowed for numerous experimental traverses across the surface in several directions. Some of the test trajectories had rapid changes of direction and altitude which where using to evaluate the robustness of the systems algorithms.

The real-time display of the output from the ILP system can be seen in Figure 12. In this image it is possible to see the mosaiced images from the DEM generation phase, the actual relative location of the aerobot, the estimated velocity and relative heading and the communication information. The black areas with in the DEM are where it wasn't possible to get correlation between subsequent images. Looking at the white markers it's possible to see that the aerobot trajectory was "U" shaped.

One of the main requirements of this study was the ability to view the gathered DEM data in 3D. Along with the DemoShell interface software a 3D DEM viewer (DEMView) was also developed, Figure 13. Using the DEMView software it is possible to measure distances between surface points. It is possible to view the terrain with either the ortho-image overlaid or with a false colour, enabling easy surface inclination assessment. This would be a useful tool if the system was being used for rover navigation, as it would be possible to measure objects and gaps, before the rover was commanded to that region.

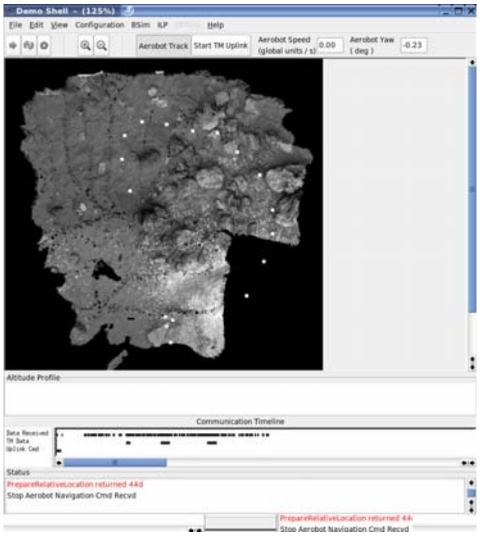

```
PrepareRelativeLocation returned 44d
Stop Aerobot Navigation Cmd Recvd
```

Figure 12, DemoShell GUI to the ILP system showing the data captured during testing at the ESTEC PUTB facility. (Image shows the mosaic of the Martian terrain).

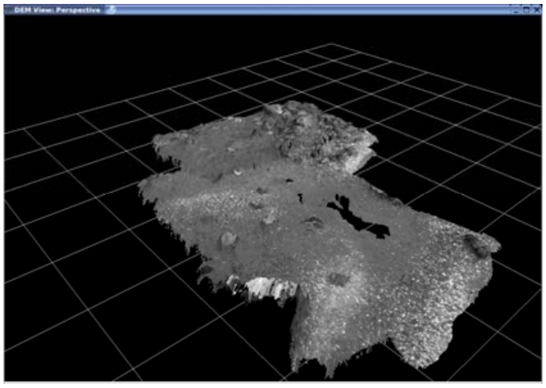

Figure 13, DEMView 3D data viewer for DemoShell.

7. Conclusions

This was an ambitious and challenging project given the number of complex functional elements that had to be developed and integrated.

The final ILP system and the test environment showed that an image based aerobot system would be capable of gathering images, localising and data handling under real world environments.

Perhaps a unique aspect of this work was the autonomous determination of science or image richness. In general terms this is a widely applicable concept. It is potentially directly relevant to the current ExoMars mission [5] where large volumes of image data may have to be compressed or prioritised based on an on-board determination of science or planning quality. In follow-on work [6] we have intergrated this concept with on-board re-planning technology [7] to demonstrate the potential benefits to this mission.

9. Acknowledgements

This work was funded under European Space Agency (ESA) Contract No.17400/03/NL/CH. The authors would like to thank Gianfranco Visentin, Head of the Automation and Robotics Section at ESTEC (ESA D/TOS) for his collaboration and technical guidance on this project.

10. References

[1]: Kremnev, R. S.; Selivanov, A. S.; Linkin, V. M.; Lipatov, A. N.; Tarnoruder, I. IA.; Puchkov, V. I.; Kustodiev, V. D.; Shurupov, A. A.; Ragent, B.; Preston, R. A., “The Vega balloons - A tool for studying atmosphere dynamics on Venus”, Soviet Astronomy Letters (ISSN 0360-0327), vol. 12, Jan.-Feb. 1986, p. 7-9. Translation.

[2]: A. J. Shaw, and D.P. Barnes, (2003). “Landmark recognition for localisation and navigation of aerial vehicles” *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas.

[3]: P.A. Shary, L.S. Sharaya, and A.V. Mitusov, 2002. “Fundamental quantitative methods of land surface analysis”. *Geoderma*, 107(1-2), 1-32.

[4]: J. Wood, (1996), The “Geomorpholological Characterisation of Digital Elevation Models”, PhD Thesis, University of Leicester, UK.

[5]: M. Van Winnendael, P. Baglioni, J. Vago, "Development of the ESA ExoMars Rover", *Proc. 8th Int. Symp. Artif. Intell., Robot. Automation and. Space*, Sep. 5-8 2005, Munich, Germany

[6]: M. Woods, A. Shaw, P. Rendell, D. Barnes, S. Pugh, D, Price, D. Pullan and D. long, “CREST Autonomous Robotic Scientist: Developing a Closed-Loop Science Exploration Capability for European Mars Missions, *In Proc. 9th* Int. *Symp.on AI in Robotic Automation and Space*, Feb 2008, Los Angeles, USA, *In Press*

[7]: M. Woods, L. Baldwin, G. Wilson and S. Hall, R. Aylett, “MMOPS: Assessing the Impact of On- Board Autonomy for Deep Space Missions” In Proc. *SpaceOps*, Rome, Italy AIAA-2006-5977