

AUTONOMOUS IMAGE BASED LOCALISATION FOR A MARTIAN AEROBOT

D. Barnes ^{a,*}, A. Shaw ^a, P. Summers ^a, M. Woods ^b, R. Ward ^b, M. Evans ^b, G. Paar ^c, M. Sims ^d

^a Dept. of Computer Science, University of Wales, Aberystwyth, SY23 3DB, UK. - dpb@aber.ac.uk

^b SciSys (Space & Defence) Ltd., Bristol, BS4 5SS, UK. - Roger.Ward@scisys.co.uk

^c Institute of Digital Image Processing, Joanneum Research, Graz, A-8010, Austria. - gerhard.paar@joanneum.at

^d Space Research Centre, Dept. of Physics and Astronomy, University of Leicester, LE1 7RH, UK. - mrs@star.le.ac.uk

Commission I, Working Group I/V

KEY WORDS: Robotics, Space, UAV, Aerial Digital Cameras, DEM/DTM, Autonomous navigation, Image-based navigation

ABSTRACT:

Balloon based aerobots have much to offer ESA's future planetary exploration programmes, e.g. high resolution mapping, landing site selection, rover guidance, data relay, sample site selection, payload delivery, and atmospheric measurement and meteorology. Aerobots could be used in a variety of configurations from uncontrolled free-flying to tethered rover operation, and are able to perform a range of important tasks which other exploration vehicles cannot. In many ways they provide a missing 'piece' of the exploration 'jigsaw', acting as a bridge between the capabilities of in-situ landers and rovers and non-contact orbiters. Technically, a Lighter than Air (LTA) aerobot concept is attractive because it is low risk, low-cost, efficient, and much less complex than Heavier than Air (HTA) vehicles such as fixed wing gliders, and crucially, much of the required technology 'building blocks' currently exist. Smart imaging and localisation is a key enabling technology for remote aerobots. Given the current lack of comprehensive localisation and communications systems, it is important that aerobots are equipped with the ability to determine their location, with respect to a planet's surface. The availability of a variety of terrain feature extraction, point tracking, and image compression algorithms, means that a self-reliant system is now achievable. We have developed a demonstrator imaging and localisation package (ILP) for a Martian balloon. This ILP system incorporates a unique combination of image based relative and absolute localisation techniques. We have demonstrated our ILP using both simulation and a real laboratory based model aerobot. The availability of both simulated and real aerobot data thus providing a comprehensive test and evaluation framework for the ILP functionality. The results from our work are presented in this paper.

1. INTRODUCTION

Whilst much attention has been given to the use of landers and rovers for planetary exploration, most notably the NASA JPL Sojourner rover (Matijevic, Bickler, *et al.*, 1997), and the current Mars exploration rovers Spirit and Opportunity (Squyres *et al.*, 2004), the use of flying robots, or aerobots, for planetary exploration represents a highly innovative concept. Whilst lander and rover technology is clearly competent at facilitating useful science, their application is terrain limited. Rovers are capable of travelling relatively small distances and much of a planet's terrain is impassable to small wheeled vehicles, aerobots in comparison have no such limitations.

For those planets and moons that support an atmosphere, flying robots are likely to provide a practical solution to the problem of extended atmospheric meteorology, planetary surface coverage for terrain mapping, and surface/sub-surface composition surveying. Not only could such devices be used for suborbital mapping of terrain regions, but they could be used to transport and deploy science packages or even microrovers at different geographically separate land sites. If tethered, then they could provide valuable navigation information for the control of the attached surface rover.

The challenge of flying a planetary aerobot encompasses mobility control and autonomous navigation in a constantly changing 3D environment. Inter-planetary distances prohibit the real-time communication of external meteorological and internal robot state data which would allow human remote control from Earth. An aerobot's long term endurance and ultimate survival can be achieved only if sophisticated autonomous flight control and navigation methods are employed. It is to address these challenges that our research is dedicated (Barnes, Summers, and Shaw, 2000).

Aerobots drifting in the vagaries of a Martian wind across its often rugged terrain, will need to be autonomous, communicating with an orbiter as chance permits. With limited resources of memory and power, the main problem will be economic storage and use of acquired images. All unnecessary imagery will need to be suppressed. Positional data will need to be provided to any scientific package carried. Although other means of navigation could be provided, using already available imagery would be the most economic. Thus our study is designed to investigate localisation to the required accuracy (commensurate with the imagery requirements), economic use of memory to store the vast quantity of images, and the problem of predicting the next communication window with an orbiter, so that the best use of storage can be made to meet the mission targets without loss of any important information. The main outputs of the mission are to be, a 3D digital elevation model of the surface (DEM), and images of the surface at various resolutions.

This project has studied the design of an Imaging and Localisation Package (ILP) for a Martian balloon. The package allowing optimal acquisition of images to reconstruct accurate models of the surface of the explored planet, and accurate localisation of the balloon with respect to the Martian surface. The ILP by means of aerobot mounted cameras and computer vision techniques can:

- acquire and store images of the surface at various resolutions,
- construct and update a 3D model (DEM) of the surface,
- 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 base of the communications budget, of the morphology of the surface and of the information content of the images (i.e. 'image richness'), which images at which resolution/compression need to be transmitted to Earth.

* Corresponding author.

2. ILP SYSTEM OVERVIEW

The ILP system is designed to control, acquire and process images from a single camera suspended from the balloon gondola. There are a number of technical challenges posed by individual tasks such as assessing image scientific content or richness, and predicting uplink windows (Woods, Evans, Ward, Barnes, *et al.*, 2004), and additionally there is the general problem of performing these tasks in a constrained real-time environment. We have investigated available cycle times versus balloon speed for a number of altitudes using a camera with a 70° field of view (FOV), and maintained image overlap of $> 70\%$. When these figures are examined within the context of a laboratory based demonstrator, available cycle time is reduced dramatically due to the very small altitude distance between the ILP balloon system and a laboratory mock terrain (typically 2.5m). Significant algorithm optimisation is required to meet these real-time demands, together with a very slow ILP balloon ground velocity, and fast image data processing.

To combat these real-time constraints we have constructed an ILP balloon gondola, see Figure 1, which contains an on-board mobile PC, hard-disk drive, wireless ethernet for ground station PC communication, laser based altimeter, digital camera and battery power supply. Image data processing on-board negates the bandwidth and time delay problems associated with first transmitting images to a remote ground based PC prior to data processing. Key to our method of aerobot localisation is the generation of a DEM of the terrain below. However, as only one camera is flown on our gondola to limit mass etc., then additional data is required by our DEM generation algorithm. This data is provided by our gondola altimeter. To emulate the effects of the Martian wind on our ILP gondola and helium filled balloon, we have designed and built two pairs of micro-propulsion units. The units within each pair are counter rotated. One pair provides vertical thrust, and these are attached to the gondola, whilst the second pair, which provide horizontal thrust, are attached to the balloon envelope about the total system centre of mass. This arrangement provides maximum camera stability when the balloon is in motion, and prevents any undesirable gondola ‘pendulum’ like motion. The small micro-propulsion unit propellers possess bi-directional pitch variable blades. These allow very low and precise balloon ground speed control to be achieved ($< 2cm/sec$). Each micro-propulsion unit is connected to the on-board PC via servo drive interface electronics. The total gondola mass is $\approx 2kg$ and neutral buoyancy is achieved with a helium inflated envelope of diameter $\approx 1.8m$. The gondola camera has been calibrated using the same rig that was used to calibrate the Beagle 2 Mars Lander PAW stereo cameras (Barnes, Phillips, and Paar, 2003). This has allowed numerous camera image mosaics of the ESTEC Planetary Test Bed (PTB) terrain to be captured for ILP demonstration purposes.

An important aspect of our ILP system is that it incorporates a unique combination of image based *relative* and *absolute* localisation techniques. Information about position and pointing of an aerobot platform is necessary to spatio-temporally assign each measurement from any science instrument or other sensor on the aerial vehicle. Here we define absolute localisation as being undertaken when, for example, a DEM generated by an aerobot ILP is spatially matched to that of a global reference DEM. Such a DEM could come from an orbiter’s data, e.g. MGS MOLA, or Mars Odyssey THEMIS data after a ‘shape from shading’ algorithm has been applied. However, DEM generation on an aerobot is computationally expensive, and needs morphological landmarks. Conventional localisation sensors may be available, but in the standard case they only provide pointing information (e.g. star/Sun trackers), or short-term information (e.g. gyros,

Doppler radar). Therefore landmarks tracking on the aerobot captured camera images could be used directly for relative localisation. This process thus filling the gap between that of short term location sensors, and that of absolute localisation using an orbiter generated DEM, such that the position and pointing of the aerobot platform is known at any time.

3. LOCALISATION - RELATIVE

The accuracy of relative navigation by image based 2D landmarks tracking depends on a set of key parameters. They can be divided into system parameters (e.g. camera resolution, FOV, camera calibration), and situation dependent parameters which induce a set of different options depending on the availability of information sources apart from the images. Our relative localisation method employs a landmark tracking algorithm that is used to follow upon the tracks of homologue points, called *Interest Points (IP)*, in two subsequent images. Corresponding displacements between the IP are then used as a data base for calibration update. To extract the landmarks from the origin image an *interest operator* is used. The essence of the calibration update method (i.e. identifying an instant camera position and pointing parameters along the aerobot path) is given as follows.

Consider two consecutive images (frames N and $N + 1$, see Figure 2) taken from the aerobot during motion. When frame N is involved in DEM generation, the 3D coordinates of the IP on frame N can be calculated from this DEM. Having these 3D coordinates and their 2D coordinates from tracking between frame N and frame $N + 1$, a calibration can be performed to obtain position and pointing parameters of frame $N + 1$. Further on, new IPs can be identified in frame $N + 1$, and using the exterior orientation from tracking calibration, these can be reconstructed in 3D. The procedure can then continue with frame $N + 2$ making use of all 3D coordinates of the $N + 1$ IPs. If no DEM is available for frame N , then 3D coordinates for the IPs can still be calculated from the frame N to $N + 1$ by tracking vectors and a scale factor from altimetry or displacement estimates (which provide information for relative orientation between frame N and $N + 1$), again allowing continuation with frame $N + 2$. A thorough analysis of this tracking method has been undertaken (Paar *et al.*, 1994), and experiments with a rover configuration and considerable horizontal displacement showed similar behaviour (Kolesnik and Paar, 1997). Based upon this research, our relative localisation method is in the range of 1% of a distance to ground d , when travelling horizontally a displacement of d .

4. LOCALISATION - ABSOLUTE

To maintain localisation accuracy over large aerobot flight distances, our ILP approach augments our relative localisation data with an additional absolute localisation approach. By matching an aerobot viewed terrain image with a region from a global image taken from orbit, it is possible to derive the aerobot’s global position and orientation. However, the practicality of this approach depends on the relative resolution of target and reference images as problems arise when trying to match images with varying resolutions i.e. matching low-resolution surface data taken from an orbiter to high-resolution data taken from the aerobot mounted camera. One solution to this problem is to abstract the information contained in the data being used. Our absolute localisation approach employs a feature extraction algorithm which is based upon the radius of curvature at a point on the surface. Six features are currently characterised, see Figure 3, and have been




Figure 1. ILP aerobot gondola together with vertical thrust micro-propulsion units. Horizontal thrust units are also shown. These are attached to the balloon envelope about the balloon plus gondola centre of mass.

colour coded for ease of visualisation (Red = Peak, Magenta = Ridge, Green = Pass, Yellow = Plane, Blue = Channel and Black = Pit). Figure 4 shows a terrain example that contains all six terrain features.

The percentage of extracted features from a scene can be varied by adjusting the slope, minimum curvature and resolution thresholds. The slope threshold determines an angle, below which the terrain is said to be horizontal. The minimum curvature threshold is an angle above which the terrain is said to be curved and the resolution threshold determines the number of surrounding data points used in the calculations. Absolute based localisation in an environment with large quantities of features can be difficult, but by automatically adjusting the slope, curvature and resolution threshold values, the number of features can be reduced, reducing the number of possible locations matched in the imaged data. Using a combination of terrain feature and gradient matching, our method for absolute localisation has proven to be robust, and the experiments conducted have been in feature rich environments with varying degrees of imaging noise.

It is therefore possible to obtain absolute localisation by matching images of a DEM of the region under inspection to regions in, for example, MGS MOLA data (Shaw and Barnes, 2003). The accuracy of the position information depends largely on the accuracy and amount of gathered data by the aerobot. The orbiter can also be used to confirm the determined aerobot position occasionally during data upload periods. This method cannot be considered solely for localisation because the position of the aerobot will frequently fall outside the orbital surface trace/track covered by the orbiter.

5. IMAGE RICHNESS PARAMETER

An important element of the aerobot function is its ability to determine the scientific interest of the images and DEMs that it collects as it traverses the planetary surface. As it can gather more data than it can offload, it is important that the prioritisation and compression of downlinked data is set according to the underlying scientific value. 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




Figure 2. Principle of relative navigation based on landmarks tracking.

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 the region. Time available for assessment depends on the speed and altitude of the aerobot.

The science members of the team defined a set of features of interest at a number of levels. At the primary 'first-pass' level were significant geomorphological features such as impact craters, channels, volcanos, 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 the 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 DEMs generated on-board by the ILP. Raw DEMs provide rich 3-D structural information which is more suited to primary feature detection than 2D ortho-images. Our core assessment is based on the determination of key morphological variables (MVs) such as gradient and various curvatures which can be derived using the first and second derivatives of a DEM. These variables 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 (Shary, Sharaya, and Mitusov, 2002). 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 (Wood, 1996). 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 being used to provide a coarse indication of the geomorphological richness of the area. Figures 5, 6 and 7 show the ortho-image for context, associated DEM, and a subsequent extraction of the 'raw' boundaries of a crater in the image.

Although the initial feature boundaries are poorly defined after the initial processing they are sufficient to provide a basic assessment which may be all that the real-time constraints allow. This is intuitively equivalent to a fleeting glance by a human geologist




Figure 3. The six terrain feature classifications.




Figure 4. 2×2 degrees test scene containing the six colour coded terrain features.




Figure 5. Context ortho-image.




Figure 6. Associated DEM.




Figure 7. Raw ridge extraction.




Figure 8. Feature segmentation.

travelling in a moving aerial vehicle. Figure 8 also shows that we are able to further segment and separate major features when more time is available by considering the spatial relationships between classified pixels. 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.

6. IMAGE STORAGE MANAGEMENT AND UP-LINK SCHEDULING

Since the time between communications with an orbiter will be variable and doubtlessly long, and the amount of imagery is dependent upon the vagaries of the wind velocity, assuming no area of surface beneath the aerobot is to be lost, image storage management is an important aspect of the study. Our image richness parameter is used to determine the priority to be given to the image, and potentially the image's most interesting areas. Compression provides a means to greatly reduce and control the actual amount of storage required. Progressive compression techniques (ECBOT wavelet) are being used to compress the images and

DEM data. Such a technique is not only more resilient against drop out but allows, for example, 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 will vary as much as the vagaries of the wind, continuous update of the capabilities of the up-link would be required and the compression ratio adjusted to roughly meet the link capability, and the memory storage requirements. However, if the memory storage requirements are not limiting, assuming an acceptable level of compression, then final compression decisions could await the actual beginning of the transmission.

Communications links with surface or low-altitude planetary exploration units are usually constrained in terms of frequency, duration and uplink downlink transmission rates. In addition, real-time commanding is not possible given lengthy RTLT durations. Given that imaging is a data intensive process, the communications link access needs to be optimised by employing intelligent use of available bandwidth and transmission opportunities. Thus the transmissions must be scheduled by predicting the timing and duration of the next up-link of data based on the aerobot's trajectory, the orbit of the relay satellite and the characteristics of the two antenna. During the transmission phase, the link must be acquired, any commands received and confirmed by the aerobot, before its image data can be transferred up to the orbiter, and their correct receipt confirmed. Only then can any images be deleted to make space for the new collection of images.




Figure 9. ILP aerobot undergoing trials with mock terrain sample manufactured from MGS MOLA terrain data.




Figure 10. Screen shot of the ILP aerobot flight simulator software with MGS MOLA based terrain.




Figure 11. Global MOLA elevation data for the Mamer Vallis crater region (32N, 340W).




Figure 12. Screen shot of Demoshell with mosaic of cropped camera ortho-images. The small white squares show the results from the aerobot relative localisation algorithm. Also shown is the aerobot 'Altitude Profile' for each relative localisation point, 'Communication Timeline' and ILP 'Status'

7. ILP DEMONSTRATOR RESULTS

Our ILP demonstrator system has two distinct modes. A hardware mode that encompasses a real balloon with gondola running the ILP software, and a balloon simulator. Both modes are interfaced to a demonstrator shell (Demoshell). Figure 9 shows laboratory trials with our ILP aerobot flying over a sample of mock Martian terrain that has been manufactured from $1m^2$ polystyrene blocks using MGS MOLA terrain data.

Using our aerobot flight simulator, 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, see Figure 10.

Our Demoshell provides a number of functions, ranging from the control of the balloon hardware and simulator, presentation of data gathered and eventual evaluation of the complete system and

ILP performance. To this end the Demoshell software encompasses enough flexibility to be able to efficiently and effectively perform all these tasks and experiments. Demoshell has to communicate with both the real balloon hardware, and the balloon simulator in order to receive up-linked DEMs and images and to download camera models and balloon trajectory information. The shell allows selection of the demonstration mode, hardware or software, provides an interface to allow selection and upload of camera, antenna and terrain models and enables the user to define the communication window function.

The Mamer Vallis crater region has been used for our investigations and a global MOLA derived elevation map is shown in Figure 11. We have 'flown' our ILP balloon over this region and a mosaic of cropped (non-overlapping) camera images is shown in the Demoshell window, Figure 12. The results from the relative localisation algorithm are overlaid onto this camera image mosaic (shown as white squares). Results from the absolute location algorithm have been obtained, and an example local DEM (created by the ILP from overlapping camera images) is shown




Figure 13. Example of a local DEM produced by the ILP using > 70% overlapping camera images. Local DEMs are input into the absolute localisation algorithm.




Figure 14. The ILP has identified (and colour coded) the corresponding terrain features for the DEM in Figure 13. The features and their respective gradients inside the cyan coloured circle are used by the absolute localisation algorithm.




Figure 15. The absolute localisation algorithm uses the global DEM data (shown in Figure 11) and this is processed for terrain features and gradients (shown here). The algorithm has found the absolute location of the local DEM shown in Figure 14, and its global position is indicated by the cyan coloured circle on the global terrain model.

in Figure 13. This has been processed for terrain features and gradients, Figure 14. Using a combination of terrain feature and gradient matching the location of the local aerobot obtained DEM region has been found in the global MOLA DEM data, Figure 15. Hence the aerobot has been autonomously localised in relative and absolute terms using our ILP methods.

8. CONCLUSIONS AND FUTURE WORK

An overview of the rationale for sending a LTA aerobot to Mars has been presented. Whilst there are many challenges, we believe that the technology ‘building blocks’ are available to launch a planetary exploration robotic balloon to Mars within the next launch opportunity windows (2009, 2011, 2013). We are using a novel combination of imaged based relative and absolute localisation techniques to ensure a robust solution to the problem of Martian aerobot localisation. We have demonstrated our ILP using both simulated and real aerobot data, thereby providing a comprehensive test and evaluation of our ILP functionality. Building upon our ILP research, we are examining also the application of a number of our developed imaging and DEM generation methods to tethered aerobot vehicles. Such technology could be integrated with future rovers and landers. In the context of a future Mars

sample return mission, for example, such a vehicle could provide invaluable landing site aerial surveying and meteorology prior to a fetch rover being sent to obtain a Mars surface and sub/surface sample. We will report on our tethered aerobot work in the future literature.

ACKNOWLEDGEMENTS

The authors would like to thank G. Visentin (ESA D/TOS), ESA Head of Automation and Robotics Section, and C. Hansen (IMT-CTM), Contracts Officer, for their collaboration with this work. This work has been funded under the ESA Contract No.17400/03/NL/CH.

REFERENCES

- Barnes, D. P., Phillips, N., and Paar, G., 2003. Beagle 2 simulation and calibration for ground segment operations. *7th Int. Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, CD-ROM proceedings.
- Barnes, D. P., Summers, P., and Shaw, A., 2000. An Investigation into Aerobot Technologies for Planetary Exploration. *6th ESA Workshop on Advanced Space Technologies for Robotics and Automation (ASTRA 2000)*, 3.6–5.
- Kolesnik, M. and Paar, G., 1997. Algorithmic Solution and Simulation Results for Vision-based Autonomous Mode of a Planetary Rover. *Proc. 7th Int. Conference on Computer Analysis of Images and Patterns*, 10–12.
- Matijevic, J. R., Bickler, D. B., *et al.*, 1997. The Pathfinder microrover. *Journal of Geophysical Research* 102, 3989–4002.
- Paar, G. *et al.*, 1994. Vision-based Navigation for Moon Landing.
- Shary, P. A., Sharaya, L. S., and Mitusov, A. V., 2002. Fundamental quantitative methods of land surface analysis. *Geoderma* 107(1-2), 1–32.
- Shaw, A. J. and Barnes, D. P., 2003. Landmark Recognition for Localisation and Navigation of Aerial Vehicles. *IEEE/RSJ Int. Conference on Intelligent Robots and Systems (IROS)* 1, 42–47.
- Squires, S. W. *et al.*, 2004. The Spirit Rover’s Athena Science Investigation at Gusev Crater, Mars. *Science* 305, 794–799.
- Wood, J. D., 1996. The Geomorphological Characterisation of Digital Elevation Models.
- Woods, M., Evans, M., Ward, R., Barnes, D., *et al.*, 2004. Developing an Autonomous Imaging and Localisation Capability for Planetary Aerobots. *5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, 509–515.