CREST Autonomous Robotic Scientist: Developing a Closed-Loop Science Exploration Capability for European Mars Missions

Mark Woods(1), Andy Shaw(1), Phil Rendell(1),), Ehsan Honary(1), Dave Barnes(2), Steve Pugh(2), Dave Price(2), Derek Pullan(3), Derek Long(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) *Dept. of Physics and Astronomy, University of Leicester, LE1 7RH, UK, dpu@star.le.ac.uk*
(4)*Dept. of Comp. Science, University of Strathclyde, Glasgow, G1 1XH, UK, derek.long@cis.strath.ac.uk*

Abstract

In common with most Mars missions, the current communications baseline for Europe's ExoMars Rover mission exhibits constrained data links with Earth, making remote operations difficult. The time taken to transmit and react to planning data places a natural limit on the amount of science exploration that can be achieved in any given period. In order to increase the potential science return, autonomous science assessment and response is an attractive option and worthy of investigation. In this work, we have integrated technologies and techniques developed in previous studies and used the resulting test bed to demonstrate an autonomous, opportunistic science concept on a representative robotic platform. In addition to progressing the system design approach and individual autonomy components, we have introduced a methodology for autonomous science assessment based on terrestrial field science practice.

1. Introduction

The forthcoming ExoMars Rover mission scheduled for launch in 2013 is the European Space Agency's (ESA) first attempt to deploy a mobile robotic platform on the surface of Mars. The mission has an Exobiology focus and will be used by scientists to search for signs of extinct or extant life on Mars [1]. The six wheeled mobile platform will weigh in the order 200kg and will have a dedicated science payload known as Pasteur which incorporates a range of panoramic, contact and analytical laboratory instruments. This will include a drill to gather sub-surface samples, ground penetrating radar, wide angle and close-up imaging capabilities and a robotic arm to deploy contact instruments. The nominal, 180 Sol science plan for ExoMars is to visit and investigate a number of primary sites which are in the order of 500m apart and gather samples for analysis through successive and iterative sample assessment. The same basic process will apply at each site i.e. broad site assessment, multiple target selection for deeper analysis, assessment, continued selection refinement and analysis through the range of position and imaging scales leading to eventual sample selection and characterisation.

Figure 1: ESA's ExoMars rover - artist's impression

At the time of writing it is expected that communication to the Rover will be established mainly via NASA's Mars Reconnaissance Orbiter (MRO) and a low-bandwidth Direct-to-Earth (DTE) link will also be available. The volumes of data required for science assessment and planning are significant and will strain the available bandwidth possibly introducing delays in the operations process. The nominal science operations plan described above has a number of discrete ground decision points where the science team must be involved e.g. in selecting a small number of candidate targets in a Wide Angle Camera (WAC) image for further analysis. Maintaining this science plan will of course be a challenge given the inefficiencies associated with the communications path.

One way of trying to sustain the proposed science schedule is to consider migrating some of the assessment and planning tasks from mission operations to the space segment in order to decrease decision making response times. This of course implies greater on-board autonomy. By carrying out autonomous science assessment and response planning on the robotic platform itself, the bottleneck associated with ground based decision making would be greatly reduced. However it was clear prior to the commencement of this study that the technology required to support this concept is not sufficiently mature. It is too early therefore to consider replacing ground decision points with an autonomous equivalent. There was however sufficient justification to evaluate the use of this approach for opportunistic field science and to offer support for lower order and automated data-gathering functions, particularly when anomalies occur. The overall goal therefore is to develop an approach which would increase the science return for such a mission.

2. Rationale and Objectives

The basic ExoMars science exploration model of repeated analysis at reasonably well seperated sites is well disposed towards opportunistic science. The terrain intervals between sites will be traversed “blind” from a science perspective in order to maintain the primary science schedule. Of course data will be gathered in an undirected way en-route and this could be used to alter primary plans if a target of sufficient interest presents itself. A better alternative would be to equip the plaform with the ability to autonomously assess the terrain with respect to science goals and if necessary plan a response which effects deeper analysis of some candidate target. Focusing on opportunistic approach serves three purposes. Firstly, the technology requirements for the science assessment which apply in this mode are less stringent than those associated with the current ground decision points as the intent is to improve on “blind” data capture. As long as tactical resource allocation and re-planning functions are available on-board, the risk to the nominal operations is negligible. It is also configurable in the sense that autonomous responses to assessment decisions can be constrained by mission operators in advance. Secondly, the solutions we develop for these problems will build a capability which could be used to augment or possibly replace the current set of ground decision points over the longer term for future missions. Finally, it is envisaged that low-level autonomous science selection and response re-planning could play an important part in recovering from certain anomalies during the nominal science schedule and prioritizing data for download. It is envisaged that this work will help develop capability which supports this concept.

A significant body of work has been undertaken in the US in the area of autonomous opportunistic science [2], [3], [4] & [5] for remote exploration. The aim of this work was to prototype methods which are directly applicable to the ExoMars current operations scenario which has its own unique attributes and constraints. As such it was part of the UK STFC CREST initiative which has sought to promote the development of technologies which will be of benefit to ExoMars like missions. This work of course may be complementary to wider research in this area. The objectives for this work were as follows:

- Establish an initial scientific methodology for the automation of science assessment and planning based on terrestrial field practice
- Prototype a system architecture which can support the concept of autonomous opportunictic science
- Prototype elements of the methodology provided by the science team in order to establish the feasibility of this approach
- Demonstrate the prototype system in a representative “Mars Yard” environment
- Use the forthcoming ESA ExoMars mission as a target and source of operations and science requirements

Our primary task was to demonstrate opportunistic science in a representative “ExoMars” environment. A reference scenario was outlined in the early phases of the study which provided a baseline for our work. The objective was to demonstrate that the mobile platform could traverse a rock field en-route to a target destination and both detect and respond to targets of scientific interest that it encountered en route. The response was to occur at two levels – first to authorize close-up or high resolution imaging of a target detected in a wide angle camera image – and second to plan and place a robotic arm if the target was sufficiently interesting. This required a number of key components including science assessment and response; re-planning and resource monitoring; and robotic arm approach and placement. In addition we required a basic system to support this scenario including ground based planning and on-board software elements such as timeline or plan execution.

In previous ESA sponsored work an autonomous science assessment capability developed for a Martian Aerobot prototype has been used as the starting point for the autonomous science aspect of this work [6]. The requirements for a rover based assessment system are very different so a new approach has been developed. The study team includes a planetary geologist and a key feature of this work is our attempt to define a framework from which to build a hierarchical scoring system for scientific evaluation based on fundamental geological features. Although ExoMars is the initial target mission for this work, the framework approach is generic and could be used by any surface element. The intent is to base this model on terrestrial geological field practice [11] whilst considering the constraints associated with robotic exploration [12]. The framework is used in turn as a basis for our autonomous science assessment and response models.

In another ESA sponsored activity the AI planning and scheduling based tactical re-planning software called Timeline Validation and Control (TVCR) has been developed to support goal based arbitration and timeline re-planning to support opportunistic science [7]. Recent work at AU provided both the test environment or “Mars Yard” and a half scale model of the ExoMars chassis E concept with a simple assembly representing a robotic ARM.

3. Test bed Framework

Figure 2 shows the architecture we have developed for the CREST demonstrator system, consisting of an autonomous science assessment component, closed loop approach and placement and an on-board planner and scheduler called TVCR. This component is at the heart of the autonomous science concept and will be used to reason about the suitability of servicing science operations requests generated by the on-board science component. The closed-loop approach and placement element will provide the basis for an autonomous implementation for more detailed science assessment request. The basic operations or usage model for the system is as follows:

- Nominal exploration timelines or plans are uplinked from the mission control centre (MCS)
- The rover executes the planned sequence which is mainly a traverse action between designated waypoints.
- At selected points the imagery collected during the traverse is assessed for science interest
- If sufficient interest is detected the science component will request a more detailed analysis via the executive and TVCR
- TVCR will assess the current plan, resource state and mission priorities before recommending a go/no-go for the new opportunistic science request
- The request may involve a close-up image activity or an actual ARM placement on a target object such as a rock or outcrop.

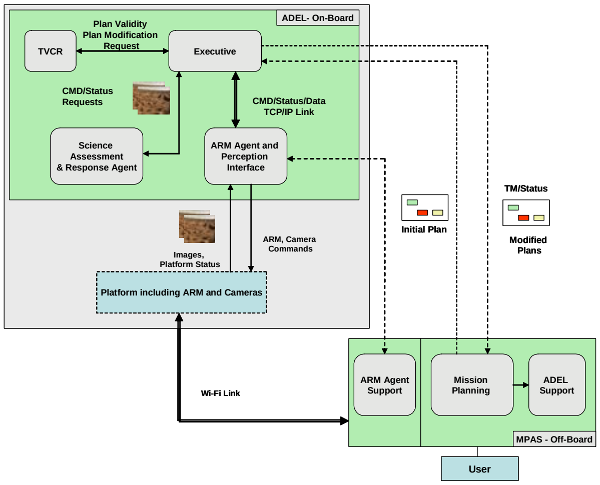

Figure 2: CREST system architecture

Each of the primary components namely Science Assessment and Response Agent (SARA), TVCR and the ARM agent is discussed in more detail in the following sections.

3. Science Assessment and Response

The SARA component is based on the underlying scientific scoring framework summarized below.

3.1 Science Assessment Framework

Effective geological fieldwork on Earth, Mars, or any other solid planetary body, relies on multi-thematic techniques (imaging, analytical and geotechnics) to unravel the often complex physical processes and geological history of an area being explored. Field observations rely on three fundamental criteria: structure, texture and composition. A first level scientific assessment can often be made using these basic attributes in conjunction with available contextual information such as regional geology/geophysics.

The ability of a human scientist or robot to identify and interpret geological criteria relies on feature/signature recognition (intelligent processing), a set of established rules (scientific weighting) and experience (learning). A methodology to evaluate the first two of these aspects has been devised in the form of a database of fundamental attributes and a numerical Science Value Score (SVS) system. Experiential learning will be a feature of a more advanced system involving a contextual model containing knowledge of the local/regional geology and the mission science log.

By adjusting the significance of the basic attributes both individually (i.e. simple laminar bedding) and in combination (i.e. complex cross bedding) one can assign an overall SVS to targets of interest. On the basis of the overall SVS one can then rank targets and respond accordingly (i.e. ignore primary target and proceed to more interesting target or some other predetermined action). Pre-assigned SVS values associated with features stored in the database can be tuned to the scientific objectives of a mission or phase within a mission. At this stage, the database is only populated with a modest number of feature types associated with each primary attribute. More detail and complexity will be incorporated as the system evolves and is systematically proven in field trails.

For this study our input data were restricted to greyscale images acquired by the rover camera system so structure and texture were the dominant attributes. However, extremes of grey scale and albedo provided a crude emulation of composition. To assist in the rover trials, realistic 3D renditions of geological outcrops exhibiting bedding features were crafted and integrated into the analogue terrain surface at AU.

3.2 Science Agent

The concepts outlined above have been used to create the basic SARA architecture outlined in Figure 3. The structure of the architecture is built around three levels of processing; candidate target area extraction, geological attribute analysis and final fusion of results factoring in geological context. The actual analysis applied depends on the scale at which the input image was taken. For example this basic process can be applied to WAC images showing whole rocks and high-resolution images which show close-up views of the rock faces.

In the first phase, input images are pre-processed in order to detect likely target candidates i.e. rocks or outcrops. To achieve this we relied on a number of basic techniques such as local thresholding, edge detection and region growing which proved sufficiently robust for our target environment. Robust rock and outcrop detection under a variety of conditions is of course not trivial and has been the subject of work by other groups [8]. Once regions of interest have been defined the next step is to carry out an attribute analysis. As Figure 3 indicates, attribute analysis is dependent on the detection and classification of primary features such as bedding or Albedo etc. The detection and classification of these features is complex and will require a long term effort to build up an extensive capability. Clearly the more features that are included the better the attribute classification.

In this instantiation we focused on developing algorithms which could detect a variety of bedding types and basic morphological analysis in order to assess target structure. For composition we considered Albedo and reflectivity. Finally for texture we sought to detect graded bedding

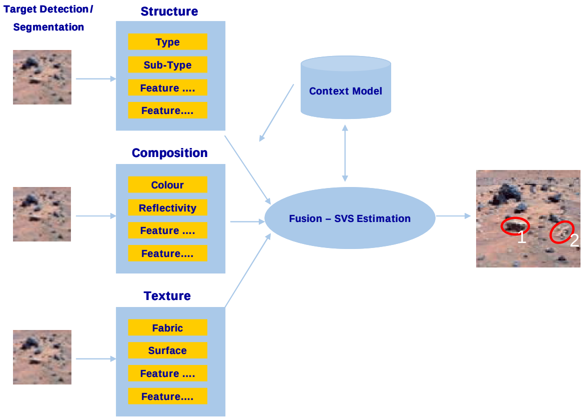

Figure 3: SARA Architecture

To detect linear bedding features we used a combination of edge detection and analysis to identify features such as layered, horizontal, vertical and cross bedding. Edge groups were assessed for orientation and size in order to build detect and characterize the linear bedding nature of each target. This assessment could of course be configured in order to alter the sensitivity of the assessment. Albedo values were estimated based on average greyscales values in areas of interest and appropriate scaling parameters. To detect graded bedding we used a circular filtering technique whereby segmented edge images were analysed for circular features and an estimation of the distribution of various scale classes used to confirm and score the degree of graded bedding present. Given the timescale of the work we did not model context in a sophisticated way so the final Science Value Score for each candidate target was simply a weighted fusion of the individual attribute scores.

4. Re-Planning Agent

Classically, in AI planning, replanning is the defined as the task of taking an existing plan, together with new goals, and constructing a new plan that achieves both the original and new goals. This problem is as hard, in general, as planning from scratch to achieve all the goals. In the context of the autonomous scientist, the replanning problem differs from this classical definition in several ways. Firstly, a characteristic of the task-oriented behaviours that drive the science missions is that it is natural to describe goals not in terms of the states that must be achieved but in terms of the actions that must be executed. This is because the science-gathering actions have only one purpose, which is to acquire the science that is the true objective of the mission, so it is natural to specify the goals of a scenario by specifying the actions that must be executed. Secondly, there are typically few choices in the actions from which the plans are to be constructed, but more choices about the ordering of the critical science-gathering actions. This makes the problem more like a scheduling problem than a classical planning problem. However, the order in which science-gathering actions are executed determines precisely which supporting actions must be executed to link these activities together into a coherent sequence. Thirdly, the most important constraint on the achievement of a successful replan is not the achievement of the goals, but the management of the limited resources.

These three factors contribute to shaping the replanning strategy in TVCR and in making a replanning strategy more efficient that to plan from scratch. In the autonomous scientist work, only the replanning functionality of TVCR is of significant interest: the other functions, including monitoring the execution of the timeline, remain active, but are not central to the focus of the research. To perform its replanning, TVCR reasons with a model of the activities that are available in the domain. This model is an action-centric model, providing declarative descriptions of the actions that the executive can be called on to perform. The model is written in Planning Domain Description Language (PDDL) [9]. The benefits of such a model are that it is relatively easy to extend with new actions, or to modify the descriptions of existing actions, in a form that closely corresponds to intuitions about the behaviours. Furthermore, PDDL is a widely used language, so its use offers access to a wide selection of existing planning and plan-manipulation tools.

The most important functionality on which TVCR depends is plan-validation [10]. This allows TVCR to take a timeline description, either supplied from the ground or partly constructed on-board, and validate it. The validation process identifies predicted flaws in the execution, indicating where in the timeline they occur and what is their cause.

A request to replan arrives at TVCR in the form of a fragment of a plan (a connected sequence of actions, but typically with unsupported preconditions) that is to be inserted into the timeline. The fragment has an associated priority value reflecting its estimated relative science value in comparison with the fragments that form the current timeline. TVCR handles the request by inserting it into a plausible slot in the timeline, identified by selecting the point at which the rover is located most closely to the site at which the fragment is to be executed, following other activities at the same site. Usually, this attempted insertion will generate plan flaws due to interactions between the newly inserted actions and the actions that were already on the timeline.

TVCR has a series of strategies available to it to handle the flaws that can arise. These include removing low priority actions when there is a shortage of resources such as power or data storage, delaying actions when there is a conflict between demands on fixed resources (such as instruments) and adding in support actions to ensure the coherence of the execution trace of the timeline. TVCR applies these strategies in a fixed order. The resolution of some flaws can create new ones and these are handled alongside the others. If the process of resolution fails to generate a valid timeline within a fixed number of iterations, TVCR falls back on a fail-safe strategy of stripping down the original timeline until the activities that remain fit within the resource envelope available and represent an executable plan. In this way, TVCR can ensure that there is always a valid timeline awaiting execution (although it might be empty) and can add opportunistic plan fragments to the timeline as they arise, while respecting the constraints on resource use and on correct execution of the plan.

5. Robotic ARM Approach and Placement and Platform

Key hardware components within in the autonomous Arm Agent and Perception Interface

(AAPI) of the CREST architecture are the panoramic cameras and zoom high resolution camera (PanCam), the PanCam pan and tilt mechanism, the robotic arm, and the locomotion chassis with associated on-board computer and electronic interfaces. Figure 4.

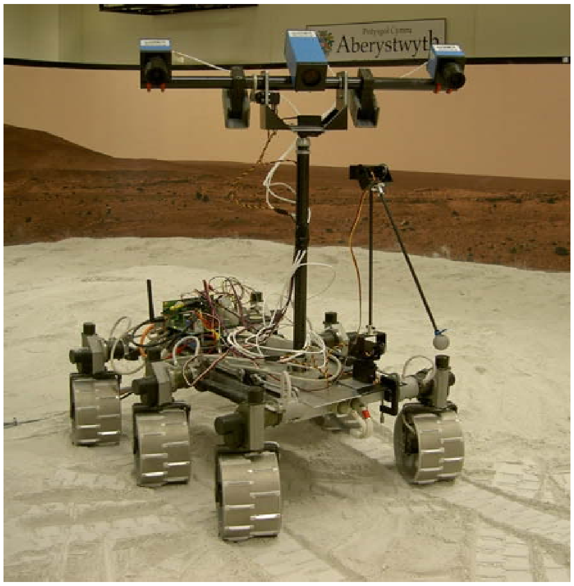

Figure 4: Shows a close-up of our demonstration rover platform

Upon instruction from the Executive, a stereo image pair is captured using the PanCam wide-angle cameras (WACs). Our demonstration scenario required an overlapping sequence of image pairs to be captured via the autonomous operation of the rover's PanCam pan and tilt mechanism. SARA then examines one image (typically the left-hand image) from each image pair.

Upon identification of a science target, the image pixel coordinates of this object (e.g. the rock's centroid) is communicated to the AAPI. SARA can request a zoom image of the rock in question to confirm a science target hypothesis. The AAPI accomplishes this by calculating a correction to the PanCam pan and tilt mechanism orientation so as to centre the science target in the zoom camera's field of view. A rock zoom image can then be captured. Using the science target WAC captured image pair and stereo triangulation, the 3D position of a science target is calculated relative to the rover. This allows the pan and tilt orientation to be calculated for the zoom image capture activity, and allows a science acquisition 'cost' to be calculated. This is based upon the power and time that would be required for the rover to traverse to the science target location.

The cost information is used by the architecture TVCR to assess the resource implications of an opportunistic science activity. If a 'go' is given, a rover traverse can occur to place the science target within arm's reach. SARA again requests a WAC image pair, performs a final science target assessment, and notifies the AAPI as to the image pixel coordinates of the final science object. The AAPI now uses its stereo triangulation and arm kinematics to confirm target reachability, and calculates an appropriate arm configuration for instrument placement.

A safe instrument placement trajectory and contact region on the science target can also be determined by generating a mini-DEM (Digital Elevation Model) of the science target. Arm placement costs are also calculated and if a final 'go' is issued by the TVCR via the Executive, then the arm (with instrument) is moved and science target contact is made.

The autonomous AAPI has required a number of algorithms to be designed and implemented including camera image distortion correction and rectification, disparity map generation, stereo triangulation, pan and tilt mechanism pointing, science target acquisition cost calculation, arm reachability, and safe instrument placement trajectory determination and execution. Our current research is focused upon an FPGA based implementation of the AAPI functionality.

7. Trials and Results

Figure 5 shows the layout of the nominal scenario. The planned activity for this simulated Sol was to traverse from A to B and carry out periodic assessments of the terrain at discrete waypoints using acquired WAC images. The assessment is therefore carried out at proximal scale. The platform passes a small rock field to its left (near centre line of figure) which are not of science interest. As nears it nears WP3 there are a set of rocks at 10 O'clock from its position. Some of these are of high scientific value as they exhibit layering features including both horizontal and cross bedding. The challenge therefore was to ignore some rocks, detect those of interest and escalate the opportunistic analysis to include high resolution imaging and arm placement.

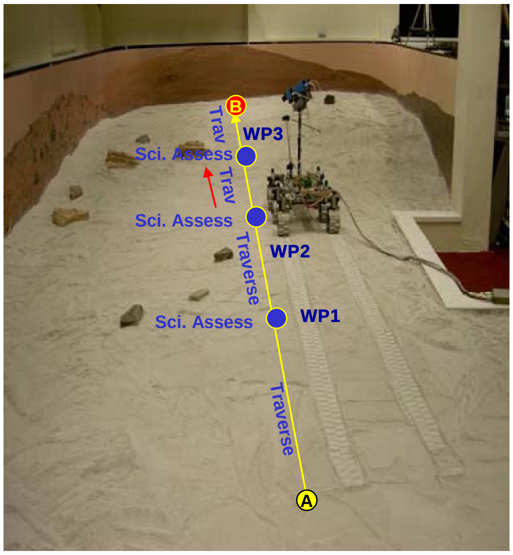

Figure 5: Demonstration scenario showing nominal route and opportunistic deviation

Figure 6 shows the how the nominal plan evolved in the course of the experiment. Images at WP2 were assessed but not deemed to be interesting enough. At WP3 a morphological analysis was carried out on an input WAC image The resulting analysis showed a high concentration of peak and ridge features which caused the SARA component to generate a request to take a high resolution image. This was assessed and authorised by TVCR.

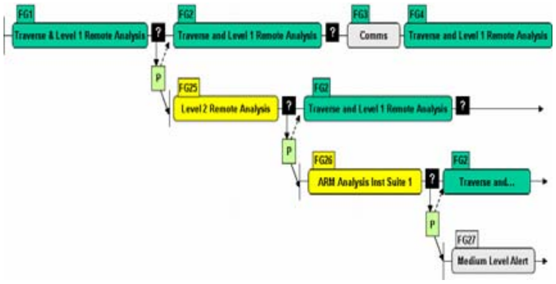

Figure 6: Nominal plan is on the first line at the top. Branches inserted as a result of opportunistic analysis and requests are shown in lines two and three.

Figure 6 shows the resulting plan deviation where a level 2 analysis fragment has been inserted into the nominal plan. Once executed, SARA carried out a structural analysis on the high resolution image.

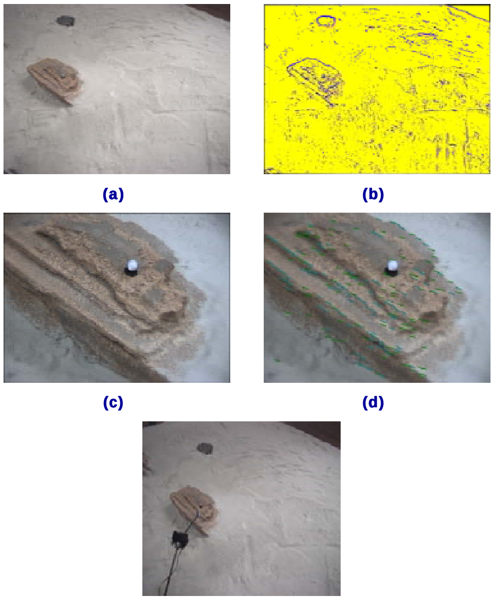
(a) (b) (c) (d) (e)

Figure 7: (a) WAC input image (b) processed morphology analysis (c) high resolution input image (d) bedding analysis output (e) showing successful arm placement. Not: Spherule for illustration only

The output shown in Figure 7(d) shows a significant concentration of horizontal bedding. SARA requested a more detailed analysis of the target via the executive. The executive queried the ARM agent to get an approximate measure of the resources required to conduct an ARM placement. It then commanded TVCR to parse this request and determine its feasibility. TVCR determined that the request could be serviced and the appropriate plan fragment was inserted in the plan. Figure 7(e) shows the resulting ARM placement. Figure 6 shows the various branch stages that occurred as a result of re-planning.

7. Conclusions

Implementing science autonomy for a robotic exploration platform is a challenging but attractive goal. It requires a comprehensive system design which co-ordinates and combines a number of complex components. This work had shown that such a system concept can be developed for a mission such as ExoMars which evolves the current operations paradigm in a practical way. There are three generic elements that are required to build this system. Firstly there is the science assessment and response. We have outlined a framework for the development an approach to autonomous field science based analysis of fundamental geological attributes. The SARA implementation demonstrated here has addressed only the basic elements of this framework and clearly the evolution of this capability is a long term prospect. However we believe that it can be deployed in an incremental way as the concept evolves. The next step would be to develop this element further in a stand-alone form.

The second critical component is a plan validation and re-planning element. This study has demonstrated the suitability of the TVCR concept for this particular task. Such a component may be used on ExoMars to support timeline optimization. It is beneficial to know therefore that it could also support autonomous science should it be deployed. The third generic component is and autonomous control of active platform elements such as a robotic arm or locomotion system. In this work we have demonstrated a concept for autonomous approach and placement of a robotic arm. In conclusion we believe that although challenging, science autonomy in some form will be essential for missions such as ExoMars.

The need to optimize the collection, selection and transmission of science data could greatly benefit from the use of such technology and may ultimately be required to ensure nominal science goals are achieved.

9. Acknowledgements

We would like to extend our gratitude to UK STFC represented by Mrs. Sue Horne, who funded this work as part of the UK CREST programme.

10. References

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

[2]: R. Castano, M. Judd, T. Estlin, R.C. Anderson, D. Gaines, A. Castano, B. Bornstein, T. Stough, K. Wagstaff, “Current results from a rover science data analysis system”, In *Proc. of the 2005 IEEE Aerospace Conference*, Big Sky Montana, US, 2005

[3]: R. Castano, T. Estlin, D. Gaines, A. Castano, B. Bornstein, C. Chouinard, R.C. Anderson, and M. Judd, (2006) Automated target selection for opportunistic rover science, *LPSC XXXVII*

[4]: D. Thompson, T. Smith & D. Wettergreen “Data Mining During Rover Traverse: From Images to Geologic Signatures”, In *Proc. 8th Int. .Symposium on Artificial Intelligence, Robotics and Automation in Space,* Munich, Germany, 2005

[5]: D. Wettergreen, N. Cabrol, S. Heys, D. Jonak, D. Pane, M. Smith, J. Teza, P. Tompkins, D. Villa, C. Williams, M. Wagner, A. Waggoner, S. Weinstein, W. Whittaker. “Second Experiments in the Robotic Investigation of Life in the Atacama Desert of Chile”, In *Proc. 8th Int. .Symposium on Artificial Intelligence, Robotics and Automation in Space,* Munich, Germany, 2005

[6]: D. Barnes, A. Shaw, P. Summers, R. Ward, M. Woods, M. Evans, G. Paar, M. Sims, Autonomous “Image Based Localisation For a Martian Aerobot”, in *Proc. 8th Int. Symposium on Artificial Intelligence, Robotics and Automation in Space,* Munich, Germany, 2005

[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

[8]: Thompson D., Castano R., “Performance Comparison of Rock Detection Algorithms for Autonomous Planetary Geology”, In *Proc. 2007 IEEE Aerospace Conference*, Big Sky Montana, US, 2007

[9] M. Fox and D. Long, "PDDL2.1: An extension to PDDL for expressing temporal planning domains", Journal of AI Research, 20, 2003

[10] M. Fox, R. Howey, and D. Long, "Validating Plans in the Context of Processes and Exogenous Events", Proceedings of The 20th National Conference on Artificial Intelligence (AAAI-05), 2005

[11] D.A.V. Stow, “Sedimentary Rocks in the Field”, Manson Publishing.

[12] D. Pullan, F. Westall, B.A. Hofmann, J. Parnell, C.S. Cockell, H.G.M. Edwards, S.E. Jorge Villar, C. Schröder, G. Cressey, L. Marinangeli, L. Richter, and G. Klingelhöfer, (2007) “Identification of Morphological Biosignatures in Martian Analog Field Specimens using *In Situ* Planetary Instrumentation”, *Astrobiology*, In press.