# CREST Robotic Scientist

**Andy Shaw[1], Mark Woods[1], Ehsan Honary[1], Philip Rendell[1], Derek Pullan[2], Dave Barnes[3], Steve Pugh[3], Derek Long[4]**

1) SciSys
Clothier Road
Bristol BS4 5SS, U.K.
Mark.Woods@scisys.co.uk

2) Space Research Centre
Dept. Physics and Astronomy
University of Leicester
dpu@star.le.ac.uk

3) Dept. Computer Science
UWA, Aberystwyth
Ceredigion, SY23 3DB.
dpb@aber.ac.uk

4) Dept. Computer and Information Sciences
University of Strathclyde,
Glasgow, G1 1XH

### *Abstract*

Due to the success of the recent NASA interplanetary space mission MER (NASA, 2007) it has highlighted the importance of using a roving science platform for exploration. Near and long-term requirements for future interplanetary missions place increasing demands on rover performance to extract maximum benefit from the large effort and funding committed to such missions. Missions are more diverse in their science objectives and require improved robustness and reliability over longer distances during surface operations.
To keep pace with these complex and evolving requirements it is essential that the level of autonomy used on future missions be increased in order to improve the responsiveness of historical operations models which are biased towards an open-loop response for high-level anaylsis and decision making. The driver of the CREST rover development initiative is the need to achieve: More accurate delivery of science instrumentation, New science opportunities, Large increases in the science data returned to Earth, More robust and reliable operations and More efficient use of operational resources. This paper presents our work and results obtained to date.

Keywords: autonomy, planning, opportunistic, rover

## 1 Introduction

The development of more sophisticated vision-based Navigation algorithms such as those planned for ESA's ExoMars rover mission will allow rovers to cover greater distances giving rise to more science opportunities. However the current reliance on open-loop science assessment and planning means that responding to these opportunities and carrying out nominal mission activities such as sample site selection is slow. Executing the engineering activities associated with science requirements in this way also slows overall mission progress. For example, open-loop approach and placement for the NASA MER rovers currently takes in the order of three days.

Instrument placement through geological feature analysis can reduce sample/measurement acquisition times to less than one day (Huntsberger *et al.,* 2005) thus enabling higher science return. The Autonomous Robotic Scientist addresses the need for greater rover autonomy in the areas of, data analysis, science planning and approach and placement. At a more abstract level, providing remote exploration rovers with the ability to detect targets of scientific interest on an opportunistic basis and rank possible sample site areas would both enable new science activity and improve the turn-around time for nominal mission goals. This work is being carried out within the context of ESA's ExoMars rover mission. Although a significant body of work has been undertaken in the US (Castano *et Al., 2005* ) (Castano *et Al., 2006* ) (Thompson *et Al., 2005)* our aim is to prototype methods which are directly applicable to the ExoMars current operations scenario which has its own unique attributes and contratints.

In order to advance these goals, the Autonomous Robotic Scientist aims to:

- Establish an initial sceintific methodology for the automation of science assessment and planning based on human field practice
- Prototype a system architecture which can support the concept of autonomous science

- Prototype elements of the methodology provided by the science team in order to establsih 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

There are a number of key technical challenges in this work which will be impossible to completely resolve in the short one year duration of the CREST activity. In some instances the possibility of developing a complete engineering solution is still an open research question. However we do aim to establish the suitability of our proposed architecture, its core components and the scientific methodology itself. It is envisaged that the capability of the system will evolve with subsequent reasarch and development.

Figure 1 shows the architecture we have developed for the CREST prototype demonstrator, 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 deliberate over the suitbaility 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 excutes the planned sequence which is mainly a traverse action between designated waypoints.
- At selected point s 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 opportuinistic science request
- The request may involve a close-up image acticity or an actual ARM placement on a target object such as a rock or outcrop.

Each of the primary components is discussed in more detail in the following sections.

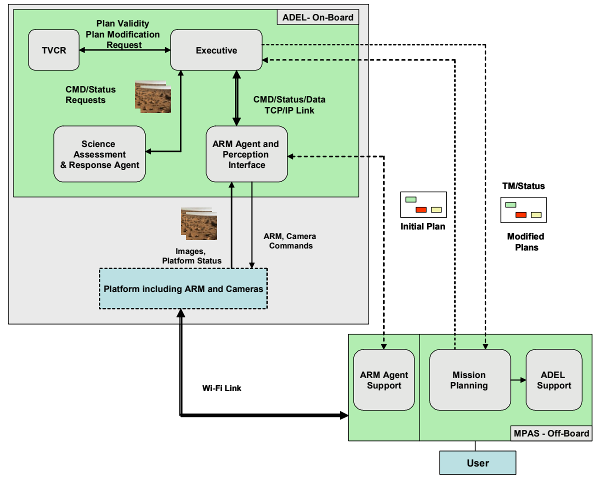


Figure 1: CREST system architecture

## 2 Planning

Presently mission planning is conducted by the system engineers and scientists on Earth, but due to the bandwidth limitations and time delays associated with interplanetary communication “up-to-date” information is not guaranteed. As a consequence a conservative scheduling approach is adopted as it is often difficult to construct robust timelines at the task level (producing recurrent safe modes and downtimes) resulting in a reduced science return.
With the introduction of an autonomous planner/scheduler on-board the rover system can alleviate the inefficiencies previously mentioned. Having direct access to the on-board state can result in greater science return though the adaptation of the current plan when anomalies have been detected. Through the use of AI planning and scheduling technologies these functionalities can be realized.
The planning system GUI, is used to construct initial schedules. Detailed plans (fragments) are constructed with this tool. A collection of these customized fragments form the entire plan. In addition, a number of constraints are also defined that present a logical relationship between these fragments. For example, one fragment may only take place after the execution of another. The designer of the plan also creates a number of opportunity fragments. These can be accessed and used by the TVCR (explained below) to repair a plan or extend it based on the recommendations made by the Science Agent.

### *2.1 System status*

A previous study called Mars Mission On-board Planner and Scheduler (MMOPS) (Woods *et al.*, 2006) showed the technology readiness of AI systems especially for

Timeline Validation, Control and Repair (TVCR). During current mission if an error in the timeline is detected the system reverts to a safe idle mode and waits until further communication with ground resulting in the loss of all subsequent tasks. Our system figure 1, can presently repair a time line if an error has been detected, by either replacing the problem task with a “standby” task (only if system resources allow and subsequent task are not effected) or by removing the problem task and continuing with the subsequent tasks, resulting in science return.

### *2.2 TVCR*

TVCR is a PDDL based planning and scheduling service which is responsible for the validation and repair of the plan. TVCR receives a plan and attempts to validate it. It uses the opportunity fragments to repair the plan based on new state information. It bears some similarities with the EO-1 CASPER planning tool although TVCR also supports a more continuous model during validation.

Figure 2, shows the TVCR architecture, the validation service uses a model of the activity sequence, this contains preconditions that must hold if an activity is to execute successfully and also the effect on the system state after completion. As sensor data from the system is fed in to the TVCR and interpreted via a system model, allowing comparisons between observed and predicted states during task execution. This is handled by the plan execution monitor. A plan failure flag is only set when the state violates the completion of the plan condition and not from slight deviations in predicted and actual.

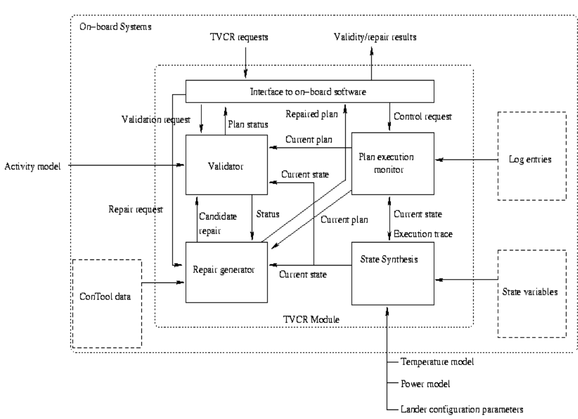


Figure 2: TVCR Architecture diagram, showing the links between the validator, execution monitor, state synthesis and repair generator.

If an error is detected, the TVCR calls the repair module, where it uses the models of the activities to predict their inclusion in the plan and the current state. These are fed back though the validator to assess candidate repairs, this is achieved through iteration as fragments that are broken due to the current state are removed and are potentially replaced with fragments that can still be executed (the extra information provided in addition to the plan). If no substitutions can be made, what is left of the original plan is executed.
As different tasks require the system to be in a certain configurations (instrument placement using a robotic arm) the TVCR automatically calculates short, special purpose moves required to move between configurations as these will also effect the timeline and resources.

### *2.2 Goal task prioritization*

During any plan reshuffle, any extra tasks are sorted with respect to priorities associated with them during the initial plan construction. These are often characterized using the scientific return and potential system resources consumption. An important feature of the CREST TVCR is the ability to adapt to incorporate new opportunistic science goals that may arise. Using an inbuilt overall mission goal and task prioritization tree any new opportunities can be integrated into the plan with the removal/postponement of lower priority tasks.

### *2.3 Fixed tasks*

The TVCR can not move specific tasks that must occur at fixed times, such as communications or when specific experiments are supposed to be conducted at a fixed time.

### *2.4 System simulation for resource validation*

As mentioned, the on-board system is capable of performing accurate self simulation to obtain resource management. Monitoring internal sensor information allows real-time modeling. This includes solar array power generation, surface traverse power consumptions etc. giving up to date prediction models that can be fed into the plan.

## 3 Data Analysis

As the main form of navigation for present rovers comes from the on board cameras, it is imperative that the quality of each image is prefect. As imperfections can lead to spurious results and misclassifications etc. but also the precious download bandwidth could be wasted. As part of the autonomous scientist image acquisition process it has to analyse each image for the following parameters.

### *3.1 Image quality*

When an image is taken it initial passes through the calibration phase, here the image has any lens distortion removed by applying a *pixel transformation matrix*. The images also pass through a *colour correction process* (generated from a colour chart) to balance the pixel levels.

After calibration the image is then checked for several standard properties;

- Exposure: not under or over exposed depending on scene luminance.
- Contrast: object distinguishability.
- Focus: sharpness of image when using a zoom camera.

If any errors are flagged in the images a “request” for new image is issued, along with the corresponding correction factor. This image is again processed for quality and if it is now deemed correct it is accepted, if on the other had a similar image is obtained even though the correction factor was applied the image is “tagged” for inspection.

### *3.2 Image richness*

Once an image has passed the quality control process the actual content can be analysed. Through image processing various geological features are identified, these include rocks, outcrops, soil types etc.
As each image is given an, image richness index the compression and importance values can be obtained and its communication priority can be calculated. If an image contains several features its compression is lower and priority higher than an image with a low feature count whose compression is higher and priority lower.

## 4 Autonomous Field Science

For purposes of rover navigation and locomotion, terrain features such as rocks, outcrops, slopes, unconsolidated sands etc are potential obstacles to either avoid or traverse over. Such features may or may not have additional scientific value but all share fundamental physical criteria such as size, shape and degree of consolidation. By recognising, evaluating and “scoring” these and other basic parameters from a scientific perspective, one has the basis on which to empower the rover with autonomous scientific reasoning.

### *4.1 Robotic Geology*

In the terrestrial context, the search for exploitable resources such as oil, gas, water, minerals, geothermal energy etc, although specific objectives, heavily relies on an initial understanding of the fundamental geology of the region being explored. Prior to any field campaign it is important to accumulate all previous data in order to establish local and regional context. On Earth this is achieved via survey data (including geological maps, satellite remote sensing, geophysical surveys and analysis of samples collected on previous expeditions). On Mars (and other planets), orbital data from previous missions are likely to be the only source of contextual information prior to landing although some ground truth (albeit inferred) may be available. Surface missions to new sites therefore have to undertake basic site investigation *in situ* with whatever payload assets are available, usually a restricted suite of instruments. Once the landing site has been characterised, human scientists can then place detailed observations into appropriate context. Robot systems should adopt the same approach during autonomous sessions by utilising feature recognition algorithms, a rule-based scientific scoring system and a working contextual model as part of the verification and “learning” process.

### *4.2 Contextual Model*

Assessing the science quality of image data will require a detailed assessment of the primitive geological features such as composition, structure and texture for the range of terrain identified in a scene. Once this problem has been addressed satisfactorily then the complex task of establishing context must be examined. Scientific context is an essential element in helping to make a final assessment and understanding of the area under survey. Ultimately it is the context which helps establish a human expert level of understanding of scene. The science team have identified essential feature extraction and analysis as the focus for this one year study however the role of a context model has also been assessed as part of this work

The contextual model describes the geological environment in which the robot “explores”. For example, if the rover landed in a volcanic region then the contextual model would start off with a presumption that the rocks should be volcanic. If the robot subsequently observes fine layering then this would be assumed to be due to ash deposition. On the other hand, if the landing site was a lacustrine deposit or an aeolian sand sheet then different interpretations would result.

In addition to being guided by geological context we need to bias autonomous decisions toward the mission objectives. So if a robot searching for evidence of water cannot decide which of two candidate targets should be considered prime and following autonomous assessment one is shown to contain hydrated minerals and the other is a basaltic ash, then a positive bias would be applied to the former. To take things further if during a mission a rock turns out to be comprised of carbonate, this would be treated as a “discovery” (assuming it had yet to be found) and subsequently override any mission objective. The contextual model could also be used to recognise rare rock types found previously during the mission during opportunistic science excursions.

### *4.3 Target classification*

Geological features often appear complex and are influenced by a huge number of variables. In the field, human geologists mentally deconstruct what they see and draw on broader contextual input (the bigger picture) to help classify geological materials and the processes that act on them.

The basic attributes used in field classification of rocks and soils are structure, texture and composition. The CREST Robotic Scientist system also uses this fundamental concept. Once data have been processed by the quality and richness routines, they are then passed on to the autonomous science agent for analysis. Feature recognition routines scrutinise the data and compare with pre-defined features specified in a scientific attribute database. The result is a cumulative science value score which can be used for autonomous science decision making.

It is unlikely that an adequate scientific evaluation could be made using single attributes in isolation even though there may be cases where this might apply. Nevertheless, it is appropriate to first consider each attribute independently by performing a feature recognition assessment of the target. Feature lists associated with each attribute are pre-defined and allocated numerical Science Value Scores (SVS) based on relative significance.

The total SVS derived for each attribute is only an indicator of "feature richness" and does not necessarily reflect the overall SVS of the target. The SVS of the target itself is derived by evaluating all the matching features for each attribute and biasing the assessment using the contextual model.

In general terms the SVS of the target is a function of a number of derived parameters:-

$$SVS = f\left(As, At, Ac, Q, B\right)$$

where;

*As* is the overall structural attribute score
*At* is the overall textural attribute score
*Ac* is the overall compositional attribute score
*Q* is a quality factor
*B* is a bias factor

Note that the quality and bias factors *Q* and *B* are intended to enhance or diminish the overall score in much the same way a human geologist may apply these criteria in the field.

Unambiguous interpretation relies on iteration since features seen from afar often look very different when viewed close up (sometimes unexpectedly so). This emphasises both the importance of detailed close-up observations and the need to incorporate re-evaluation into the onboard autonomous routines.

### *4.3.1 Structure*

The most obvious structural form is layering or stratification, a term used in reference to sedimentary rocks but it can be applied to volcanic and metamorphic deposits exhibiting layered structures. Where thickness is implied, units display either "bedding" (> 1 cm) or "lamination" (< 1cm). At all scales the basic geometric parameters are the same. Depending on the material, bedding can often be readily identified remotely, especially if enhanced due to differential weathering and natural illumination geometry. Closer up, thin beds and laminations sometimes require additional aids such as controlled illumination and surface preparation (i.e. splitting, grinding and sometimes polishing). Combined geometries can be very informative and in some cases have high scientific value. Figure 3 shows two examples of rock structures at different scales.

Figure 3: Examples of geological structures at different scales. Left: Bedding features of two geological units with an undulating interface. Note the compositional variation within the lower unit and between upper and lower units. Hammer approximately 25 cm. Durness, NW Scotland. Right: Ancient sedimentary structures in a hand specimen equivalent in age to Martian Noachian rocks. Note the cross bedding feature at the top left portion and the lensoid feature centre right. Field of view approximately 5 cm across. Pilbara, Western Australia. Images courtesy of GSPARC and the Planetary Analogue Field Studies Network (PAFS-net).

### *4.3.2 Texture*

The textural properties of rocks are dependent on particle grain size and distribution, grain morphology and overall fabric (how grains are orientated and packed). Although these properties can only be determined macroscopically, some generic aspects are applicable to remote observation of larger geological features. Figure 4 shows two examples of rock textures at different scales.

Figure 4: Examples of geological textures at different scales. Left: Basaltic lava (few days old) showing characteristic ropey texture (pāhoehoe). Kilauea, Hawaii,

USA. Field of view approximately 3 m across. Right: Hematite concretions and remnant casts in iron oxide deposit. Jura, Switzerland. Field of view approximately 5 cm across. Images courtesy of GSPARC and the Planetary Analogue Field Studies Network (PAFS-net).

### *4.3.3 Composition*

The geochemical and mineralogical make up of rocks is perhaps the most demanding of attributes to define. Weathering and alteration processes can subtly or radically change both the chemistry and/or mineralogy of rocks and soils so there is much reliance on contextual data to assist in the interpretation of analytical measurements. Thankfully, initial clues regarding composition can be obtained from imaging data. In fact target selection must be made on the basis of remote sensing involving a combination of spatial and spectral imaging techniques since only close up surveys of pre-selected targets will benefit from analytical measurement. Figure 5 shows two examples of compositional variation at different scales.

Figure 5: Examples of geological compositional variation at different scales. Left: Bleaching in aeolian sandstones due to mobilisation of iron. Distance to first redox interface approximately 300 m. Valley of Fire, Nevada, USA. Right: Olivine phenocrysts in vesicular basalt. Kilbourne Hole, New Mexico, USA. Field of view approximately 8 cm across. Images courtesy of GSPARC and the Planetary Analogue Field Studies Network (PAFS-net).

### *4.4 Example scenario*

Consider a rover stationed at an initial waypoint and three pre-selected candidate targets (A, B and C) located some distance away on higher ground (Figure 6). The primary objective for the mission is the search for life (i.e. ExoMars) and the contextual model describes the current location as a volcanic plain with debris fields adjacent to ridges. The targets were selected from wide angle and high resolution stereo survey data on the basis of the following:-

A = Small vertical promontory, dark (black), mottled texture

B = Large boulder, smooth, rounded, high albedo, otherwise featureless

C = Potential outcrop, bluish, slight evidence for bedding

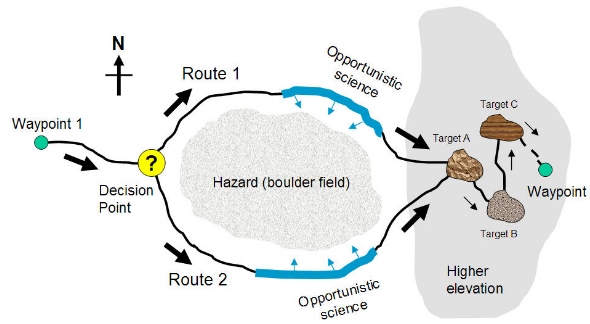


Figure 6: An example traverse with opportunistic science en-route. Scale is arbitrary but the distance between waypoint 1 and waypoint 2 could be say 10 m in the case of a planetary yard or 50m in the case of a field locality.

Each target (A, B or C) is a potential final waypoint for this excursion and the final choice will depend on the outcome of close-up scientific assessment and engineering requirements made autonomously. The sequence is A-B-C. Waypoint 2 is defined by the navigation team as a suitable intermediate waypoint in the case of an unresolved outcome (i.e. the robot cannot decide between targets on grounds of both science and engineering). Otherwise the rover will remain at or return to the prime target following scientific scrutiny.

Two routes are considered viable in order to reach the location of the targets due to the presence of a large boulder field. Route 1 tracks to the North and route 2 tracks to the South. The decision which route to take is autonomously is confirmed at the new vantage point (decision point). Remote sensing from waypoint 1 suggests the boulder field to be of no scientific interest. However, an optional wide angle/high resolution imaging survey of the hazard at the decision point may enhance a decision based on navigation/engineering grounds only. Whichever route is chosen the rover is programmed to initiate a localised survey along occluded regions of the boulder field as indicated. In other words the rover goes into “opportunistic science mode” along segments of route 1 or route 2 as shown. In addition, the rover’s “brief” on this occasion is to restrict activities to imaging and only store data that pass the feature recognition algorithms. An extreme “discovery threshold” is assigned whereby only a high opportunistic SVS would cause the rover to curtail the traverse and await instructions. If this does not happen then the rover proceeds to the target locality and performs close-up surveys of each target.

Possible scenarios and outcomes that could be demonstrated using this example include:-

- A salt deposit is detected halfway along route 1 (rover stores data and proceeds to target locality)
- A carbonate is discovered near end of route 2 (rover safes itself and waits for instruction)
- Target C displays fine cross bedding (rover stays at C)
- Target A, B and C turn out to be the same volcanic material as waypoint 1 (rover proceeds to waypoint 2)
- Target A is confirmed as an outcrop and is rich in hematite (rover returns to A)

These scenarios form the basis of future verification experiments that will be conducted initially within the Planetary Analogue terrain Laboratory (PATL) at the University of Wales, Aberystwyth.

## 5 Experimental System and Results

Consultation with the science team had revealed a strong desire on their part for the system capability to be developed and validated incrementally in a bottom up type approach. In practice this means developing algorithms which can detect and recognise each feature that partially defines the three main attributes.

Figure 7 shows the individual rover system components (rover chassis, camera mast and manipulator) at the PATL. This system will form the base from which all "real" experiments will be conducted. The soil (DLR Mars Soil Simulant-D) has the same physical properties as those experienced on the Martian surface and together with both rocks (not in image) acting as obstacles and for corer/grinding purposes, and a sample of "science rocks" (which have undergone laboratory analysis) will generate a realistic Martian environment and operation scenarios.
The experimental camera setup consists of two fixed focus $60^0$ FOV cameras (PanCams) and a third which is variable zoom and focus. Figure 8 shows some images from the initial analysis system at work, here we see the system capturing a series of images from the PanCams with various defects; under exposure, over exposure etc. each time the system calculated the problem and retook the image with a correct camera setup, the test was also carried out on the zoom camera along with focus defects.

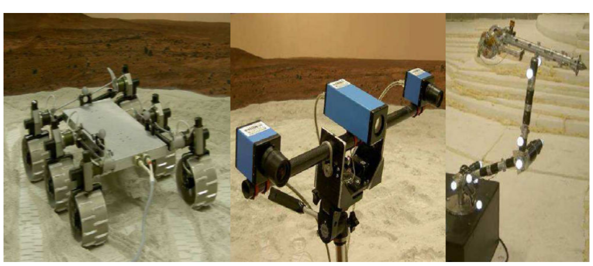

Figure 7: Rover chassis, panoramic and zoom camera with pan/tilt mechanism and manipulator. Images courtesy of UWA.

Figure 8: Initial results from the image quality testing phase, left image show under exposure, right image shows over exposure and centre image shows correct exposure. The images were taken in the PATL and contain the test science rocks.

When an image is obtained that passes the quality criteria it is then analysed by the image richness and feature extraction functions. Figure 9 shows a test image that has been used during initial experimentation; this shows synthetic scientific rocks in the UWA PATL. Initially the image is segmented to extract the rocks and features from the soil. The method of segmentation used is one of graph based (Felzenszwalb, 2004) smoothing watershed, the result from which was used along side a geological analysis (figure 10) to select any potential targets of scientific interest. In this example the result can be seen in figure 11.
Figure 12 shows the results of the structural and textual investigation of the selected rock. From the results the rock was classified as being planar bedding with horizontal layers, as shown by the line markings.

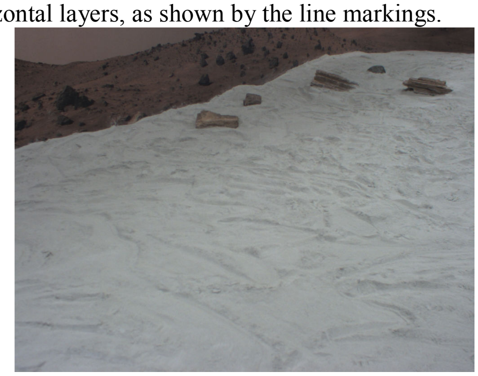

Figure 9: Image of the UWA PATL with realistic scientific rocks, used for testing.

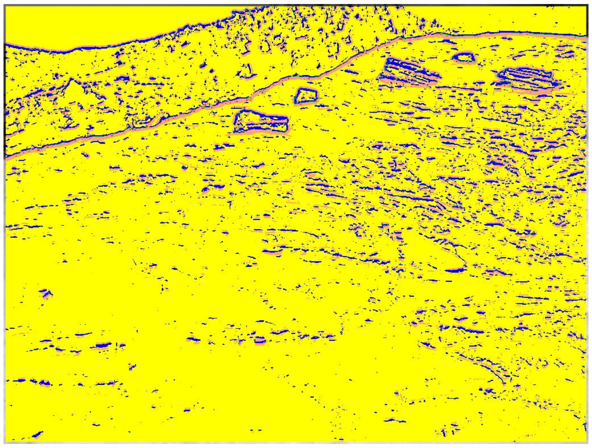

Figure 10: Geological feature analysis of the UWA PATL.

Figure 11: Potential candidate selected for further investigation. This represents the rock on the right hand side in figure 9.

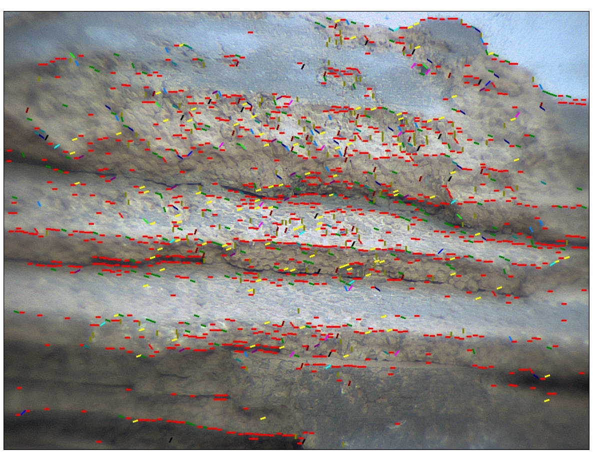

Figure 12: Close up of the selected rock, analysed for structural composition. The result showed planar bedding, primarily in the horizontal direction.

## 6 Conclusion

This paper has presented the ongoing work for the CREST autonomous scientist. The underlying model of the system architecture for planning and scheduling has been generated. This includes the architectures for the user interface, automated system planner, rover control architecture, Opportunist Science Agent, camera positioning, image quality analysis, image segmentation and object classification. Initial experiments in the planning, system status monitoring, image capture and quality control have yielded promising results. The mechanism for the object classification is now being implemented.

## Acknowledgments

We would kindly like to acknowledge the support of PPARC who have funded this work as part of the CREST Rover program.

## References

NASA (2007) Mars Exploration Rover Mission. http://marsrovers.jpl.nasa.gov/overview/

Huntsberger, T., Cheng, Y. , Stroupe, A., Aghazarian H. (2005) Closed Loop Control for Autonomous Approach and Placement of Science Instruments by Planetary Rovers. *IEEE/RSJ International Conference on Intelligent Robots & Systems*, Alberta, Canada,

Pullan, D. (2006). Scientific Autonomy for Planetary Rovers. Internal technical report, Space Research Centre, Department of Physics and Astronomy, University of Leicester.

Felzenszwalb, P. F. and Huttenlocher, D. P. (2004) Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision, Volume 59, Number 2, September 2004.*

Woods, M., Baldwin, L., Wilson G. and Hall S. (2006) MMOPS: Assessing the Impact of On- Board Autonomy for Deep Space Robotic Missions. *International Conference SpaceOps*, Rome, Italy AIAA-2006-5977

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

Castano, R., Estlin, T., Gaines, D., Castano, A., Bornstein, B., Anderson, R.C., Judd, M., Stough, T., and Wagstaff, K., (2005) Science alert demonstration with a rover traverse science data analysis system, *LPSC XXXVI.*

Thompson, D.R., Smith, T., and Wettergreen, D., (2006), Autonomous detection of novel biologic and geologic features in Atacama Desert rover imagery, *LPSC XXXVII.*