Autonomous Science Target Identification and Acquisition (ASTIA) for Planetary Exploration

Dave Barnes, Stephen Pugh and Laurence Tyler

Abstract— We introduce an autonomous planetary exploration software architecture being developed for the purpose of autonomous science target identification and surface sample acquisition. Our motivation is to maximise planetary science data return whilst minimising the need for ground-based human intervention during long duration planetary robotic exploration missions. Our Autonomous Science Target Identification and Acquisition (ASTIA) architecture incorporates a number of key software components which support 2D and 3D image processing; autonomous science target identification based upon science instrument captured data; a robot manipulator control software agent, and an architecture software executive. ASTIA is being developed and tested within our Trans-National Planetary Analogue Terrain Laboratory (PATLab). This provides an analogue Martian terrain, and a rover chassis with onboard manipulator, cameras and computing hardware. Experimentation results with ASTIA and our PATLab rover are presented.

I. INTRODUCTION

A major mission driver for space exploration is to maximise science data return whilst minimising ground-based human intervention and hence associated operations costs. Future robotic exploration such as the ESA ExoMars mission [1] (launch 2018), and the eventual Mars Sample Return (MSR) [2] mission will require rovers to travel further and faster than has been achieved to date. The current NASA Mars Exploration Rover mission has shown the need to reduce the number of full sol (Martian day) command cycles required to accurately place an instrument upon a terrain object once ground based scientists have identified this science target. Greater rover autonomy is an essential requirement if full sol command cycles are to be reduced. Going a stage further, we envisage the deployment of scout rovers capable of both autonomous science target identification *and* science sample acquisition. Such autonomous rovers could be utilised to identify and cache science samples as a precursor to a subsequent MSR mission.

The research presented here builds upon previous work [3] funded by the UK STFC. This work demonstrated autonomous science target identification and rover arm placement within our Trans-National Planetary Analogue Terrain Laboratory (PATLab). Since this work we have developed a new knowledge-based approach to science target identification, together with improvements to the calibration and control of our rover robotic arm and pan and tilt (P&T) hardware. What has emerged is a new software architecture design called ASTIA - Autonomous Science Target Identification and Acquisition, and this paper presents our ASTIA implementation progress to date. ASTIA comprises a number of (agent-like) software modules which are described here. We present more detail on a new module called KSTIS (Knowledge-based Science Target Identification System), together with an overview of the additional modules that have allowed us to perform end-to-end science target identification and rover arm placement trials. The results from this work are presented.

II. ASTIA BACKGROUND

Current research into autonomous systems for planetary exploration includes studies into rock detection and target prioritisation [4], feature detection [5], and novelty detection [6]. Typically, isolated rocks serve as potential science targets with the aim of assigning scientific parameters such as albedo, texture and colour, together with parameters such as a rock's shape and size. In some cases spectral signature and fluorescence indicators constitute analytical inputs. Results of this nature can be used to identify targets of interest and unexpected objects, and to characterise an exploration site.

A notable body of work is OASIS, the On-board Autonomous Rover Science Investigation System [4]. This has been designed to enable a rover to identify and react to serendipitous science opportunities such as dust devils, clouds and novel rocks that the rover has not seen before. OASIS analyses data that the rover captures, and then prioritises this data based upon established target attributes. It may also schedule new observations of interesting targets. The criteria for prioritisation are set to be appropriate to the current environment and science goals. OASIS currently uses greyscale (single filter) images for its rock identification and analysis, concentrating mainly on rock shape, size and albedo.

It is interesting to compare the target evaluation processes undertaken by previous autonomous science research with those processes undertaken by a human field geologist. Given that we are attempting to emulate the expertise possessed by a human planetary scientist, we realised that a human field geologist [7] typically assesses a potential science target in terms of its *structure* (e.g. geometric shape, scale, orientation and form), *texture* (e.g. luster, relief, grain size, shape, and sorting), and *composition* (e.g. colour, albedo, specularity

Dave Barnes and Stephen Pugh are with the Department of Computer Science, and Laurence Tyler is with the Institute of Mathematics and Physics, Aberystwyth University (AU), Penglais, Aberystwyth, SY23 3DB, UK. dpb@aber.ac.uk

and mineralogy). This approach has been used with good results in SARA (Science Assessment and Response Agent) [3]. SARA accumulates a numeric score value for potential science targets in an image, concentrating mainly on rock morphology at various scales. Whilst it is possible to represent some target attributes by a *singleton* value, many attributes are difficult to represent in such a manner. We realised that many science target attributes can be likened to a *fuzzy linguistic variable* [8] (e.g. the "roundness" of a rock, or its "distinct" cross-bedded structure). This led us to adopt a knowledge-based approach and we decided to represent the human geologist domain expertise as a fuzzy-rule set. This knowledge representation approach is fundamental to our KSTIS module.

Previous research has also addressed the problems associated with autonomous arm placement, i.e. given a target rock, the operation of moving an instrument and contacting the science target using autonomous arm control methods. A large body of work has been undertaken in this area [9] [10], and is generally described as SCIP (Single-Cycle Instrument Placement), or SCAIP (Single Command Approach and Instrument Placement). The main driver for this work has been the desire to maximise the science data return rate by limiting the number of required command cycles for each individual instrument placement operation. An added benefit is minimisation of the ground-based operator workload. The SCIP goal is to autonomously approach and place an instrument on multiple features of scientific interest in a single command sequence uplink. Vision-based target tracking techniques are key to the general SCIP approach, and 2D feature-based visual servoing has been used to keep a rover's navigation cameras foveated onto a science target while commanding the rover to move directly towards the given target. It should be noted that with the current SCIP work, it is a ground-based scientist(s) who identifies and selects the desired science target from a previously captured Panoramic Camera (PanCam) image.

Given the promising results emerging from research into autonomous science and autonomous arm placement, we believe that it is timely to combine both of these areas into a single advanced rover capability. Our resultant ASTIA architecture has been designed to facilitate this integration.

III. ASTIA SYSTEM OVERVIEW

The ASTIA system is directed towards planetary exploration missions such as ExoMars and MSR, and makes use of typical rover hardware. Key hardware components of the ASTIA architecture include: a pair of wide angle cameras (WACs) for stereo imaging; a high resolution camera (HRC) for detailed target analysis; a P&T mechanism, which together with the cameras form the PanCam unit; a robotic arm for deploying close-up or contact science instruments, and a rover locomotion chassis with associated on-board infrastructure. The key software components within the designed ASTIA architecture are shown in Fig. 1. Each rectangular box represents a software agent, and the lines indicate the flow of relevant information between agents.
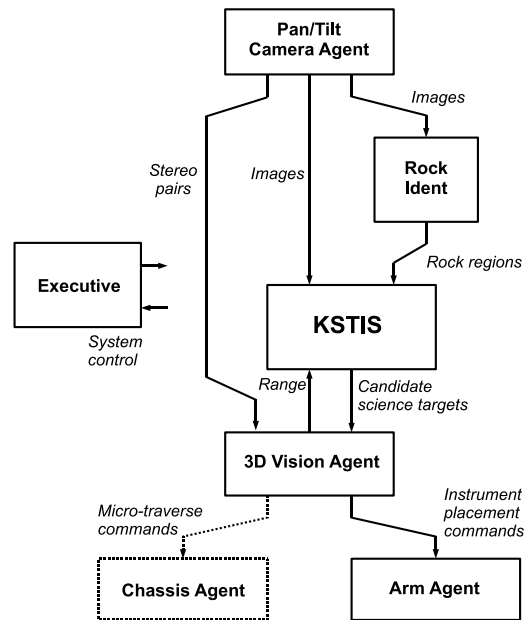

Fig. 1. Schematic diagram of the ASTIA architecture. Dotted lines denote work areas not addressed in this paper.

Note that the Chassis Agent (denoted by dotted lines in the diagram) is not addressed in this paper.

The Executive represents the operation sequencing and decision-making component of ASTIA. In a real mission scenario, the Executive would be closely linked with the onboard mission operations planning and resource management subsystems. Upon instruction from the Executive, The Pan/Tilt Camera Agent captures one or more stereo image pairs of a possible target site using the WACs, and passes them to other agents for analysis.

The Rock Identification Agent examines one image (typically the left-hand image) from each stereo pair. Rock regions within the image are identified and information about the size, location and centroid of each candidate rock is produced. KSTIS then applies domain expert knowledge to assess each candidate rock to identify the 'best' science target. The image pixel coordinates of this object (i.e. the rock's centroid) are then passed to the 3D Vision Agent. Additionally, KSTIS may request a higher magnification image of a candidate rock from the HRC to verify its choice of science target. Using the candidate science target image pair and the pixel coordinates from KSTIS, the 3D Vision Agent applies stereo triangulation to calculate the 3D position of the science target relative to the rover. If a zoom image is required by KSTIS, the Pan/Tilt Camera Agent can use this information to centre the science target in the HRC field of view and capture a suitable image. Knowledge of the 3D position also allows a science acquisition 'cost' to be calculated, based primarily upon the power and time that would be required for the rover to traverse to the science target location. This cost information can used by the Executive to assess the resource implications of a science activity, especially if the target is currently out of reach. If

the projected resource usage is acceptable, a rover traverse may be scheduled to place the science target within reach of the arm.

Once a traverse has completed, further stereo image pairs may be captured and passed to the Rock Identification Agent and KSTIS for a final science target assessment, and the 3D Vision Agent is notified of the image pixel coordinates of the chosen science target. The 3D Vision Agent and the ARM Agent then use stereo triangulation and the arm's kinematics model respectively to confirm target reachability. An appropriate arm configuration, instrument placement trajectory and contact region on the science target are determined. This process may involve 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 Executive, then the arm (with attached instrument) is moved, and science target contact is made.

IV. ROCK IDENTIFICATION AGENT (RIA)

A computationally minimal region growing algorithm has been developed to identify potential science targets. Each pixel in a WAC image is first polled to see if it currently has a region assignment. If not, a pixel object is created and is passed to a function which then examines all its neighbouring regions and finds the closest region for the pixel to join. If no region is found to be close enough, the pixel will stand on its own and create a new region. After this assignment stage all neighbouring regions are examined, and very close regions (currently regions with a difference of average pixel value less than 20) are merged. The next stage is an examination of all pixel regions that contain a small number of members: these are merged with their closest neighbour. The region data is then converted into an image for further processing. Here the regions are examined as objects with a uniform background and all adjoining objects are joined, numbered and labeled. This isolates the rock targets from the background image and reconstructs large rocks that were split into separate regions by the region growing algorithm. Note that the current algorithm has been tailored towards identifying strewn boulders as found in a boulder field. This constraint has both helped speed up development and reduce some of the target identification complexity. Additionally RIA can determine the centre-of-area for each identified rock region, and this can be used by the 3D Vision Agent to generate a 3D science target for the Chassis and Arm Agents.

V. KSTIS

A. Target Region Fuzzy Input Parameter Assignment

KSTIS is responsible for processing the images to determine the *Science Value (SV)* of any identified target. KSTIS is still in development and we foresee continued collaboration with our planetary geologist [7] (see section XI). For the current implementation, there is some need for human interaction during the image assessment stage, however there are four (automatically) calculated parameters: albedo, colour, whiteness and roundness. As the albedo is

approximated without any chemical knowledge or knowledge of the ambient lighting, the whiteness calculation is exactly the same, therefore both whiteness and albedo are represented on a sliding scale between 0 (black/very low albedo) and 255 (white/very high albedo). As it was desirable for colour to be represented by a single value, it was decided that the best way was through the use of a hue value taken from the hue, saturation and luminosity (HSL) colour space. The images taken by the Pan/Tilt Camera Agent are standard RGB images so it was necessary to convert the colour space representing the target area to HSL. The H value was chosen as on a scale of 0 to 1 it represented all available colours. The "roundness" indicator was determined by examining the eccentricity of the rock. This was represented by a number between 0 and 100, with 0 being "angular" and 100 being "very round". Other fuzzy-linguistic inputs processed by KSTIS include: "Surface", "Sphericalness", "Roundness", "Disk-likeness", "Rod-Type", "Scale", "Stratification", "Curviness" and "Lenticularity".

B. Fuzzy Rule-Base Implementation

The implementation takes the form of three fuzzy logic rule-bases: **Structure**, **Texture**, and **Composition**; one responsible for each of the three examined attributes. These rule-bases utilise Mamdani's fuzzy inference method [11], a number of membership functions, and a collection of rules. The combined output is then de-fuzzified using Centre of Gravity (COG) defuzzification. This returns a *crisp* number which represents a rock's *SV*. The implemented rules and membership functions have been developed through extensive collaboration with our *Domain Expert* [7]. This has led to a group of membership functions which model the way that the expert's interest in certain features developes. Trapezoidal functions were used where a range of inputs could be viewed as satisfying the membership criteria; for example, in the **Structure** rule base the *Scale* input utilises five trapezoidal functions to allow ranges of thickness to fully satisfy the membership (note that thin lamination can range from $2-3\ mm$), however the *Curviness* input also utilises two trapezoidal membership functions, but has in addition two Gaussian membership functions. The Gaussian functions model an input that has one 'fully' satisfying value and outside of that membership the degree of membership degrades slowly (see Fig. 2). Mixed membership functions (e.g. trapezoidal plus Gaussian) often proved to be a useful way forward when attempting to represent the diversity of domain expertise required here. Rules were developed in a similar way, i.e. during collaboration a quantification of the *SV* of certain geological features was produced (based upon the ESA ExoMars science goals). The rule base was developed through use of these data together with an appropriate mapping of crisp *SV* scores produced by the planetary geologist so as to generate the requied *Degree of Membership (DOM)* outputs.

Structure: Basic geometric forms are considered here. The most obvious form is layering or stratification, a term normally used in reference to sedimentary rocks but which
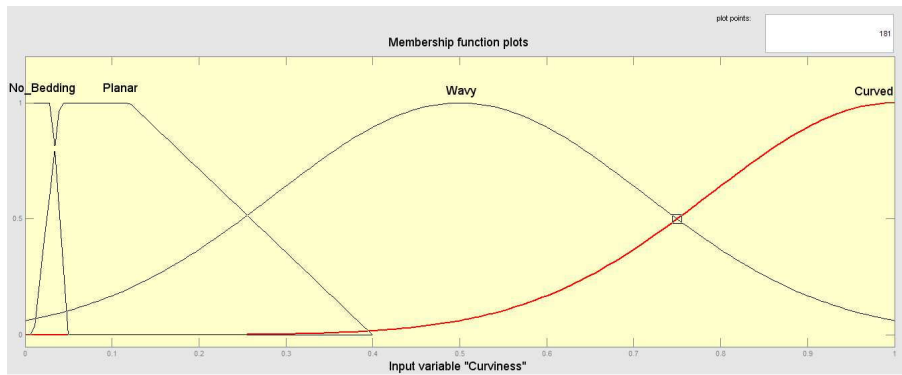
Fig. 2. Diagram displaying the membership functions associated with the *Curviness* input within the **Structure** Rule-Base.

can also be applied to volcanic and metamorphic deposits exhibiting layered structures. Where thickness is implied, units display either bedding ($> 1\ cm$) or lamination ($< 1\ cm$). This applies to all scales in the same way. Four membership functions were developed for this rule-base: *Scale*, *Stratification*, *Curviness*, and *Particularity*. The *Scale* value represents the perceived thickness of the layering or stratification, with inputs ranging from very thin lamination $< 2\ mm$ to very thick bedding $> 300\ mm$. *Stratification* is an indication of how continuous the structure is, whether it forms a continuous pattern or a broken or disjoint bedding structure. The *Curviness* of a feature is indicated by the inputs: No Bedding, Planar, Wavy or Curved. The final input is *Particularity*, in this case the input is whether the feature is lenticular or not. The Structure Rule-Base contains *23 rules*.

Texture: The textural properties of rocks are dependent on particle grain size and distribution, grain morphology and overall fabric (how grains are oriented and packed). Although these properties can only be determined at relatively close range, some generic aspects are applicable to remote observation of larger potential targets. The Texture Rule-Base contains five membership functions: *Surface*, *Sphericalness*, *Roundness*, *Disk-likeness* and *Rod Type*. The *Surface* input has a range of 0 to 100 which identifies the surface of the target as being Dull, Rough, Pitted, Polished or Bumpy. *Sphericalness*, is used to indicate the presence of a spherical or equant pattern on the surface of the target. *Roundness* refers to the shape of the overall target, as does *Disk Likeness* and *Rod Type*. The Texture Rule-Base contains *70 rules*.

Composition: This is the geochemical and mineralogical make up of rocks. It is perhaps the most demanding of attributes to define as weathering and alteration processes can subtly or radically change both the chemistry and/or mineralogy of rocks and soils. This means that there has to be much reliance on contextual data to assist in the interpretation of analytical measurements. Initial clues regarding composition however can be obtained from image data. The Composition Rule-Base contains three membership functions: *Hue*, *Albedo* and *Whiteness*. As previously stated, without analytical measurements of the target it is not possible to know its exact composition, but measurements like these implemented here
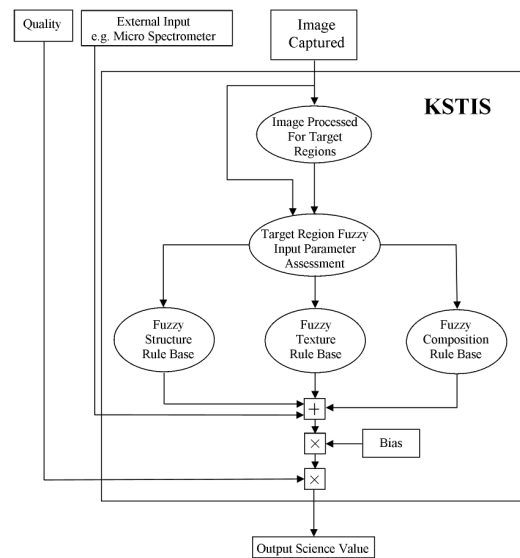

Fig. 3. Diagram of the KSTIS architecture (based upon [7]).

can give an indication as to the rock or soil make up. Colour is the first attribute measured and it is represented by *Hue*, as a single value, rather than the more common RGB triple. *Albedo* and *Whiteness* are inherently very similar as the higher the whiteness of the rock the more reflective it is. The Composition Rule-Base contains *15 rules*.

VI. 3D VISION AGENT

The current ASTIA stereo triangulation algorithm requires a simplified epipolar geometry to be observed, and hence any captured camera images have to be rectified. Rather than implement 'yet-another-disparity' algorithm, we wished to investigate the performance of a state-of-the-art approach that showed good performance when compared to other algorithms, and was able to deal with occlusion problems (a situation that is quite probable in a Martian 'rock garden'). We based our disparity map generation upon the cooperative algorithm for stereo matching and occlusion detection [12].

We found that this algorithm performed well provided that good (close to solution) minimum and maximum pixel disparity values were known a priori. The major problem with such an algorithm is the large computation time. Whilst this may not be an issue for terrestrial applications, when using disparity algorithms onboard an autonomous rover it must be noted that processing memory and power are very limited (of the order of 256 Mb memory, and 100 MHz clock rate!). Once a disparity map had been generated, stereo triangulation was performed using the obtained science target *left_image_x, y* and *right_image_x, y* pixel coordinates and the camera extrinsic parameters which were obtained during camera calibration. This resulted in the 3D position of the science target relative to the camera origin.

VII. ARM AGENT

A Helmert transformation[1] was used to transform the candidate science target locations from camera origin 3D space to arm base 3D space. The required parameters for the transformation were obtained from a calibration procedure which imaged a marker at the end of the robotic arm in various positions within the area of operation. Our Vicon MX tracking system (see section VIII) was used to obtain an accurate 3D position for the marker, corresponding to each image. A fitting algorithm was applied to this data to find the best parameters for the Helmert transformation.

The ARM Agent contains an inverse kinematics model of the arm, mapping 3D science target positions relative to the rover into arm joint angle values. Should a target rock DEM also be available from the 3D Vision Agent, then the Arm Agent can additionally process this data and determine the 'best' instrument placement sites on the target rock (e.g. rock planar regions where instrument-head/rock collisions can be avoided).

The arm used for the field trials was a 3 DoF demonstration device of limited accuracy constructed using radio control model servos, with no joint feedback information. In the absence of a full deflection model for this arm, an empirical calibration of the joint angle offsets was made over the work area of interest. This calibration technique was used previously with the Beagle 2 arm, and is described in [13]. The arm joint angle adjustment given by the calibration procedure was incorporated into the Arm Agent.

VIII. EXPERIMENTAL SETUP

A new Planetary Analogue Terrain Laboratory (PATLab) has been created at AU. The aim of the PATLab is to allow comprehensive mission operations emulation experiments to be performed. Such trials and experiments are essential when learning how to deploy and use a robot science instrument for a given mission, and hence maximise the return of high-quality data. The PATLab includes a $50\,m^2$ landscaped terrain region composed of Mars Soil Simulant-D (from DLR, Germany). The terrain includes an area for sub-surface
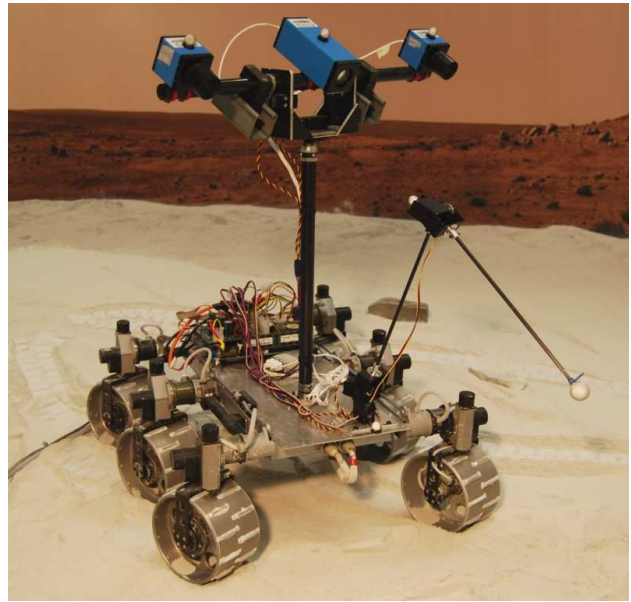

Fig. 4. The AU half-scale ExoMars-based rover chassis with 3 DoF arm and PanCam instrument. The grey spheres are passive markers used by our Vicon system to obtain arm placement metrics during ASTIA trials.

sampling and a collection of 'science target' rocks that have been fully characterised. The PATLab is heavily instrumented and its data and control facilities are available remotely via high-speed network links.

The PATLab supports a half-size rover chassis which is based upon the ESA ExoMars rover Concept-E mechanics [1], (Fig. 4). The rover has 6-wheel drive, 6-wheel steering, and a 6-wheel walking capability (thus 3 DoF per wheel). The rover supports a panoramic camera instrument and a 3 DoF robot arm, in addition to onboard computing and communication facilities.

Using COTS cameras we have built a panoramic camera instrument which emulates the proposed ExoMars PanCam [14]. Our PanCam supports two Wide Angle Cameras (WACs) with a baseline separation of $500\,mm$, and a High Resolution Camera (HRC) mounted centrally. Image capture and machine vision processing algorithms have been implemented and these can run using the rover on-board computer or remotely. A P&T mechanism attached to a mast structure on our rover chassis allows control over camera direction.

During PATLab experiments the position and orientation of the rover chassis, robot arm and PanCam P&T mechanism can be measured using our Vicon MX motion capture system. Using twelve specialised infra-red cameras, the Cartesian position of reflective markers placed anywhere within the PATLab terrain region can be tracked in real-time (typically $120\,Hz$) with *sub-millimetre* accuracy.

In the current design the ASTIA Executive co-ordinates the autonomous operation and decision-making of the ASTIA system in an integrated fashion. Since ASTIA is under development and currently spans several different computer systems, only some parts of the Executive are implemented at present. Others are simulated by a combination of scripts

---

[1]Helmert Transformation, named after Friedrich Robert Helmert, 1843-1917, is a method often used in geodesy to produce distortion free transformations in 3D space from one datum to another.

and some manual intervention. A basic resource calculation has been implemented. Based upon a priori information regarding motor speeds and power consumption for the rover chassis motors and the arm joint servo-mechanisms, science activity cost values can be calculated in terms of the time and power required to execute a traverse to a science target and/or an arm placement.

The rover PanCam, Arm and P&T unit were calibrated prior to conducting the ASTIA trials. The current P&T unit is constructed using radio control model servos and has a limited pointing accuracy of about $\pm0.41°$ in pan and $\pm0.84°$ in tilt ($\pm1sd$) due to both the intrinsic servo resolution and mechanical play in the joints. To partially compensate for these errors, the Vicon system was used to measure the pan and tilt values more accurately after moving the P&T unit each time, and the measured rather than commanded values were passed to the rest of the ASTIA system.

IX. RESULTS AND DISCUSSION

Experiments were performed to assess our current ASTIA implementation. These included Pan/Tilt Camera Agent, 3D Vision Agent and Arm Agent tests to measure the accuracy of arm placement from stereo imagery, together with RIA and KSTIS experiments. An 'end-to-end' integrated test was performed whereby a science target site was selected by KSTIS and the arm end-point moved to this target using the other ASTIA software agents.

Arm Agent: The accuracy of arm end-point placement from stereo imaging was tested by using Vicon markers as substitute science targets. Thirteen different marker positions within the arm's working envelope were imaged using both WACs and the resulting 3D position transformed from camera to arm coordinates. The Arm Agent was then used to produce joint angles which were applied to the arm. The resulting arm end-point positions were measured and compared to the target positions. Ten target positions were used to adjust the camera-to-arm Helmert transform parameters, resulting in a mean Euclidean position error of $12.695\,mm \pm 5.793\,(\pm1sd)$. The three remaining target positions had measured position errors of 13.814, 13.155 and $15.961\,mm$ respectively. Previous work has shown the calibrated placement accuracy of the demonstration arm to be $3.58\,mm \pm 1.79\,(\pm1sd)$. The residual error of approximately $9\,mm$ in mean position is largely attributable to mechanical play in the PanCam P&T unit and uncertainties in the 3D position reconstruction and coordinate transformations.

Rock Identification Agent: The RIA was used on a distant or 'standoff' image captured by the left WAC. The image output by the process can be seen in Fig. 5. Of the potential 16 rock targets 13 were successfully identified. Three rocks were missed; this was due to a combination of their small size in the image and a relatively low contrast between them and the soil background.

KSTIS: The output image from the rock identification process (Fig. 5) was analysed by KSTIS. The 13 identified rock targets were assessed and each target given a Science Value. The values assigned were as follows: $Rock\,1 = 56.75$;
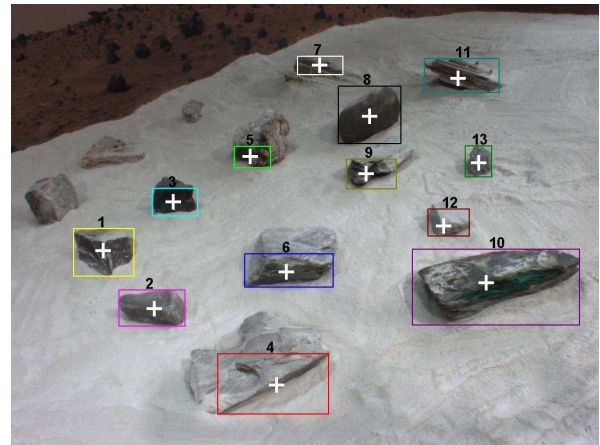

Fig. 5. Distant image, taken from a standoff distance of about $4\,m$, with 13 potential targets automatically identified and labeled. The centroid ('+') of each target has been identified for rover traverse purposes.

TABLE I

EXAMPLE OF THREE TARGET ASSESSMENTS FOR FIG. 5

Property Name	Rock 8	Rock 10	Rock 11
Surface	5	35	35
Sphericalness	100	100	10
Roundness	100	50	20
Disk Likeness	100	10	10
Stratification	5	10	10
Lenticularity	0	0	0
Rod Type	100	89	89
Whiteness	0.183	0.21	0.18
Scale	0	1.5	3
Curviness	0	0.4	0.1
Albedo	46.55	81.43	53.2
Hue	0.1	0.35	0.105
Science Value (SV)	45.77	129.6	109.6

$Rock\,2 = 51.5$; $Rock\,3 = 44.86$; $Rock\,4 = 64.36$; $Rock\,5 = 62.33$; $Rock\,6 = 66.85$; $Rock\,7 = 93.06$; $Rock\,8 = 45.77$; $Rock\,9 = 98.78$; $Rock\,10 = 129.6$; $Rock\,11 = 109.6$; $Rock\,12 = 51.27$, and $Rock\,13 = 51.04$. The values produced a rank order so the most valuable science target could be identified. Table I shows the detailed results of 3 example rock assessments. The most interesting rock was artificially enhanced by the addition of green ('chlorophyll') colouring. This helped to properly exercise the KSTIS rule-base and provide an expected high *SV* target. The results obtained were verified by our *Domain Expert*.

End-to-End Trial: The rover was moved to within arm working distance of the target rock identified by KSTIS (Rock 10, Fig. 5). A second, 'near' image of the target rock was captured and processed by the RIA to yield an updated centroid (Fig. 6. Note that Rock 10 in Fig. 5 is labelled here as Rock 7). The 2D centroid position was converted into a 3D target position by the 3D Vision Agent, and finally to a set of arm joint angles using the Arm Agent. Note that a $15\,mm$ stand-off from the target point was introduced to the commanded arm position to avoid an arm end-point/rock collision. Fig. 7 shows the achieved final arm end-point position. Although accurate distance measurements were not available for this trial (there was no Vicon marker on the
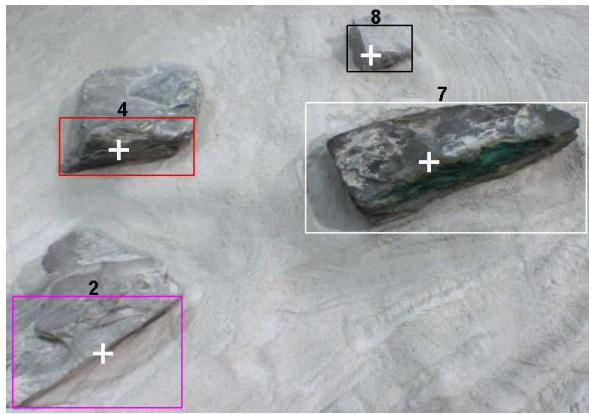
Fig. 6. Near image (cropped), taken from arm placement range ($< 1.5\ m$). The image was passed through the RIA to find the centroid of the target to be sampled. Rock 7 is the target rock here, which corresponds to Rock 10 of the original assessment. The centroid ('+') is used as the science target point for instrument placement.
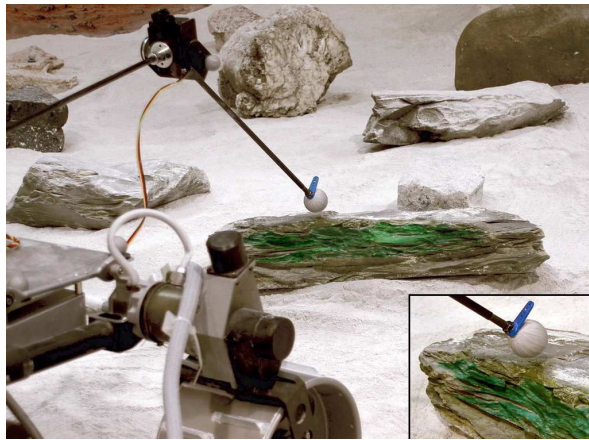

Fig. 7. Final arm end-point (science instrument) position during end-to-end trial of ASTIA. Inset is a close-up from a different angle.

calculated rock centroid), the arm end-point was estimated to be $< 2\ cm$ from the selected science target point, i.e. commensurate with the previous Arm Agent results and the introduced $15\ mm$ stand-off value.

X. CONCLUSIONS AND FUTURE WORKS

An autonomous planetary exploration software architecture has been designed for the purpose of autonomous science target identification and surface sample acquisition. Whilst the architecture implementation is work in progress, we have performed a number of trials with our current ASTIA software, and laboratory hardware. ASTIA combines both autonomous science target identification *and* autonomous arm placement, and this has been demonstrated. The results obtained have shown the integrity of our KSTIS knowledge-base, and support our decision to adopt a fuzzy-rule set approach to represent a human geologist's domain expertise. Future KSTIS work will focus upon completing the automatic extraction of the required fuzzy linguistic inputs from captured camera data.

Our arm placement accuracy results are commensurate with our previous arm trials, but there is room for improvement. Planned future work includes replacing our current P&T unit with a more accurate and precise COTS unit, with the aim of improving our overall ASTIA end-to-end accuracy. Similarly, we plan to eventually replace the 3 DoF demonstration arm with a 5 DoF precision model. Additionally we are working with the 5 DoF Beagle 2 development model (DM) arm [13] which we plan to integrate into our ASTIA setup.

XI. ACKNOWLEDGMENTS

REFERENCES

[1] Barnes, D. et al., The ExoMars rover and Pasteur payload Phase A study: an approach to experimental astrobiology, *Int. Journal of Astrobiology*, vol. 5, no. 03, 2006, pp 221-241.

[2] Schenker, P., S., Huntsberger, T. L., Pirjanian, P., Baumgartner, E., T., Tunstel, E., Planetary Rover Developments Supporting Mars Exploration, Sample Return and Future Human-Robotic Colonization, *Autonomous Robots*, vol. 14, nos. 2-3, 2003, pp 103-126.

[3] Woods, M., Shaw, A., Barnes, D., Price, D., Long, D., Pullan, D., Autonomous Science for an ExoMars Rover-Like Mission, *Journal of Field Robotics*, Special Issue on Space Robotics, Part II, vol. 26, issue 4, 2009, pp 358-390.

[4] Castaño, R., Estlin, T., Anderson, R.C., Gaines, D.M., Castaño, A., Bornstein, B., Chouinard, C., Judd, M., Oasis: Onboard autonomous science investigation system for opportunistic rover science, *Journal of Field Robotics*, vol. 24, no. 05, 2007, pp 379-397.

[5] Gulick, V., Morris, R.L., Ruzon, M.A., and Roush, T.L., Autonomous image analysis during the 1999 Marsokhod rover field test, *J. Geophys. Res.*, 106(E4), 2001, pp 7745-7763.

[6] Thompson, D.R., Smith, T., and Wettergreen, D., Autonomous detection of novel biologic and geologic features in Atacama Desert rover imagery, *In Proc. 37th Lunar and Planetary Science Conf., XXXVII*, League City, USA, 2006, CD-ROM Proc..

[7] Pullan, D., Analogue Studies for In Situ Surface Planetary Exploration, PhD thesis, Univ. of Leicester, UK, 2009.

[8] Zadeh, L., et al., *Fuzzy Sets, Fuzzy Logic, Fuzzy Systems*, World Scientific Press, 1996, ISBN 9810224214.

[9] Huntsberger, T., Cheng, Y., Stroupe, A., Aghazarian, H., Closed loop control for autonomous approach and placement of science instruments by planetary rovers, *In Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, IROS2005*, Edmonton, Canada, 2005, pp 3783-3790.

[10] Pedersen, L., Kunz, C., Sargent, R., Madison, R., Backes, P., Bajracharya, M., Clouse, D., Nesnas, I., Performance Evaluation of Handoff for Instrument Placement, *In Proc. AIAA Space 2006*, San Jose, USA, 2006, AIAA-2006-7385.

[11] Mamdani, E.H., Applications of fuzzy logic to approximate reasoning using linguistic synthesis, *IEEE Trans. on Computers*, vol. 26, no. 12, 1977, pp 1182-1191.

[12] Zitnick, C. L., Kanade, T., A Cooperative Algorithm for Stereo Matching and Occlusion Detection, *IEEE Trans. On Pattern Analysis and Machine Intelligence*, vol. 22, no. 7, 2000, pp 675-684.

[13] Barnes, D., Phillips, N., Paar, G., Beagle 2 simulation and calibration for ground segment operations, *In Proc. 7th Int. Symp. on Artificial Intelligence, Robotics and Automation in Space, iSAIRAS'03*, NARA, Japan, 2003, CD-ROM Proc..

[14] Griffiths, A. et al., Context for the ESA ExoMars rover: the Panoramic Camera (PanCam) instrument, *Int. Journal of Astrobiology*, vol. 5, no. 03, 2006, pp 269-275.