ON-BOARD PLAN MODIFICATION FOR OPPORTUNISTIC SCIENCE

Derek Long[1], Mark Wood[2], Andrew Shaw[2], Derek Pullen[3], David Barnes[4], and David Price[4]
[1] *University of Strathclyde, Glasgow, UK*
[2] *SciSys Ltd, Bristol, UK*
[3] *Leicester University, UK*
[4] *University of Aberystwyth, UK*

ABSTRACT

In this paper we report on a demonstration system using a rover chassis in a Mars-analogue environment, performing science target identification and evaluation in order to support opportunistic science experiments. To support this functionality requires the ability to reason about the impact of proposed operations on the structure and integrity of an existing plan, adapting or modifying the plan to allow for new operations. This function is performed by the plan management subsystem, TVCR. In this paper we describe the context in which TVCR performs this function and outline the techniques it uses to perform its role.[1]

1. INTRODUCTION

In this paper we report on an experiment that was recently completed as a demonstration for a collaborative research project, funded by the UK Science and Technoloy Facilities Council as part of their CREST programme. The intention behind the project was to put together several technologies to create a physical rover platform and use it, on a Mars-analogue surface, to show that it is possible to recognise geological science targets of interest and autonomously repond to them by deploying appropriate evaluation techniques and, possibly, more specific investigative tools. Others have also explored this same theme, particularly in the OASIS project [CEA$^+$07] which has pursued a very similar objective. The OASIS work has a similar architecture to the Autonomous Robotic Scientist we describe here, coupling image processing to identify potential targets with on-board planning capability to determine whether and how the actions that would be required to exploit possible targets can be inserted into the plan.

In this paper we describe the plan-modification and validation system built to support the Autonomous Robotic Scientist as well as, more briefly, the supporting subsystems that analyse and evaluation images, perform selection of targets for close-up examination and manage automatic arm deployment for close-up science data collection. To achieve these processes requires that the nominal mission plan, previously developed on the ground, be interrupted and modified. This process involves establishing the possibility of providing necessary resources, initially merely for evaluation of science targets and, subsequently, for the deployment of a camera for close-up investigation. The general structure of this process is one in which a nominal plan must be adjusted or modified to allow introduction of new activities into the mission timeline. To do this will require that existing activities be adjusted, possibly just by delaying them, but in more complex cases by abandoning lower priority activities in favour of high priority science target opportunities.

We present aspects of the planning process, including the model we used, the framework in which plan validation and modification takes place and the mechanisms required to support the process of plan modification.

2. SYSTEM ARCHITECTURE

Current procedures for planning operations of remote space missions represent a compromise between the conservative use of vulnerable systems and attempting to meet the ambitious demands of trying to acquire the most scientific data at least cost. This compromise is significantly impacted by communication constraints which mean that operations are typically planned on the ground for a window of hours or days ahead and then downlinked to the craft for subsequent unsupervised execution. There are several consequences of this process: firstly, with very limited capacity on-board to react to unexpected results of actions, the standard response to significantly off-nominal behaviour is for the craft to enter safe mode,

---

[1] This paper overlaps considerably with the paper of the same title presented at the International Workshop on Planning and Scheduling in Space (IWPSS'09) and also a journal paper covering the experiment in more detail [WSB$^+$09].

Proc. 'IJCAI–09 Workshop on Artificial Intelligence in Space', Pasadena, California, US
17–18 July 2009 (ESA SP-673, September 2009)

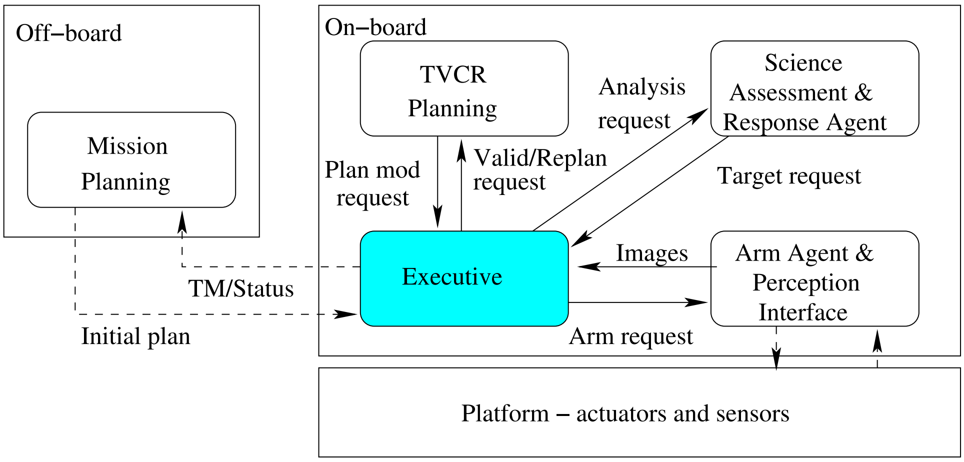

Figure 1. Architecture of the Autonomous Science Evaluation and Acquisition System

and secondly, the science gathering operations are limited to targets that are already known to the scientists at the point of planning. Both of these limitations have motivated work exploring the introduction of intelligent on-board systems, offering on-board diagnosis and replanning or plan repair and on-board science evaluation and response. The former is focussed on attempting to respond to problems during plan execution, recovering from operational failures by replanning in an attempt to avoid losing operations that might be unaffected by minor execution problems. The latter work, in contrast, is concerned with attempting to exploit opportunities to acquire scientific data that might otherwise be lost. In this paper we focus on the latter problem, although we do so by adapting a tool used in former work exploring plan repair [FLB+06].

The architecture of our demonstration system (figure 1) is based around three key components: a science assessment and response agent (SARA), an arm agent and perception interface (AAPI) and a timeline validation, control and replanning system (TVCR). The control system for the rover receives an initial plan from mission planners on Earth. During execution, TVCR is then used to monitor the execution process, responding to plan failure by proposing plan modifications to the executive. The executive can choose to accept the modifications or to behave more cautiously, perhaps entering safe mode. During normal operation, the executive will also request, at intervals, a science assessment to be performed on image data acquired from the navigation cameras. This analysis is performed by SARA. If SARA determines that there is a possible target, further evaluation will be requested and the executive can then request that the necessary resources for this evaluation can be scheduled by TVCR. Assuming further evaluation is performed and determines that a target of appropriate priority has been found, further replanning requests will be issued to TVCR to plan science operations. If a plan can be constructed to include the operations, then requests will be issued by the executive, according to the plan, to deploy instruments. These requests are handled by the AAPI, which solves the arm placement problem and deploys the appropriate instruments. The AAPI uses closed-loop control of the arm to resolve the approach and placement problem.

In this work we describe the details of the plan modification machinery, TVCR, but we first give an overview of the component subsystems that provide the context in which TVCR operates.

3. THE SCIENCE ASSESSMENT AND RESPONSE AGENT

The ability to autonomously identify and evaluate potential science targets is a key objective to improve the capability of future remote space systems. Indeed, even near space operations can be significantly enhanced if on-board science identification and evaluation can be achieved. EO-1 [CST+04] is an exciting example of what this technology can offer, in this case performing on-board evaluation of images to identify cloud cover in order to respond to poor data quality by automatically rescheduling new observations. Different science missions will present different potential science targets and each will require appropriate recognition systems. In this project we focussed on recognition of geological targets, working from a planetary rover base.

Science target selection is non-trivial, requiring a combination of experience and careful trading of choices between alternative targets based on anticipated science value and assessment of risk in trying to reach a target and deploy instruments on it. The problem is made harder for ground-based human scientists, since the choice has to be made with limited and very partial information: the images from which

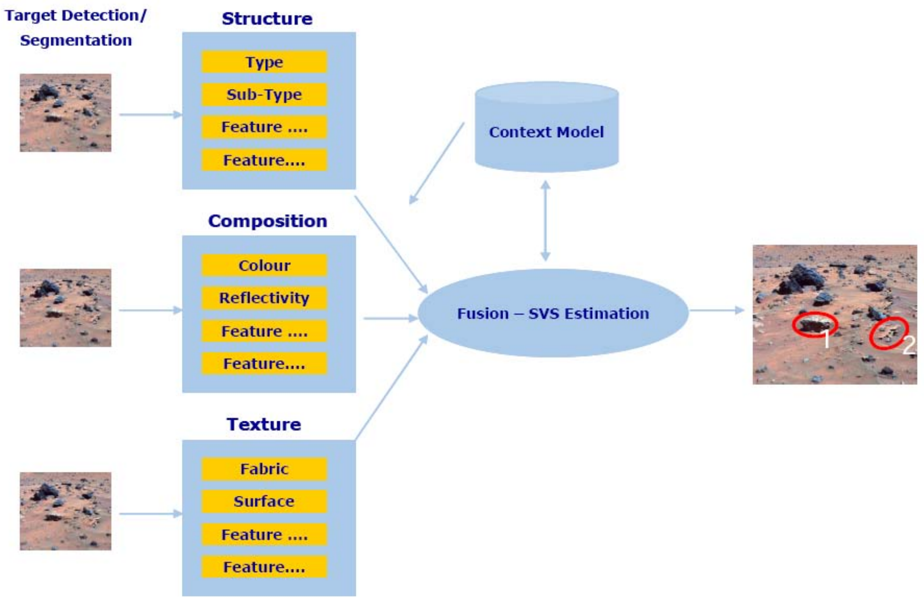

Figure 2. SARA: Target assessment is built up through the analysis of individual attributes. Context is included to qualify the final score.

selections must be made are typically the panoramic camera images of the terrain taken from the rover itself. Occlusion and image quality mean that many possible targets are simply not considered. Visual assessments of targes are best performed at a variety of scales from remote to close-up, since the key features that influence selection of targets, including texture, structure and composition, are hard to identify at long range. To achieve this without any on-board decision-making requires slow and painstaking work in planning short traverses for the rover, collection of further images and long delays while these are communicated, choices made and plans communicated back. Even then, targets can be missed because they were invisible from the vantage points at which the rover collected image data for the ground-based personnel. This is one of the reasons that opportunistic selection of science targets when the rover passes close by possible candidates could lead to discovery of important new scientific data.

The Science Assessment and Response Agent, SARA, uses image analysis to extract clues that indicate the possible value of candidates. Features characterising geological aspects of candidate targets were identified, representing structure (shape, scale and orientation), texture (luster, relief and grain) and composition (albedo, colour and mineralogy). Images captured by the on-board navigation cameras are analysed to identify possible targets. Candidate targets are then evaluated by summing weighted scores associated with these elements, taking their scores from the image analysis, to arrive at a science value score (SVS). Weights are dependent on the regional geology of the area being explored and the expectations and interests of the scientists: these allow scientists to set the overall objectives for opportunistic science gathering and focus on particular types of investigation. The weighted sum promotes feature-rich targets, since each new feature accumulates with others to increase the SVS for the target.

Image data can be acquired at various ranges from a possible target, and the range impacts on the way that the target is evaluated and the response that might be expected. Three ranges are considered: proximal (∼100cm), macroscopic (∼10cm) and microscopic (∼1cm). Initial target selection, performed by ground-based scientists, is performed at even longer ranges, but the function of SARA is evaluation of potential targets close to the rover during traverse operations. Long traverses are hazardous to perform without an initial survey and it is unlikely that such operations will be devolved to on-board control in the near term. Therefore, the range of our proximal observation area is deliberately chosen to be within the range of micro-traverse and arm placement.

An interesting challenge in the development of autonomous capabilities is to manage the process of gradual migration of responsibility from ground-staff to on-board systems. This is a process that depends on the incremental development of trust and this, in turn, shapes the extent to which on-board systems can direct operations. In our study, we only considered assessment of image data at proximal and macroscopic ranges. In part, this was a consequence of constraints on time and resource within the project, but it is also the case that microscopic imaging depends on the deployment of instruments, which represents a significant step in the release of control to on-board systems. At this stage, it seems unlikely that this degree of autonomy will be considered, particularly in early stages of near-term planned rover deployments.

Image interpretation includes an initial segmentation phase, to identify individual candidates within the image, followed by a feature extraction process [SB03]. In this work we focussed on successful extraction of features indicating *bedding*, which is potentially associated with sedimentary rock formation by deposition of layers and possible water erosion. However, the techniques are designed to be extensible to include other features.

4. TIMELINE VALIDATION, CONTROL AND REPLANNING

Our approach to the plan management problem is to exploit *plan fragments*, constructed by mission planners on the ground, parameterised by variables that are then instantiated with waypoints as science evaluation targets are considered and, as targets are identified, with those targets themselves. These instantiated plan fragments can be thought of as prepackaged plans, although they are typically not entirely self-contained, since they must be linked to the plan fragments that precede or succeed them in the structure of a modified plan. TVCR uses plan validation and systematic plan editing functions, alternately, to find a plan modification that achieves as many high priority activities as possible, while not exceeding resource bounds or breaking constraints between the plan elements.

Plan fragments are described as a collection of constrained actions, where the actions are themselves drawn from a domain model, capturing pre- and post-conditions for primitive actions available to the system. The base model is constructed in PDDL+ [FL06]. Constraints in a plan fragment determine sequencing, maximum and minimum separations, and possible constraints on the interactions between the plan fragment and other fragments (which might include dependency, mutual exclusion or ordering requirements). One reason for constructing plan fragments is that they allow operational constraints on the structure of plans to be captured in ways that are more convenient and more accurate than placing them into the domain model. For example, although it is logically possible to perform a rock-grinding operation before taking an image of the surface of a rock, it is methodologically unsound and, therefore, operationally constrained not to happen. However, it is possible to imagine situations in which this operational constraint might be overridden, relying only on the logical constraints in the domain (say, the grinder might be being tested after an attempted recovery). A second key benefit of plan fragments is that they encapsulate the results of what would otherwise be constructed by search.

Each plan fragment has an associated priority value, indicating the importance attached to the execution of the fragment by the missions planners. During plan repair, the priorities are used to determine which plan fragments should be dropped first in order to achieve an executable plan. We assume that there is a simple minimal plan that is guaranteed to be executable, which is the plan in which all fragments are removed and the single action instructing entry into safe mode is placed into the plan, together with the action committing to interaction in the next communication opportunity. Reduction is performed by iterative greedy elimination of plan fragments, using priority as the metric. Note that priorities simply provide an ordering on fragments — their absolute values are not interpreted as utilities, so there is assumed never to be a situation in which a combination of lower priority fragments might be preferred to a higher priory fragment. This greatly simplifies the problem of deciding which elements to maintain and which to drop, which is a combinatorial problem sharing structure with Knapsack and Orienteering problems.

Plan editing operations can remove fragments, move fragments (in order to exploit availability of temporally dependent resources) and reorder fragments. In general, the removal or reordering of fragments can leave breaks in the causal linkage of the activities that form a plan, where preconditions of a fragment are no longer satisfied in the state in which the system is left by the execution of the preceding fragments. In our application, there are key activities that ‘glue’ fragments together and which we do not see as part of the fragments themselves. These activities move instruments to specific locations or place instruments in particular modes. Selection of the correct sequence of activities to link together fragments is a planning problem, but it is highly specialised and constrained, so can be solved by a very simple means-ends analysis, apart from the planning of arm movements between specific joint-configuration waypoints. The latter is treated as a symbolic path-planning problem (the actual kinematics required to move the arm between waypoints are solved separately, in the AAPI) and solved by a simple shortest-path algorithm.

Once an executable plan is achieved, resources might have been released by the removal of activities from the original plan. Along with the fragments that form the original plan, missions planners can define additional plan fragments we call opportunities that may be inserted into the plan if the resources become available (subject to the constraints that might hold between these fragments and those in the plan). To simplify the problem of determining whether an insertion might be possible, each fragment has an associated estimated cost in terms of power, time and other resources. They also list the instruments they rely on. This information allows a rapid initial pass to eliminate candidate fragments for which there is insufficient resource budget or for which instrument requirements cannot be met. In our previous experiments, the numbers of opportunities have been small and it is a lightweight task to consider insertion using a greedy iterative incremental selection process.

4.1. Opportunistic Science

The need to address the existence of opportunistic science targets has limited impact on the TVCR behaviour described above: exploiting an opportunity amounts to inserting a new plan fragment into the current plan, if possible. This is a slight change to the interpretation of the way in which TVCR is expected to respond to events, since the opportunity is not a failure in an existing plan, but rather a possibility to enhance the current plan. It is straightforward to handle this by invoking the extension phase described above, restricting attention to the newly recognised opportunity. There are several reasons why this process is complicated by the discovery of these targets during execution of a plan.

1. The location of newly discovered potential targets will typically not be a waypoint, but will lie between waypoints on a traverse. This means that the location has no identifier and does not appear in the problem description for the current plan.
2. The target itself is not named in the problem description.
3. No parameters that describe the characteristics of the tasks involved in interacting with the new target are given in the problem description, so it is impossible to check whether actions that depend on availability of resources can be applied, since their preconditions are not instantiated.
4. Obviously no fragments referring to the target will appear in the opportunities library.

All of these points represent problems for symbolic reasoning about the planning problem, since the symbolic description does not contain the means to articulate, still less solve, the new problem. The resolution of these problems depends on a component of TVCR responsible for translating the current state, as it is determined from the history of the plan execution coupled with sensor data, into a symbolic form for use in the plan management. When a potential science target is spotted during a traverse, the initial state for plan management is constructed by creating a new location name representing the current location of the rover, splitting the traverse that was being carried out into two parts, joining the current location to the original start and end of the traverse. The costs associated with these traverses can be estimated based on the cost of the traverse that has already been completed and the anticipated cost of the original traverse. A location is created to represent the potential target, with characteristics based on the estimated distance to the target.

Plan fragments that refer to the new target can be instantiated from the library of existing fragments, by allowing fragments to contain parameters that are instantiated during plan management. This is an extension of the original capabilities of TVCR which was restricted to managing instantiated fragments, but it is a simple extension since the only new instantiations are those that are linked to the newly created object names (the new waypoint and new target). The most significant challenge in this is to associate appropriate priority with newly instantiated plan fragments. These are based on the SVS for the target science opportunities determined by SARA.

As analysis of a target progresses, new characteristics will be added to those already determined and science value might be adjusted. These changes must also be incorporated into the initial state for plan management, so the initial state is created incrementatlly reflecting the developing picture of the science target opportunity.

4.2. Operational Process

The sequence of events, as opportunities are considered, is as follows:

1. A candidate target is spotted in an image at proximal range.
2. TVCR is called to determine whether a further science evaluation activity is possible. This can be determined by requesting insertion of a new fragment corresponding to the instantiated relevant library fragment template. At this stage, the only information that is known about the target is its position relative to the rover and the only activity that is sought is capture of a zoom pancam image. The impact of this activity will be to cause the rover to halt while the pancam is targeted and the image captured. The parameters that govern the use of this fragment are the resource demands. These values are approximated, conservatively, by the Arm Agent, AAPI, and used to determine whether the fragment can be inserted without preventing completion of any higher priority activities. In contrast to the plan repair function, it is no longer reasonable to assume that the opportunities are lower priority than anything already in the plan. TVCR might be faced with the need to consider reducing the plan before extending

it, by removing fragments with lower priority than the opportunity. This allows for the possibility that ground-based planners might include low priority operations that will be completed if nothing better arises during the traverses, but will be dropped in favour of opportunistically identified targets.

TVCR responds with a 'go/no go' signal and modifies the plan accordingly.

3. Assuming that the response is 'go', the Arm Agent completes deployment of the pancam to obtain a zoom image. SARA evaluates the zoom image to decide whether macroscopic investigation should be considered. If the target generates a high science value score a new fragment is instantiated corresponding to approach, instrument deployment and data capture. The initial state is constructed to support the necessary resource estimates associated with actions that might be required for these fragments to be executed.

4. TVCR is again called to determine whether this fragment can be inserted. In order to do this, it requires an assessment of resource demands for each of the activities. An estimate for the arm deployment costs is provided by the Arm Agent using an approximation of the kinematics it uses to solve the actual deployment problem. It is important that this should be a lightweight estimate, in case the cost of assessment of the opportunity threatens the availability of resources to complete other activities in the plan.

   If TVCR determines that the new fragment can be added to the plan a 'go' signal is returned, together with the modified plan, and the executive instructs the AAPI to proceed to approach and perform instrument placement.

5. At completion of the experiment, the plan continues with a return to the original path and the continuation of the interrupted traverse. In principle, this fragment of the plan might have a low enough priority that, if resources are limited, it is displaced from the plan. This, of course, can only happen as a consequence of a deliberate selection of an appropriately low priority by the mission planners on the ground, but it offers the team a way to pause the rover at a site of particular scientific interest through the next communication window, so that a ground-based assessment can be completed and an appropriate follow through decided.

5. FIELD TRIAL

Figure 3 shows the set-up used in testing. The tests were carried out on a Mars analogue surface, using a half-sized scale model of the ExoMars chassis, with purpose-built camera mounts and a simple dummy arm. The area of the test site is clearly rather limited, which has the beneficial effect of making trials of reasonable duration, since traverse times are short, but it limits the size of the plans we could use with TVCR. In the tests on the physical device, TVCR is presented with a very simple problem involving one iteration of the complete cycle described in section 4.1. By altering the placement of a planned communications action, which has highest priority, we confirmed that TVCR could accept or reject attempts to perform opportunistic science operations, with both 'go' and 'no go' cases being possible at each decision point.

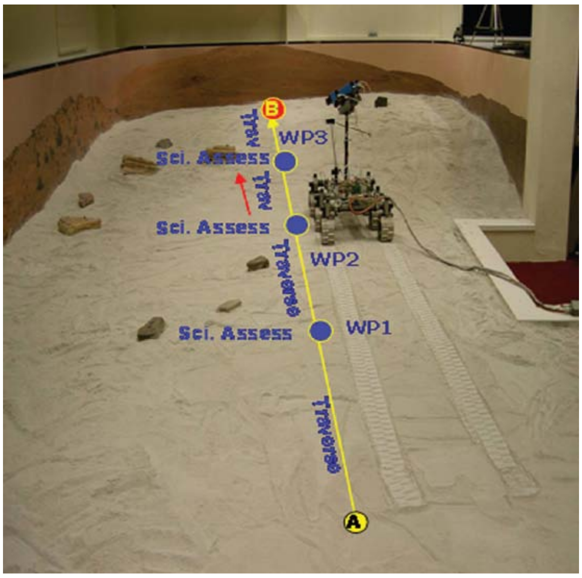

Figure 3. Field trial experimental set up.

Although the scenario we envisage includes possible micro-traverse operations to approach the target, we did not have an autonomous navigation component implemented in this field test. Therefore, the micro-traverse operation was performed manually, with appropriate time allowance in the plan to allow for manual operation.

In stand-alone tests, TVCR is able to handle more complex scenarios involving combinations of plan repair and opportunistic science insertion, including cases where opportunities are identified during a traverse rather than at a waypoint. By simulating various operating conditions, it is also possible to create scenarios in which early parts of a plan execute more efficiently than predicted, releasing resources that can then be used to exploit opportunities that would be excluded by constraints on future operations were these unexpected efficiencies not to have materialised. An interesting problem that arises generally in opportunistic science is the question of whether a greedy approach to science acquisition is the most effective strategy: in particular, should accumulated resources be spent on an opportunity now or saved for the possibility of a better opportunity later? Answering this question properly depends on having some model of expected rewards, based on probabilistic assessments of the chances of finding better opportunities later. This could be based on a probability distribution modelling the potential values of possible undiscovered targets along the unexplored parts of planned traverses. In reality, such models are unlikely to be available, so it seems that a more pragmatic solution is necessary: if a good target has been identified and there is resource available to exploit it, then it seems appropriate to take the opportunity.

TVCR is built on the VAL validator software [HLF04], which is not written for deployment on-board, but has been developed incrementally as an academic research software tool. Coupling it with the systems in the rest of the trial was achieved using CORBA, making communication relatively expensive. Fortunately, the trials were not hampered by the limited computational resources likely to be available on real mission hardware.

6. CONCLUSION

Our experiments have generated valuable data for the design of ExoMars, in all aspects. Development of the TVCR subsystem to support on-board plan repair and modification continues to be a strategic priority and reimplementation of the system as a verified and deployable software component is planned for the next 12 months.

The translation of progress in automated planning into the control of autonomous space systems is challenging. The challenges are not only technical, but also operational: it is critical that the process achieves growing levels of trust and experience has to be accumulated by both researchers in planning and operations staff who are faced with the prospect of interacting with systems with increasing levels of on-board autonomy. It is clear that the pressure for improving efficiency will lead to a steady adoption of technologies such as these, but it is equally clear that a leap from manual planning to fully goal-based on-board autonomous planning is impossible. TVCR represents an acceptable compromise, offering the means to perform constrained replanning and plan repair on-board, but within parameters that are tightly defined by missions planners on the ground.

REFERENCES

[CEA+07] Rebecca Castaño, Tara A. Estlin, Robert C. Anderson, Daniel M. Gaines, Andres Castano, Benjamin Bornstein, Caroline Chouinard, and Michele Judd. Oasis: Onboard autonomous science investigation system for opportunistic rover science. *J. Field Robotics*, 24(5):379–397, 2007.

[CST+04] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, R. Lee, D. Mandl, S. Frye, B. Trout, H. Hengemihle, J. D'Agostino, S. Shulman, S. Ungar, T. Brakke, D. Boyer, J. Van Gaasbeck, R. Greeley, T. Doggett, V. Baker, J. Dohm, and F. Ip. The EO-1 autonomous science agent. In *Proceedings of AAMAS-04*, 2004.

[FL06] M. Fox and D. Long. Modelling mixed discrete-continuous domains for planning. *Journal of AI Research*, forthcoming, 2006.

[FLB+06] M. Fox, D. Long, L. Baldwin, G. Wilson, M. Woods, D. Jameux, and R. Aylett. On-board timeline validation and repair: A feasibility study. In M. Giuliano, editor, *Proceedings of 5th International Workshop on Planning and Scheduling in Space*, 2006.

[HLF04] R. Howey, D. Long, and M. Fox. VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *Proceedings of 16th IEEE International Conference on Tools with Artificial Intelligence*, 2004.

[SB03] A. Shaw and D. Barnes. Landmark recognition for localisation and navigation of aerial vehicles. In *Proceedings of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2003.

[WSB+09] M. Woods, A. Shaw, D. Barnes, D. Price, D. Long, and D. Pullen. Autonomous science for an exomars rover-like mission. *Journal of Field Robotics*, 26(4):358–390, 2009.