Expressing Campaign Intent to Increase Productivity of Planetary Exploration Rovers

Daniel Gaines, Gregg Rabideau, Gary Doran, Steve Schaffer, Vincent Wong, Ashwin Vasavada, Robert Anderson

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109
{*firstname.lastname*}@jpl.nasa.gov

Abstract

Achieving consistently high levels of productivity has been a challenge for Mars surface missions. While the rovers have made major discoveries and dramatically increased our understanding of Mars, they often require a great deal of effort from the operations teams and achieving mission objectives can take longer than anticipated. Missions have begun investigating ways to enhance productivity by increasing the amount of decision making performed onboard the rovers. Our work focuses on the use of goal-based commanding as a means of more productively operating rovers. In particular, we are working on ways to convey the intent that operations team use to conduct science campaigns to the rover so that it can guide the rover in creating high quality plans and in identifying its own goals based on operator guidance. In addition to informing future surface exploration missions, this work is relevant for a wide range of applications in which operators must interact with a robotic system with limited communication opportunities.

Introduction

Maintaining high levels of productivity for the Mars exploration rover missions is highly challenging. While the Curiosity operations team has made significant accomplishments with the rover, doing so often requires a large amount of human effort in planning, coordinating, sequencing and validating the development of command products for the rover. Further, limitations in communication opportunities and anomalies on the vehicle can cause delays in accomplishing the team's objectives. These productivity challenges can result in the under-utilization of the vehicle's resources. These productivity challenges are anticipated to increase as our aging fleet of sun-synchronous orbiters are replaced by non-sun-synchronous orbiters, which do not provide a consistent pattern of "end-of-day" downlink relays.

The Jet Propulsion Laboratory has been exploring options for addressing these productivity challenges including conducting an extensive study of productivity factors in the Mars Science Laboratory mission (Gaines et al. 2016) and investigation into incorporating onboard activity scheduling for the Mars 2020 mission (Benowitz 2016).

Copyright © 2017, California Institute of Technology. U.S. Government sponsorship acknowledged.

In this paper we discuss work that is continuing the investigation into increasing the amount of decision making performed onboard rovers for future planetary exploration missions. In particular, we are leveraging existing technology and developing new technology to enable rovers to be more goal-directed, including following goals provided by ground operations as well as identify their own goals under the guidance of operators. We believe that this approach to goal-directed behavior will enhance surface mission productivity in the following ways:

Reducing operator effort: A goal-based interface presents a higher-level, more intuitive interface to rovers compared to the current highly-detailed command sequence interface. We believe that developing and validating command products with a goal-based interface will require less time and effort for the operations team.

Increasing rover resource utilization: The limited decision making of current rovers results in the operations team making highly conservative predictions of the amount of resources, e.g. time and energy, required to perform activity and results in significant amount of unused vehicle resources. By increasing onboard decision making, the rover can use knowledge of current vehicle resources to make more informed decisions about the goals that can be accomplished.

Reducing reliance on ground-in-the-loop cycles: Current operations relies on frequent interactions with the rover to maintain high levels of productivity in which the team assesses the rover's latest state and provides the detailed command products that direct the rover in accomplishing mission objectives. In contrast, a goal-based interface allows the team to provide objectives to the rover with reduced knowledge of the rover's state. In addition, by providing appropriate guidance, the rover is able to identify its own goals to accomplish mission objectives, further reducing the reliance on ground-in-the-loop contacts.

Although goal-based commanding has not been used on planetary rovers, it is of coarse a well-established form of interaction with robots and has also been used in other forms of space missions (Muscettola et al. 1998; Chien et al. 2005). In addition, the Opportunity and Curiosity rovers have a restricted form of goal selection in which they are able to

select targets for follow-up observations based on scientist guidance (Francis1 et al. 2016).

Our focus in this paper is providing rovers guidance on what goals to work on when:

1. the set of proposed goals over-subscribe vehicle resources and the planner must select a subset of goals to accomplish, and
2. the rover has a surplus of resource or has entered a new area that ground operators have not yet seen and should identify its own set of goals to pursue.

Our approach is to derive this guidance from the intent the science team uses when developing science campaigns. We go into more detail on campaign intent in later sections, but in general, we view intent as specifying relationships among goals. These relationships are used to determine the value of including a set of goals in a plan and, in some cases, how, or more specifically, when goals are accomplished. For example, the science team may be interested in collecting samples of a certain type of rock formation or performing a certain type of observation every X meters the rover drives. In the former case, the intent would specify the type of goal of interest (sampling a formation) and the value of accomplishing goals of this type (e.g. 2 or 3 samples is very important, additional samples are nice but less important). The latter case indicates a preference to periodically collect an observation and would indicate how important it is to accomplish the goals within a given tolerance of the indicated periodicity.

While our initial motivation for expressing campaign intent was for science objectives, this type of guidance is also relevant for many types of engineering maintenance activities performed by the rover. For example the team performs periodic activities to monitor various rover subsystems and dump system information. There is a cadence that must be followed in collecting this information, but there is flexibility in the exact timing of the activity. We have found that guidance for these engineering activities is similar to guidance for science campaigns.

In the next section we provide background on rover operations to help establish context. We then describe examples of rover science campaigns and how we derived campaign intent from these examples. Next we describe the specific semantics we developed to represent campaign intent and discuss how we are using this intent as guidance to enable the rover to generate high quality plans and identify its own goals.

Background on MSL Mission Operations

We begin with a brief overview of some important facets of MSL operations to provide context for the case study. This is not a comprehensive description of MSL operations, rather a description of some important aspects to help frame the case study.

One of the challenges of surface missions is the degree to which operations are impacted by a priori unknown and changing environmental conditions. While orbital imagery provides valuable information to guide activity, it does not capture all the conditions that affect the rover. For example, while orbital imagery may indicate that exploring a particular region is promising to achieve a science objective, the specific science targets are not known until additional data is collected from the rover itself.

As such, surface operations must be reactive and respond to the results of activity carried out during the previous sol (Martian day). This daily planning activity is referred to as "tactical" operations and is patterned after the tactical operations developed for the Mars Exploration Rovers (Mishkin et al. 2006).

MSL operations augments this tactical process with "strategic" and "supratactical" phases (Chattopadhyay et al. 2014). Strategic planning focuses on developing long-term plans, typically spanning weeks or months, to achieve high-level objectives. Examples of strategic planning include the development of strategies for exploring a large geographical area or a high-level traverse route for reaching a distant objective. The supratactical stage provides a bridge between the long-term strategic plan and the day-to-day, highly reactive tactical process. The process is designed to coordinate the complex science instruments and manage the constraints and resources required to conduct campaigns.

An Example Sol in the Life of the Rover

To provide an idea of how the team operates the rover, Figure 1 illustrates an example sol of rover activity. This is an example of a typical drive sol, derived from an actual Sol 780 command products. Following are some key aspects of the sol.

The plan for each sol begins with an "Uplink" window in which new commands products may be sent to the vehicle from Earth. There are various downlink windows throughout the sol in which the rover uses orbiter relays to send collected data back to Earth. While there are multiple downlink windows, certain downlinks have increased importance based on the time that data in the relay will reach operators. If data from a relay will reach operators by the start of the next tactical planning shift, then they relay is termed "decisional" because data from the relay can be used to make decisions in for the rover's next plan. Which relays are considered decisional depends on the relative timing between Earth and Mars along with latencies in the orbiter relay process. In Figure 1 the starred "MRO Relay" represents the decisional relay for this sol. It is important to realize that for this plan, only the data collected prior to this pass could be used to inform the next plan. While the remaining data will eventually be sent to Earth and may be used to inform future plans, it will arrive too late to inform the next plan.

Another important aspect of Figure 1 is how the team structures the sol into "blocks" of activity. For example, the main portion of the rover's day consists of a Pre-Drive Science block, a Drive with Mid-Drive Imaging bock and a Post-Drive Imaging block. The block structure organizes activity into related groups and allows a "Master" sequence to enforce timing between these major types of activity. The latter has to do with uncertainties in predicting the duration of activity in the plan. Due to environmental conditions such as lighting, scene content and terrain, the time to perform imaging and drive activities varies. The team uses the

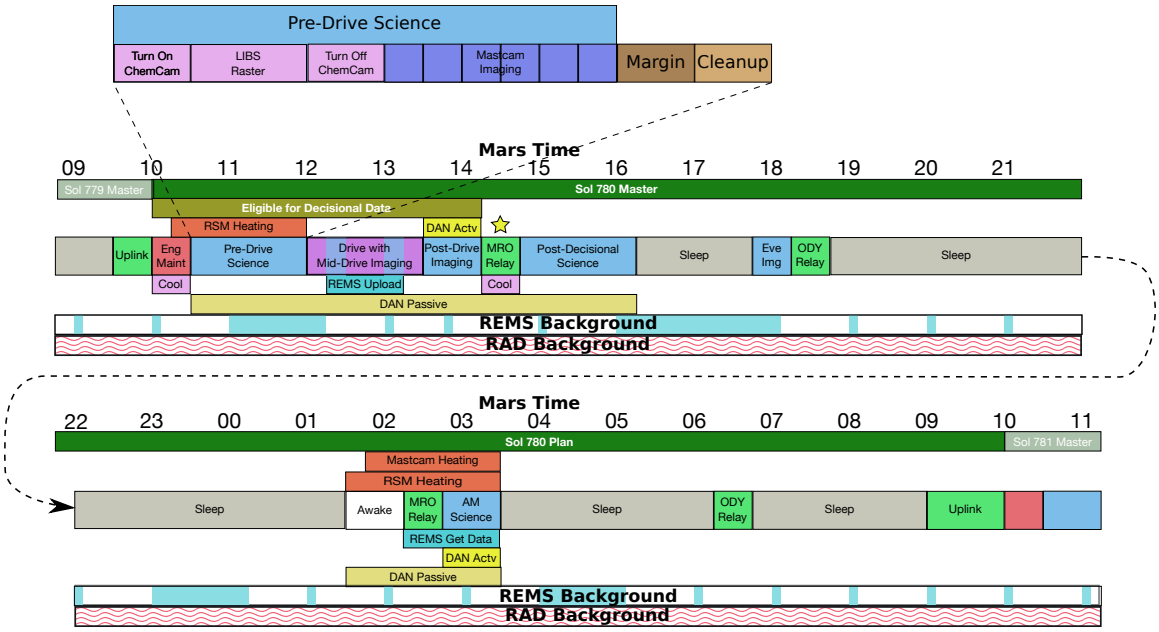

Figure 1: Example sol in the life of the rover.

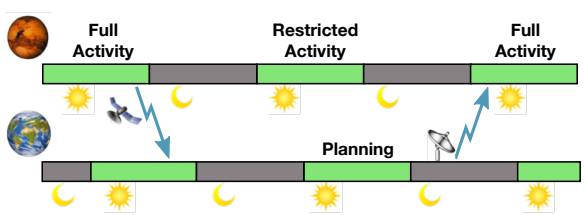

Figure 2: Mars activity vs. Earth planning.

block structure to ensure that if activity in one block runs longer than expected, it can be cut off to avoid interfering with subsequent activity. To protect against loss of data, the team builds "Margin" into each block, to allow activities to run longer than predicted. To deal with cases where durations exceed allocated margin, the team also sequences "Cleanups" after each block, to ensure that any activity is finished before the start of the next block.

Restricted Sols

The vast majority of the surface mission is conducted with the team restricting operations to daytime hours on Earth. The consequence is that the operations team is often out of sync with the activity of the rover on Mars. Figure 2 illustrates the impact this can have on the data available to the team during planning. In the diagram, the end-of-day relay from the rover arrives on the ground late in the Earth day. The team waits until the next Earth day to begin planning. Meanwhile the rover is waking up for its next Mars day without a new set of command products from Earth. By the time the team has completed the tactical process, they must wait for the subsequent Mars morning to uplink the products to the vehicle.

This often limits what the team can command the vehicle to do during the middle sol of Figure 2. If the vehicle were allowed to make significant changes to its state, in particular driving to a new location, this would significantly limit the types of activities the team could command on the subsequent sol. These limited activity sols are referred to as "restricted sols" because the latency of data often restricts the type of activity the team can perform.

A similar situation arises when the team takes days off for weekends and holidays. In these cases, the team will create plans that span multiple sols (aka multi-sol plans). Again, activities that result in significant changes to vehicle state are limited since they will impact the activity that can be done in later sols of the plan.

Given the current way in which we design and operate rovers, restricted sols are a major detractor from mission productivity. For example, with current surface operations, when the rover drives to a new location it must wait for imagery collected at this location to be sent to Earth and for the science and engineering teams to analyze the data and identify the specific set of activities to perform at the location to meet their current mission objectives. If the mission is in a restricted time period, this results in an entire sol in which the rover waits for these new activities.

Overall, 41% of sols on the MSL mission are restricted sols. This percentage is expected to be much higher if the mission were to rely on a highly eccentric relay orbiter such as the MAVEN orbiter.

Resource Prediction

An additional challenge to surface operations is that it is difficult for ground operators to predict the time and energy that will be required to perform these activities and the consequence of over-subscribing resources is severe (e.g. safing

the vehicle if energy is over-subscribed). As such, the team makes conservative estimates which almost always results in significantly under-utilizing available vehicle resources.

In the campaigns we analyzed in our MSL case study, we estimated that the rover could have conducted an additional 3 to 4 hours of activity each sol of the campaigns with the energy that went unused (Gaines et al. 2016). This would have resulted in a dramatic increase in productivity.

Identifying Campaign Intent

The previous section described the significant loss of productivity experienced by surface missions due to restricted sols and the challenge predicting resource usage. Our objective in this work is to increase the autonomy of rovers so that they can remain productive even in situations of reduced contact with human operators. Our approach is to convey mission intent to the rover so that mission operators can provide guidance to the rover even if they do not know the specific state the rover will be in when it receives the guidance. Further, operators will be able to provide a collection of goals that have the potential for over-subscribing available resources and let the rover select a high value subset based on actual available resources.

To enable more autonomous operation of a Mars rover, we aim to capture the intent behind activities planned for the rover during its service of a scientific campaign. Ultimately, the intent of all Mars rover activities is to advance our scientific understanding of the planet. However, by breaking up this larger intent into smaller, well-defined components, we can make some of this knowledge accessible to planning software for decision making. With the intent knowledge carried onboard, along with software that can use it, the rover can be more productive during times when ground interaction is limited.

The first step in this process is to understand the types of science and engineering intent that drives surface missions. We looked at several MSL campaigns in an effort to identify common relationships between planned activities and the objectives they are meant to achieve. Specifically, we investigated relationships that influence the inclusion or exclusion of an activity, as well as relationships that influence the timing of the included activities. From our initial investigation, we found three general types of relationships: samples of a class, temporally periodic observations, and sampling based on changes in rover or environmental state. In this section, we describe each of these types of relationships and the MSL examples that motivate them.

Sampling from a Class

The first relation we discuss is "samples of a class". In this relationship the operations team has some class in mind that they wish study and they want the rover to collect observations of examples of this class. This type of relation may be used to identify general areas to study or it may result is the selection of specific targets.

For example, during the Pahrump Hills Walkabout campaign the team performed a reconnaissance loop with the high level objective of studying a rock formation named

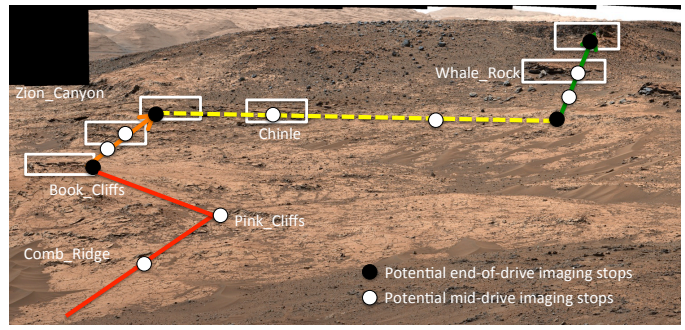

Figure 3: Curiosity's planned route for Pahrump Hills Walkabout Pass 1 at start of campaign.

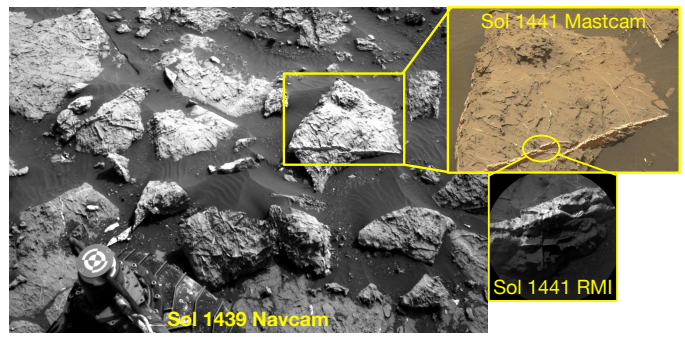

Figure 4: Catabola, an example vein target (NASA/JPL-Caltech/MSSS/LANL/CNES/IRAP/LPGNANTES/CNRS/IAS).

Murray formation, the strata recognized as lower Mount Sharp (Stack et al. 2015). During the strategic planning phase of this campaign, the team used imagery from the rover to identify resistant beds and other examples of unique rock textures that they wished to explore to develop a better understanding of the area. Figure 3 shows these identified areas, white boxes, as well as the initial route that was planned to visit them. These white boxes are an example of general areas that represent samples of class (e.g. resistant beds) the team is interested in studying.

This type of sampling from a class frequently occurs at a more local scale. For example, the team is often interested in studying veins that run through rocks (Nachon et al. 2015). Figure 4 shows the Catabola target, a vein identified in imagery acquired following a drive on Sol 1439. On Sol 1441, the operations commanded the rover to acquire the corresponding Mastcam and ChemCam data which resulted in the detection of high levels of boron (Jet Propulsion Laboratory Press Release 2016a). Notice that although the imagery containing Catabola was first collected on Sol 1439, it was not until Sol 1441 that the follow-up observations were collected. This is because the team was in restricted sols during this time and the rover spent Sol 1440 collecting untargeted observations without ground interaction.

Figure 5 shows another example of sampling from a class. In this case, the team identified two targets that were good examples of light-toned rocks, Elk and Lamoose. These

Figure 5: Navcam image taken after the Sol 991 drive, showing the Elk and Lamoose targets.

turned out to be particularly interesting targets as follow-up observations indicated they contained high levels of silica. The study of the high-silica targets lead to the conclusion that the introduction of silica represented one of the most recent water-rock interactions observed in Gale crater (Frydenvang et al. 2016). The data would also contribute to the study of the subsequently visited Bridger Basin area (Gasda et al. 2016).

In each of these examples, there are typically several instances that would serve as examples of a class. Given the limited time and resources available as well as other objectives the team has, the operations team must select a subset of candidate targets to perform follow-up observations. We would like to explicitly capture this notion of sampling from a class and allow operators to convey the value of collecting multiple samples from the class to enable the rover to make similar trade-offs. For example, it may be important to collect at least 3 samples of a class. Additional samples are nice to have, if additional resources are available, but with reduced additional value.

Temporally-Periodic Sampling

Other relationships between activities result from the desire to sample the Martian terrain and atmosphere at regular time intervals. These sampling intervals can determine how many observations are needed, as well as preferences on the temporal separations between observations. For example, throughout the mission, a variety of periodic observations are collected to monitor the Martian environmental conditions. Example monitoring includes periodically acquiring images of the sun at different times of day to measure atmospheric opacity (Mason et al. 2017). The SAM (Sample Analysis at Mars) is used to perform period sampling of the atmosphere with the intent to monitor the seasonal evolution of methane in the atmosphere (Webster et al. 2015).

While our primary focus on capture intent is for increasing science productivity, there are also many type of engineering activities that have similar intent relationships. For example, the team performs maintenance of different rover subsystems and collects detailed dumps of various vehicle telemetry on periodic cycles. The team also performs a variety of instrument calibration activities including periodic viewing of calibration targets for ChemCam (Chemical Camera), APXS (Alpha Particle X-Ray Spectrometer) and MAHLI (Mars Hand Lens Imager) calibration targets, to name just a few. While these maintenance activities may consume time and resources that would otherwise be used for immediate science, they must be considered to keep the rover safe and healthy for future science activities.

For both science and engineering periodic activities there is a preference on scheduling observations at a particular cadence, but there is flexibility in deviating from a precise interval to allow these periodic activities to be inter-mixed with other objectives vying for the same pool of rover resources. Thus, when capturing periodic intent, we want to express the preferred period as well as how value decreases as particular observations deviate from the preferred timing.

State-Based Sampling

The final type of intent relation we consider are requests that are based on changes in state that occur as the rover operates on the Martian surface. For example, as the rover traverses across the landscape there is interest in collecting certain types of observations including using navigation cameras to perform clast surveys (Yingst et al. 2013) and using the DAN (Dynamic Albedo of Neutrons) instrument to search for signs of subsurface water (Litvak et al. 2013). Due to the complexities of current operations practices the team is limited to acquiring these observations at the end of drives, which have high variance in their lengths, rather than at optimal distances to support systematic sampling. Part of our objective in increasing rover autonomy is to make it easier for the rover to perform these types of surveys closer to their ideal locations.

In recent operations, Curiosity has been climbing Mount Sharp to reach a layer of hematite, as seen in Figure 6. The team is interested in performing systematic surveys along the route in order to study variations that occur up the slope. In this case, the sampling strategy is based on change in elevation rather than strictly distance.

There are also engineering maintenance activities that correspond to changes in rover state. It is well known that the rover's wheels have suffered damage during its explorations. The engineering team conducts a periodic wheel wear monitoring activity based on distance traveled (Jet Propulsion Laboratory Press Release 2016b). The team also has the rover perform an attitude update activity in which the relative location of the sun is combined with accelerometer data to update the rover's attitude knowledge. This attitude

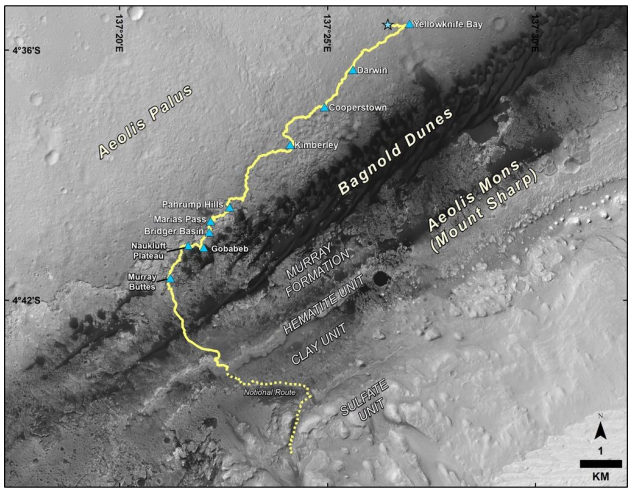

Figure 6: Curiosity's planned route to reach hematite layer (NASA/JPL-Caltech/Univ. of Arizona).

update activity is also scheduled as a function of distance traveled.

As with temporally periodic activities, this state-based periodic activities also come with a preferred cadence with flexibility about the actual timing.

In summary, we have reviewed several MSL science and engineering campaigns to identify patterns of relationships between activities and the objectives they serve. We identified three common types of relationships:

Sampling from a class: goals are selected based on how well they exemplify a class, there is typically an increase in value as more examples are collected, but with diminishing returns after a certain number of samples is reached.

Temporally-periodic sampling: goals are selected and scheduled based on a periodic temporal relationship, there is typically a preference on the cadence but with some amount of allowed flexibility in specific timing.

State-based sampling: goals are selected and scheduled based on changes in state of the rover and/or terrain, there is typically a preference on the cadence but with some amount of allowed flexibility in specific sampling.

In the next section, we discuss these types in more detail, including how they are implemented to capture intent and enable autonomous plan generation and repair.

Expressing Campaign Intent

We define a new type of planning construct, called a Plan Campaign, in order to guide automated planning algorithms towards solutions that better satisfy high-level science campaigns. These Plan Campaigns impose relationships between activities in the plan as a way of capturing the intent of performing those activities. In addition, a Plan Campaign provides an assessment of the current plan, indicating how well it is satisfying the constraints and relationships of the

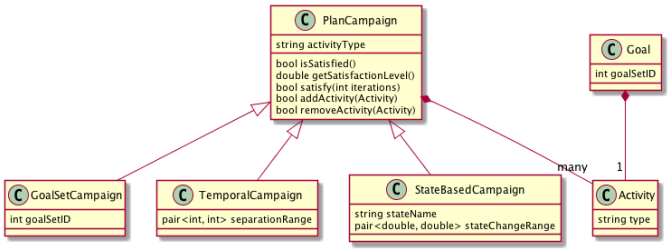

Figure 7: ASPEN classes for the three campaign types.

campaign. This satisfaction level, for example, may be defined in terms of the number of activities in the plan compared to the maximum requested by the Plan Campaign. A Plan Campaign also provides a set of satisfaction methods, which define ways that the plan might be changed to better satisfy the campaign.

We implement three types of Plan Campaigns in the ASPEN planning system (Fukunaga et al. 1997), corresponding to the most common relationships identified in the previous section. They are:

Goal Set Campaign: activities are scheduled based on a specific group of goals that request them. This type can be used to implement the "Sampling from a class" relationships.

Temporal Campaign: activities are scheduled based on temporal separation constraints. This type can be used to implement the "Temporally periodic sampling" relationships.

State-based Campaign: activities are scheduled based on state-change separation constraints. This type can be used to implement the "State-based sampling" relationships.

Each plan campaign type has its own definition for satisfaction level, and its own satisfaction method. A partial class diagram can be seen in Figure 7.

A Goal Set Campaign is a request to include in the plan some of the goals from a predefined set, such as those to sample a rock classification. In ASPEN, a goal is a request for an activity with specific parameter settings. For example, one rover goal may be to observe a particular rock at a close distance using stereo cameras. A goal might also generate new goals, such as when the intent is for the rover to collect measurements of some rock outcrop, and the particular location of the measurement can be generated onboard after approaching the outcrop and acquiring more detailed images not yet available on the ground for targeting. Each goal can be assigned an ID to specify the goal set to which it belongs. The Goal Set Campaign then references the same ID to request the set of goals. It also has parameters to specify the minimum and maximum number of goals to select from the set, which become constraints in ASPEN. The level of satisfaction increases as more goals are added to the plan. We use this type of Plan Campaign, for example, to capture a set of targeted observations for the rover that may all investigate the same type of rock formation. Multiple Goal Set Campaigns could be used for different types of formations, with

each competing for time and resources in the plan. For example, driving from one formation to another will take time and energy, but may provide more value than additional observations at the current location.

A Temporal Campaign is a request to include activities at regular time intervals. Again, a minimum and maximum number of activities can be provided. This type of campaign is very similar to the "Repeat" campaign type used to schedule Rosetta science observations (Chien et al. 2015). In addition, a minimum and maximum temporal separation is specified and captured as an ASPEN constraint to ensure an even sampling. Environmental monitoring, such as measuring dust levels in the atmosphere, is often specified in this way. In this case, the satisfaction level will be a function of not only the number of activities, but how well they are spaced.

The last, called a State-based Campaign, is the type of campaign implemented to support sampling based on rover state. In planning systems, state predictions are often plotted on a timeline. This prediction can then be used to determine where to schedule activities that were requested by the State-based Campaign. For example, as we include drives in the plan, we can predict the total distance traveled along a timeline. If a request was made to sample every 100 meters, these sampling activities can be placed on the timeline where the drive distance changes by 100m. A range of acceptable spacings can be provided using minimum and maximum parameters in the State-based Campaign definition. This, for example, would allow a campaign to request samples between 90m and 110m apart. Note that the drive, which may have been added to support a different campaign, may need to be interrupted to perform the sampling. Again, campaigns will compete for time and resources, and the satisfaction level of each campaign can be used to make planning decisions for requested activities. As with the Temporal Campaign, the State-based Campaign satisfaction level will partly depend on activity spacing. In this case, however, the spacing is driven by the change in a particular attribute of the rover state. Requested activities can be moved along planned state changes, or other activities can be used to explicitly change the planned state. For example, a drive could be added solely to increase spatial sampling, if no specific target location is provided.

All three types of Plan Campaigns have been implemented as extended features in the ASPEN planning and scheduling system. Problem-specific campaigns definitions are specified along with activity definitions in the ASPEN Modeling Language (AML). Once a set of campaigns have been provided, ASPEN scheduling functions can be used to generate campaign activities, repair campaign constraint violations, or optimize campaign satisfaction.

Using Campaign Intent

In this section, we discuss how captured campaign intent can be used for activity planning, as well as autonomous goal generation.

Activity Planning

One of the primary reasons for capturing and expressing campaign intent is to provide guidance for on-board plan generation and repair. During plan generation, the objective is to create and schedule activities that best satisfy all campaigns according to their preferences and relative priorities. During plan repair, the objective is to change the plan to better serve the campaigns in light of new state information. While plan generation may be done at regular times of relatively low activity (e.g. during the night), plan repair will most likely be triggered during execution when a new plan is needed quickly to keep the rover busy. We describe the algorithms implemented for generating and repairing plans based on a set of input campaigns that have been expressed in the manner discussed in the previous section.

Our approach to plan generation is based on branch-and-bound search. Here, various options (i.e. branches) are created from a partially generated plan, starting with the empty plan. As they are generated, the various options are evaluated and pruned based on a threshold (i.e. bound) on plan quality. Specifically, we compute the best possible quality of any plan that could be built from a new partial plan, and compare this to the worst possible quality of partial plans that have already been considered. If the new option can do no better, then that option, and all possibilities that extend from it, are pruned. The quality metric used for pruning is also used to periodically sort the options under consideration. After sorting, the partial plans that have the highest potential quality are expanded first. This is often referred to as "best-first" search. Quality of a partial plan is based on the level of satisfaction for the prioritized set of campaigns. Finally, after being selected, a partial plan is expanded forward in time. For periodic campaigns (Temporal and State-based), this means that the earliest activity is added first, while the next expansion will schedule the next activity based on the campaign separation constraint. For non-periodic campaigns (Goal Set), the partial plan is expanded multiple times, one for each goal scheduled to occur after the periodic activities scheduled so far. In the end, the search will evaluate all possible combinations of including or excluding campaign activities, as well as all possible orderings of those included.

While this can be time-intensive, it is guaranteed to find the optimal plan as defined by the campaign preferences and priorities. Further, a time limit can be placed on the search, to ensure that a plan is returned in a timely manner. Although this may not be a globally optimal plan, it will enable the rover to continue to be productive, and can be adjusted by more time-efficient repair strategies.

For plan repair, we use a greedy, local search algorithm to make fast improvements to the plan (Rabideau, Engelhardt, and Chien 2000). This algorithm can be run iteratively as the plan executes and new state information is received. On each iteration, the campaigns and their preferences are evaluated. From this, one preference is greedily selected based on its prioritized contribution to plan quality. An example might be a pair of activities occurring hours apart that are contributing to a high-priority campaign that is requesting periodic activities with a one hour separation.

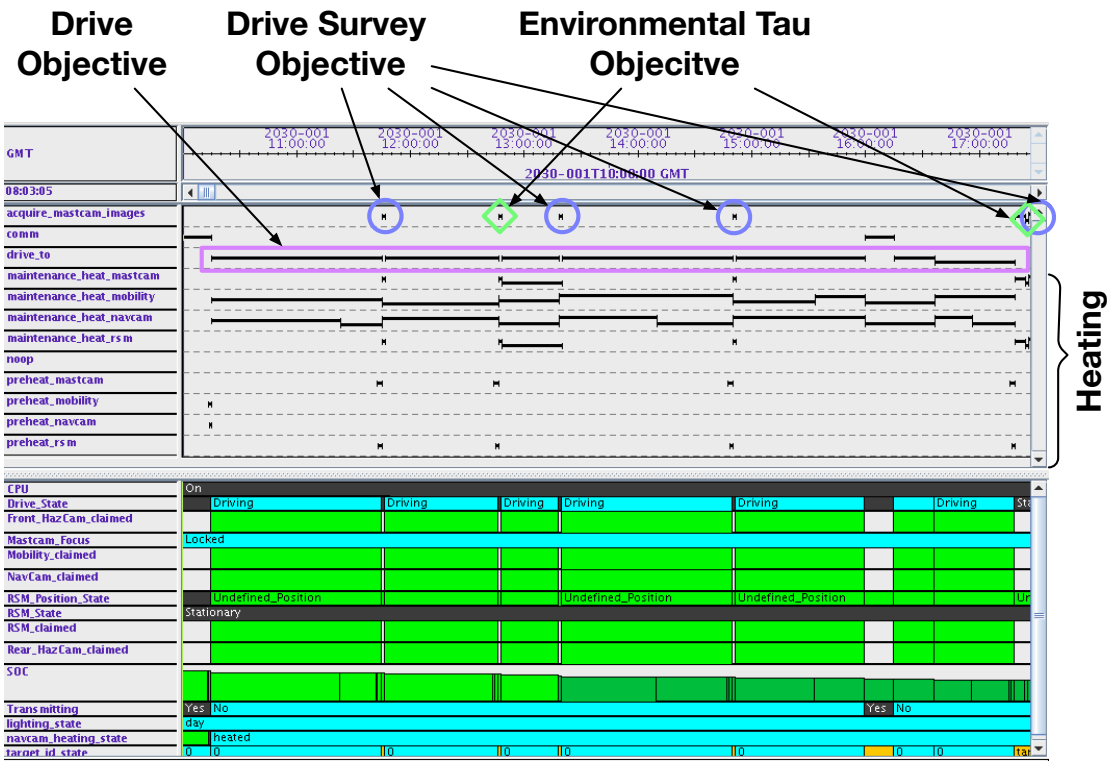

Figure 8: Example generated plan illustrating a long-range drive objective that was split up to support two different types of campaign objectives.

Once a poor-performing preferences is selected, an attempt is made to modify the plan to better satisfy the preference. For the periodic example, this would simply mean moving one of the activities closer to the other. As another example, consider a state-based campaign that is requesting observations every 100 meters as the rover drives. If a drive takes longer than expected, future observations can be postponed to better match the 100m separation preference. While local improvements to quality issues may be sub-optimal, the response time can be much shorter, making this method more suitable during execution.

Figure 8 shows an example plan generated by our system. The planning model is derived from the activity model used for MSL rover operations and includes important aspects of the mission such as science activities, communication windows and device heating. The example illustrates how the plan generator uses provided campaign relationships to coordinate rover activity.

For this example, the rover was given a long range drive objective along with two different campaign relationships: acquire environmental tau (atmospheric opacity measurement) observations every 3 hours (Temporal Campaign) and perform a mid-drive survey activity every 75 meters (State-Based Campaign). The resulting plan shows the drive objective being paused at different points to support interleaving these campaign activities.

Autonomous Goal Generation

In addition to plan generation and repair, the other use of campaign objectives in our system is to identify new goals for the system based on scientist guidance. This is applicable in cases where the operations team does not have up to date information of the area around the rover but want the rover to continue performing productivity activities. As discussed in the Background section, this situation can arise during restricted sol phases of the mission.

A high-level campaign goal can be used to generate more specific goals using onboard software. For example, suppose scientists are interested in remote-sensing, compositional measurements of a rock formation seen previously and known to exist in a region the rover is approaching. Using the TextureCam software (Thompson et al. 2012), scientists can train a model to detect the rock formation using labeled examples of the formation in previous navigation camera images (Figure 9, left). Then, upon driving into the new region, the rover can run TextureCam onboard to compute a probability map of the regions most likely to contain that rock formation (Figure 9, center). The probability map can be used to select the best locations for measurement, as well as the likelihood that each measurement satisfies the scientific intent of measuring the rock formation (Figure 9, right). Each measurement becomes a new goal, and the planner can use the probability information to reason about the tradeoffs between acting upon the various generated goals.

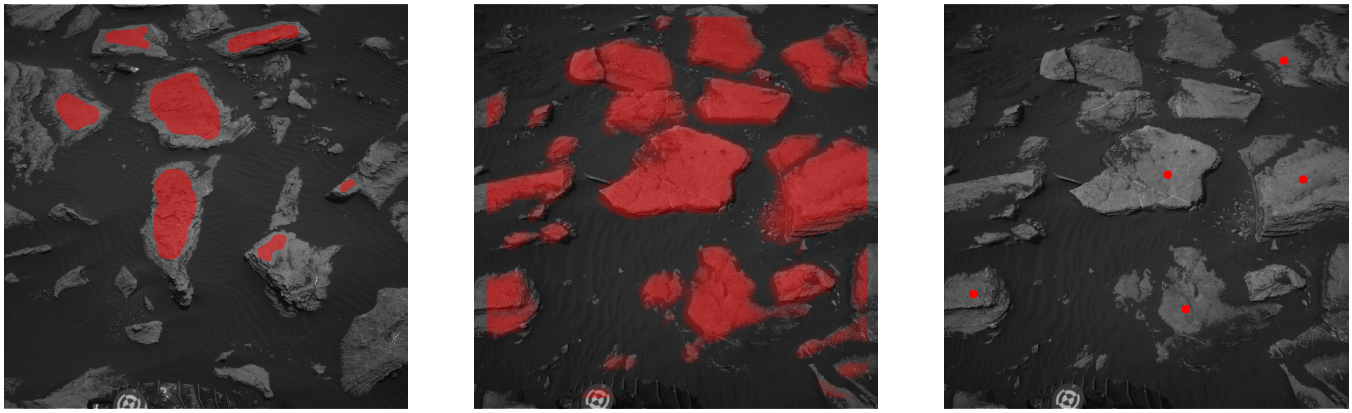

Figure 9: An example showing how scientists can use TextureCam to express intent to autonomously generate new goals on board. The left image shows hand-labeled regions of a geological formation of interest. The center image shows the estimated probabilities that regions in a new image are of the same formation, given a model trained from labels. The right image shows the top five software-selected locations for diverse observations of the rock formation, each corresponding to a new goal for the planning system.

Related Work

Shalin, Wales, & Bass, (2005) conducted a study of Mars Exploration Rovers operations to design a framework for expressing the intent for observations requested by the science teams. Their focus was the use of intent to coordinate planning among human operators and the resulting intent was not captured in a manner that would be conducive for machine interpretation. Our approach codifies some of the fields in their framework in a way suitable for the rover. In particular, the authors defined a "Related Observations" field as a way for scientists to identify relationships among different observations, which need not be in the same plan. Our work on campaign intent can be seen as a way of defining a specific semantics to these types of relationships to facilitate reasoning about these relationships by the rover.

Their framework also includes information that we agree is essential for effective communication among operators but that we do not currently express to the rover. For example, the "Scientific Hypotheses" field is used to indicate what high-level campaign objective is being accomplished by the requested observation. We are not yet providing these higher-level campaign objectives to the rover, though it is an interesting area of future research.

Mali (2016) views intent as a means for a user to place constraints on the types of plans a planner is allowed to produce such as only generating plans that have at most one instance of a class of actions or that plans must limit the use of a particular action. The primary role of our use of intent is to allow the planner to assess the value of achieving a given set of goals. However, some of our campaign intent does imply constraints and preferences on how, or more specifically, when goals are accomplished. For example, the periodic campaign intent specifies a timing relationship among goals and a preference on how close to comply with the desired timing.

There are some similarities between our campaign definitions and those used for Rosetta science planning (Chien et al. 2015). Both use campaigns to express requests for variable-sized groups of observations with relationships and priorities. Rosetta plans covered much longer time periods (e.g. weeks) and required more complex temporal patterns, such as repeating groups of observations. But observation patterns were primarily driven by the predictable trajectory of the spacecraft, allowing relationships to be expressed as temporal constraints. This is not sufficient for rovers, where many observations are dictated by the rover location and surrounding terrain, and the duration of many activities cannot be accurately predicted. State-based and goal set relationships more accurately represent some of the science intent found on surface missions.

There have been a variety of autonomous science systems deployed or proposed for rovers include AEGIS system running on the Opportunity and Curiosity rovers (Francis1 et al. 2016), and the SARA component proposed for an ExoMars rover (Woods et al. 2009). These systems allow the rover to identify targets in its surroundings that match scientist-provided criteria. The introduction of campaign relationships broadens the scope of the type of guidance that scientists can provide these systems, allowing scientists to express the amount of observations they would like for their different objectives along with the relative priorities of the high-level objectives.

The ProViScout project has similar objectives to our work (Paar et al. 2012). ProViScout is an integrated system to conduct planetary scouting and exploration. It includes autonomous science capabilities to enable onboard identification of science targets. Similar to our approach, the system selects follow-up observations for identified targets and submits these requests to an onboard planner to determine if there are sufficient resources to accomplish these new objectives. The campaign intent concepts we have developed would also be applicable to ProViScout as a way to increase the expressivity for providing scientist intent to the rover.

There is an active area of research in intent recogni-

tion (Sukthankar et al. 2014). The general goal of this area is to identify the objectives of other agents (human or otherwise) from observations of the agents' actions. In contrast, in our work, it is acceptable for users to explicitly identify their intent, rather than require the system to attempt to infer intent. Indeed, there is interest in the operations team to clearly document their intent for the purpose of communication among teams and as a record of what activity was planned for the rover and why. As such, rather than try to infer user intent, our objective is to increase the expressivity of the rover's interface in order to more closely reflect mission intent.

Conclusions

We have discussed a formalism for encoding aspects of science campaign intent in a manner suitable for reasoning by a rover. Campaign intent specifies relations among goals. These relations provide guidance to the rover to enable it to select a high-value subset of goals to accomplish among a set of goals that oversubscribe available resources. Further, they provide guidance to the rover when it identifies its own goals to work on.

We have begun implementing this campaign intent framework within the ASPEN system and integrating it with a research rover at JPL. Over the next years we will be conducting mission-relevant, multi-sol scenarios with the rover at the JPL Mars Yard to evaluate its ability to support productive operations with limited ground-in-the-loop interactions.

Acknowledgments

This research was conducted at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Jet Propulsion Laboratory Research and Technology Development program.

References

[Benowitz 2016] Benowitz, E. 2016. Concepts for the mars 2020 rover onboard scheduler. In *Proceedings of the 2016 Workshop on Spacecraft Flight Software (FSW-16)*.

[Chattopadhyay et al. 2014] Chattopadhyay, D.; Mishkin, A.; Allbaugh, A.; Cox, Z. N.; Lee, S. W.; Tan-Wang, G.; and Pyrzak, G. 2014. The Mars Science Laboratory supratactical process. In *Proceedings of the SpaceOps 2014 Conference*.

[Chien et al. 2005] Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davies, A.; Mandl, D.; Frye, S.; Trout, B.; Shulman, S.; and Boyer, D. 2005. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication* 2:196–216.

[Chien et al. 2015] Chien, S.; Rabideau, G.; Tran, D.; Doubleday, J.; Nespoli, F.; Ayucar, M.; Sitje, M.; Vallat, C.; Geiger, B.; Altobelli, N.; Fernandez, M.; Vallejo, F.; Andres, R.; and Kueppers, M. 2015. Activity-based scheduling of science campaigns for the rosetta orbiter. In *Proceedings of IJCAI 2015*.

[Francis1 et al. 2016] Francis1, R.; Estlin, T.; Gaines, D.; Doran, G.; Gasnault, O.; Johnstone, S.; Montano, S.; Mousset, V.; Verma, V.; Bornstein, B.; Burl, M.; Schaffer, S.; and Wiens, R. C. 2016. AEGIS intelligent targeting deployed for the Curiosity rover's ChemCam instrument. In *Proceedings of the 47th Lunar and Planetary Science Conference*.

[Frydenvang et al. 2016] Frydenvang, J.; Gasda, P.; Hurowitz, J.; Grotzinger, J.; Wiens, R.; Newsom, H.; Bridges, J.; Gasnault, O.; Maurice, S.; Fisk, M.; Ehlmann, B.; Watkins, J.; Stein, N.; Forni, O.; Mangold, N.; Cousin, A.; Clegg, S.; Anderson, R.; Payr, V.; Rapin, W.; Vaniman, D.; Morris, R.; Blake, D.; Gupta, S.; Sautter, V.; Meslin, P.-Y.; Edwards6, P.; M.Rice; Kinch, K.; Milliken, R.; Gellert, R.; Thompson, L.; Clark, B.; Edgett, K.; Sumner, D.; Fraeman, A.; Madsen, M.; Mitrofanov, I.; Jun, I.; Calef, F.; and Vasavada, A. 2016. Discovery of silica-rich lacustrine and eolian sedimentary rocks in Gale crater, Mars. In *Proceedings of the 47th Lunar and Planetary Science Conference*.

[Fukunaga et al. 1997] Fukunaga, A.; Rabideau, G.; Chien, S.; and Yan, D. 1997. Towards an application framework for automated planning and scheduling. In *International Symposium on Artificial Intelligence, Robotics and Automation for Space*.

[Gaines et al. 2016] Gaines, D.; Doran, G.; Justice, H.; Rabideau, G.; Schaffer, S.; Verma, V.; Wagstaff, K.; Vasavada, A.; Huffman, W.; Anderson, R.; Mackey, R.; and Estlin, T. 2016. Productivity challenges for Mars rover operations: A case study of Mars Science Laboratory operations. Technical Report D-97908, Jet Propulsion Laboratory.

[Gasda et al. 2016] Gasda, P. J.; Frydenvang, J.; Wiens, R. C.; Grotzinger, J. P.; Watkins, J. A.; Stein, N.; Edgett, K. S.; Newsom, H.; Clark, B.; Anderson, R.; Bridges, N.; Clegg, S.; and Maurice, S. 2016. Potential link between high-silica diagenetic features in both eolian and lacustrine rock units measured in Gale crater with MSL. In *Proceedings of the 47th Lunar and Planetary Science Conference*.

[Jet Propulsion Laboratory Press Release 2016a] Jet Propulsion Laboratory Press Release. 2016a. Boron in calcium sulfate vein at 'Catabola,' Mars. `http://mars.nasa.gov/multimedia/images/2016/boron-in-calcium-sulfate-vein-at-catabola-mars&s=2`.

[Jet Propulsion Laboratory Press Release 2016b] Jet Propulsion Laboratory Press Release. 2016b. Routine inspection of rover wheel wear and tear. `http://www.jpl.nasa.gov/spaceimages/details.php?id=PIA20334`.

[Litvak et al. 2013] Litvak, M.; Mitrofanov, I.; Behar, A.; Boynton, W.; Deflores, L.; Fedosov, F.; Golovin, D.; Hardgrove, C.; Harshman, K.; Jun, I.; Kozyrev, A.; Kuzmin, R.; Lisov, D.; Malakhov, A.; Milliken, R.; Mischna, M.; Moersch, J.; Mokrousov, M.; Nikiforov, S.; Sanin, A.; Shvetsov, V.; Starr, R.; Tate, C.; VİṪret'yakov; Varenikov, A.; and Vostrukhin, A. 2013. The water bulk distribution along MSL Curiosity traverse measured by dan instrument. In *Proceedings of the European Planetary Science Congress*, volume 8.

[Mali 2016] Mali, A. D. 2016. Expressing user intent in planning by instance rewriting. In *Proceedings of the 2016 IEEE 28th International Conference on Tools with Artificial Intelligence (ICTAI 2016)*.

[Mason et al. 2017] Mason, E.; Lemmon, M.; de la Torre, M.; and Smith, M. 2017. A quick look estimation of optical depth measurements from the rover environmental monitoring station ultraviolet sensors. In *Proceedings of the Sixth International Workshop on the Mars Atmosphere: Modelling and Observations*.

[Mishkin et al. 2006] Mishkin, A. H.; Limonadi, D.; Laubach, S. L.; and Bass, D. S. 2006. Working the Martian night shift: The MER surface operations process. *IEEE Robotics and Automation Magazine* 13(2):46–53.

[Muscettola et al. 1998] Muscettola, N.; Nayak, P. P.; Pell, B.; and Williams, B. C. 1998. Remote Agent: to boldly go where no AI system has gone before. *Artificial Intelligence* 103(1-2):5–47.

[Nachon et al. 2015] Nachon, M.; Mangold, N.; Cousin, A.; Forni, O.; Anderson, R. B.; Blank, J.; Calef, F.; Clegg, S.; Fabre, C.; Fisk, M.; Gasnault, O.; Kah, L.; Kronyak, R.; Lasue, J.; Meslin, P.-Y.; Moulic, S. L.; Maurice, S.; Oehler, D.; Payre, V.; Rapin, W.; Sumner, D.; Stack, K.; Schrder, S.; and Wiens, R. 2015. Diagenetic features analysed by ChemCam/Curiosity at Pahrump Hills, Gale crater, Mars. In *Proceedings of the European Planetary Science Congress*, volume 10.

[Paar et al. 2012] Paar, G.; Woods, M.; Gimkiewicz, C.; Labrosse, F.; Medina, A.; Tyler, L.; Barnes, D. P.; Fritz, G.; and Kapellos, K. 2012. PRoViScout: a planetary scouting rover demonstrator. In *Proceedings of SPIE Vol. 8301 ntelligent Robots and Computer Vision XXIX: Algorithms and Techniques*.

[Rabideau, Engelhardt, and Chien 2000] Rabideau, G.; Engelhardt, B.; and Chien, S. 2000. Using generic preferences to incrementally improve plan quality. In *International Conference on Artificial Intelligence Planning Systems (AIPS 2000)*.

[Shalin, Wales, and Bass 2005] Shalin, V. L.; Wales, R. C.; and Bass, D. S. 2005. Communicating intent for planning and scheduling tasks. In *Proceedings of HCI International*.

[Stack et al. 2015] Stack, K. M.; Grotzinger, J. P.; Gupta, S.; Kah, L. C.; Lewis, K. W.; McBride, M. J.; Minitti, M. E.; Rubin, D. M.; Schieber, J.; Sumner, D. Y.; Beek, L. M. T. J. V.; Vasavada, A. R.; and Yingst, R. A. 2015. Sedimentology and stratigraphy of the Pahrump Hills outcrop, lower Mount Sharp, Gale Crater, Mars. In *Proceedings of the 46th Lunar and Planetary Science Conference*.

[Sukthankar et al. 2014] Sukthankar, G.; Goldman, R. P.; Geib, C.; Pynadath, D. V.; and Bui, H. H. 2014. An introduction to plan, activity, and intent recognition. In Sukthankar, G.; Goldman, R.; Geib, C.; Pynadath, D.; and Bui, H. H., eds., *Plan, Activity, and Intent Recognition,*. Elsevier.

[Thompson et al. 2012] Thompson, D. R.; Abbey, W.; Allwood, A.; Bekker, D.; Bornstein, B.; Cabrol, N. A.; Castano, R.; Estlin, T.; Fuchs, T.; and Wagstaff, K. L. 2012. Smart cameras for remote science survey. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space*.

[Webster et al. 2015] Webster, C. R.; Mahaffy, P. R.; Atreya, S. K.; Flesch, G. J.; Mischna, M. A.; Meslin, P.-Y.; Farley, K. A.; Conrad, P. G.; Christensen, L. E.; Pavlov, A. A.; Martn-Torres, J.; Zorzano, M.-P.; McConnochie, T. H.; Owen, T.; Eigenbrode, J. L.; Glavin, D. P.; Steele, A.; Malespin, C. A.; Jr., P. D. A.; Sutter, B.; Coll1, P.; Freissinet, C.; McKay, C. P.; Moores, J. E.; Schwenzer, S. P.; Bridges, J. C.; Navarro-Gonzalez, R.; Gellert, R.; and Lemmon, M. T. 2015. Mars methane detection and variability at Gale crater. *Science* 347(6220):415–417.

[Woods et al. 2009] Woods, M.; Shaw, A.; Barnes, D.; Price, D.; Long, D.; and Pullan, D. 2009. Autonomous science for an ExoMars roverlike mission. *Journal of Field Robotics* 26(4):358–390.

[Yingst et al. 2013] Yingst, R. A.; Goetz, W.; Hamilton, V.; Hipkin, V.; Kah, L.; Madsen, M.; Newsom, H.; Williams, R.; Bridges, J.; Frias, J. M.; and King, P. 2013. Characteristics of pebble and cobble-sized clasts along the Curiosity rover traverse from Sol 0 to 90. In *Proceedings of the 44th Lunar and Planetary Science Conference*.