Communicating Intent for Planning and Scheduling Tasks

Valerie L. Shalin
Wright State University
Dayton, OH USA 45435
valerie.shalin@wright.edu

Roxana C. Wales
SAIC NASA Ames Research Center
Moffett Field, CA USA 94035
rwales@mail.arc.nasa.gov

Deborah S. Bass
NASA Jet Propulsion Laboratory
Pasadena, CA USA 91101
Deborah.S.Bass@jpl.nasa.gov

Abstract

This paper concerns the relation between people, technology and cognitive work, in planning and scheduling remote, temporally extended telerobotic action. We combined a computational theory of planning with ethnographic methods to design a framework for expressing the intent behind requested science observations for the 2004 Mars Exploration Rover mission. The evolution and usage of this framework substantiates a distinction between the purposes and methods of work, consistent with a computational view of planning. However, we also identify several other properties of intent that acknowledge it's role as a boundary object in coordinating collaborative planning, including the identity of the source of a request and it's priority. In addition, scientists developed artifacts to compensate for the limitations of the intent framework, including the ability to express plans that unfold across arbitrary units of time, and objects of scientific interest that have spatial extent.

1 Introduction

Much of the research literature on human computer interaction in dynamic domains assumes that the computer provides telemetry regarding a current situation, while the operator executes a physical procedure that changes this situation. The human operator evaluates the match between a representation of the current state and a representation of the goal. If the states match, the comparison will be re-executed after some period of time. If the states do not match, a physical procedure changes the state of the world. The role of the display is to mediate between the execution of a procedure and perception in the world for human controllers (Norman, 1990).

In many domains, the computer also participates in the process that defines the goal, or the planning process (Shalin & McCraw, 2003). Shalin (in press) describes the processes in computer-based real-time planning, in which the plan requires periodic re-synchronization with incoming telemetry about real-world conditions. In the present paper. we examine the processes associated with off-line, albeit time-pressured planning, with a single opportunity to evaluate the plan for feasibility. We emphasize the need to support multiple stakeholders, or special interest groups, and the type of plan representation that supports plan evaluation and modification in a distributed, collaborative work setting. This process is illustrated in Figure 1.

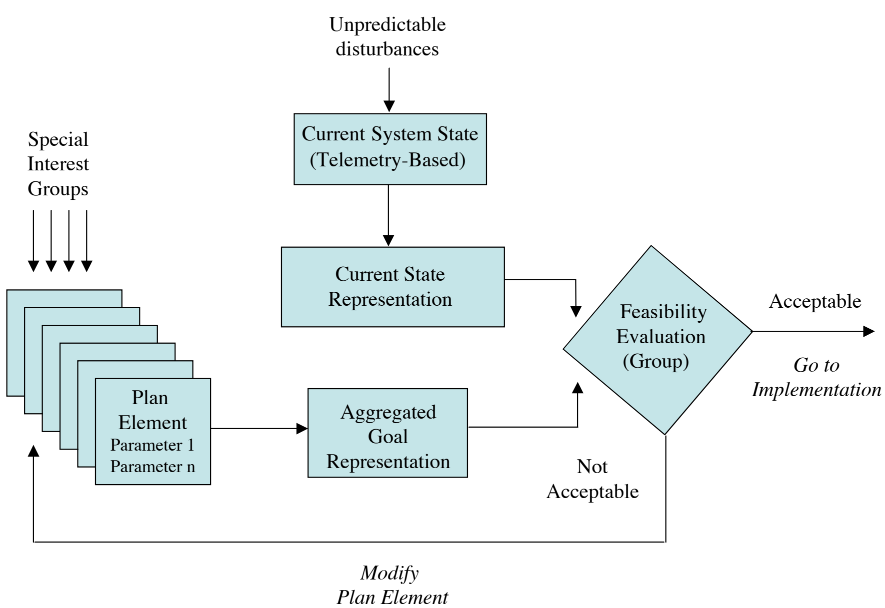

Figure 1: Comparing an aggregated goal representation with current state to determine feasibility.

Starting at the left of the figure, special interest groups submit separate plan elements for aggregation as a group goal. The group evaluates this aggregated goal for feasibility in the current circumstances, and in the event that the goal is not feasible, must modify the goal. While it might be possible to adjust the parameters within individual plan elements in order to generate a more feasible aggregated goal, the number of parameters is likely greater than the number of plan elements. Because the number of possible adjustments is a factorial function of the number of candidate elements for adjustment, modification by deleting an entire plan element is substantially more efficient than modification of individual parameters.

Figure 1 illustrates the science planning process that the authors helped design for NASA's 2004 Mars Exploration Rover Mission. In the body of this paper we describe the representations of plan elements, (e.g, take a picture of a rock) that supported this process and some of the lessons learned from practice. We begin with a brief description of the relevant background, including the Mars Exploration Rover mission and computational theories of planning.

2 Background

2.1 Mars Exploration Rover Mission

NASA's Mars Exploration Rover (MER) Project successfully landed two spacecraft carrying identical rovers on opposite sides of the surface of Mars in January 2004 and conducted two nominal surface missions for 90 Martian days, called 'sols'. The scientific purpose of this mission was to ascertain the presence of water and its influence on the Martian terrain. Each day a team of planetary scientists and rover engineers produced a planned set of activities and associated commands that were sent to the rovers so that they could collect data on the following morning. The 19-hour planning process began with the receipt of data from the spacecraft and scientific evaluation. Scientists then designed a set of observations that they wanted the rover to conduct. An engineering team implemented this set in rover command language, radiated to the spacecraft for execution the following morning.

Wales (in preparation) designed the science facilities used during the planning process. Figure 2 (left) shows science team members, seated in their Theme groups, in the process of determining the focus for the next sol. Signs hung above clusters of tables identified the Theme group in question. Following a general overview, groups and individuals refined proposed scientific observations considering the available resources, timing restrictions on the use of instruments and the presence of any engineering restrictions.

Figure 2: *(left)* Science team members identify the focus of the plan for the day, before splitting into small design groups. *(right)* Science team members negotiate final plan contents in the SOWG meeting before handing the plan to the engineering team for implementation.

Figure 2 (right) shows the science team (called the Science Operations Working Group, or SOWG) conducting negotiations and modifications on the final content for uplinking. The SOWG is seated around a large U-shaped table, with large screen displays at the front and side of the table. Pictures from the Martian surface appear on two of the screens. In a typical SOWG meeting the remaining screens contained a representation of proposed activities, generated during the theme group planning process.

The design of the representation for the intent behind a science plan and its elements is the focus of the present paper. The challenge of documenting intent arises for two reasons. First, scientists cannot express their requests in the low level programming language that the Rover uses. Instead, scientists must express their requests at a higher level of abstraction, leaving implementation (and hence implementation details) to specialists with whom they might have little interaction. Further, although requests may be conceptualized in isolation, they will be executed with other requests, and in a particular context. The other requests and the prevailing context may require modification of the original request, while preserving its original intent. A representation of intent is therefore critical to the accomplishment of scientific goals.

2.2 Computational Theories of Planning

The theoretical goal of this work was to examine and extend a pre-existing approach to the representation of intent (Sewell & Geddes, 1990), based on skeletal planning in MOLGEN (Friedland & Kedes, 1985), adapted to enable computational associate systems to monitor human intent. For the purposes of this paper, this representation will be referred to as "the motivating representation". Friedland used this representation specifically to design experiments in molecular biology. Shalin et al. (1997) has incorporated this representation into numerous task analyses, where it has been successfully used to describe performance, develop tests of job knowledge and design interfaces. This project comprises the first attempt to use the representation as a boundary object (Star, 1989) to facilitate communication between human collaborators with different areas of expertise.

The expectation from the present effort was not that humans would learn a computer code, but rather that lessons learned in the formal representation of plans might help in the identification of issues in the expression of intent. Some of the lessons previously learned from the approach adopted here include:

1. **A distinction between purposes and methods.** Purposes correspond to desired states of world. Methods correspond to procedures for achieving purposes. The distinction permits the collection of alternative methods for achieving the same purpose, which otherwise may have nothing at all in common. As an example of this *many-to-one* relationship, going to a restaurant and cooking at home are two methods for satisfying hunger. The distinction between purpose and method also acknowledges the potential relationship between one method and multiple purposes. This *one-to-many* relationship between a method and purposes is substantiated by the common experience of initiating an action (e.g., going to the kitchen), and forgetting the reason for doing so (was it to check on a roast or get a glass of water?). Commanding the Mars Rover is method based: Humans tell the Rover exactly what to do. If commands were purpose based the Rover would decide among alternative methods.

2. **The need for multiple levels of abstraction**. Events can be represented at many different levels of abstraction. No single level is best for expressing meaning. Instead a task is situated within a multi-level hierarchy, which Geddes called a *plan-goal graph*. As a graph implies, intent has no context independent, basic-level grain size, particularly concerning time. At one level of analysis, intent may be executed instantly, while others could take hours, days, weeks or years.

3. **The expression of pre-and post-conditions**. Each method includes pre-and post-conditions to assist in the control and sequencing of action. Pre-conditions include the initial state of a system. When used to control action, pre-conditions consist of qualitative properties of the world that have become true upon the execution of a preceding action, which are then available to support the execution of a following action. By commanding the Rover via methods, humans retain responsibility for monitoring qualitative conditions in the environment.

3 Research Approach

The domain posed a both a practical and theoretical challenge. The practical challenges included the novelty of the domain for the scientists. Although they were internationally known experts in their field, telerobotic operations, bringing laboratory instruments to the Martian surface, lacked even an Earth analogue. As a result, these scientists did not have an established practice for organizing or describing the necessary work. Further, our recommendations for describing requested science observations had to fit within a tool, Science Activity Planner (SAP) that was already being developed for reviewing photographic data and pointing the robotic instruments. From a theoretical perspective, we know of no other attempts to bridge the gap between the intent of an activity and the code that will be used to conduct that activity. We combined computational theories of planning described above with ethnographic methods to identify the required representation for the intent behind a science request.

Ethnographic methods allow researchers to focus their attentions from a variety of social, cognitive and technical perspectives that mirror the complexity of a domain. (Wales et al., 2002; Forsythe, 1999; Jordan, 1996; Nardi, 1996; Bloomberg et al., 1993). Participant observation is a primary data collection method of ethnographic work. One of the authors was the science operations systems engineer and later the deputy science team chief for MER and had daily access to on-going mission design work. The other authors were tasked by NASA to provide human-centered computing work systems design recommendations to the mission, spending extended periods of time over three years at the Jet Propulsion Laboratory (JPL). Our team brought a cross-disciplinary perspective to the research, drawing on backgrounds in geology, cognitive psychology and cultural psychology/anthropology.

We learned many of the intricacies of the rover instruments. After each test and training opportunity, we analyzed the data from SAP as well as from field notes to develop an understanding of the cognitive, linguistic, referential and software needs relative to science intent. During the mission, we continued our data collection, taking field notes and photographs, making video tapes of meetings and collecting copies of mission science activity planning print outs. In the present paper we describe the approach to documenting intent that we implemented and focus on the additional compensatory artifacts that scientists developed to assist in the expression of intent.

4 Results

In the following subsections we first describe the initial representation for intent that we examined in a 2002 field test. We identify modifications to this representation for the 2004 Mission, and then provide evidence regarding its adequacy for the communication of intent during the Mission.

4.1 Initial Intent Representation

Using planning tools designed for this mission, scientists submitted requests called observations, consisting of a set of activities. Activities are specific building blocks that translate into computer commands the spacecraft will understand. Activities are defined in an Activity Dictionary, which specifies a standard set of parameters with a range of values that must be identified for each instantiation. Activities also link to resource models, so that scientists can examine the demands of their requests.

Table 1 presents fields or attributes contained in a first attempt at a representation for the intent behind a requested observation, drawn from the motivating representation. However a number of properties of the representation arrived at during the course of this project were not theoretically motivated. JPL's Science Operations Support Team proposed some of these (Table 2). Further, some of the originally proposed fields were not retained (Table 3).

Table 1
Retained Fields Drawn from the Motivating Representation

Resulting Information/Purpose—(e.g., soil particle size) Corresponds to the immediate low level goal associated with the method in question. "Resulting information" helped focus attention on the appropriate level of analysis.
Scientific Hypotheses—(e.g., Determination of soil and rock thermophysical properties) Provides a placeholder for indicating the "higher-order" purposes behind the observation, and by contrast, helps to define resulting information as a "lower-order" purpose.
Notes—Alerts downstream instrument specialists of assumptions and preconditions associated with the current request. This field subsumed an originally proposed field called method, because the domain experts could not grasp what a method might be.
Related observations—Allows for the linking of observations across sols. This field allows for the grouping of multiple subgoals in a method. Thus, related observations could be used to group together the two parts of a thermal inertia study (during hot and cold periods) that happened to be described in two observations, for example, across two sols.

Table 2 presents the additional fields that JPL engineer/scientists suggested. These fields expose interesting oversights in the motivating representation for intent. The role of an intent representation as a boundary object in distributed work necessitates some of these fields, such as originators and points of contact, and names for the observation. The priority field reflects the need to plan in the absence of good information about preconditions and resources, and anticipates the negotiations that will occur later. In contrast, the motivating representation reflects an assumption that intent is whittled down to something manageable by known and available pre-conditions. That is, the set of intents is assumed realistic—a decidedly unrealistic assumption in the present context.

Table 2

Intent Fields not Originating from Motivating Representation

Plan ID—Identifies the SAP directory structure containing the observation in question.
Custodian— Identifies a primary point of contact for clarifying the contents of the requests if necessary.
Theme Group—Identifies the design group membership of the requestor. This field allows the same custodian to be a member of multiple theme groups. This additional field may help to identify the purpose of the request in the absence of more explicit indications.
Observation Name—Provides a meaningful file name for each request, so that scientists can anticipate the contents of a file without opening it up to review it's contents.
Priority—Identifies the initial priority for a group.
Creation Date—Assists readers in recalling the context that motivated the request, and may be important for interpreting the request or the resulting data.
Downlink Rationale—Justifies Direct to Earth Transmission, for example to obtain results needed for planning next sol.

The field "Plan ID" in Table 2 reveals an assumed constraint in the design of the planning system. The directory structures are organized by sol, and theme group within a sol. A directory structure based on sol makes it extremely cumbersome to express an intent that is executed across sols, because the individual pieces must be spread across different directories. This example points to the influence that a presumed unit of analysis has on the nature of human-machine interaction.

Table 3 presents the fields that we suggested, but resulted in controversy, largely because they appeared to complicate the representation. Some of them have been since reinstated, supporting the view that it is not the representation of intent that is complicated, but the intent itself. A method field, key to the motivating representation, was originally questioned. Later in this paper we present the evidence that caused it to be reinstated.

Table 3

Suspended Controversial Intent Fields

Method—Indicates the general manner in which instruments are to be combined in the observations e.g., compare thermal emission of soil in crater at coldest and warmest ambient temperature
SOWG Priority/STG Priority—These two fields acknowledged the potential differences in priority between the source Science Theme Group and the Science Operations Working Group as a whole.
Status—Enables the theme groups to track the status of their requests.
SOWG Chair approval—Acknowledges the conclusion of the SOWG meeting.
PUL approval—Acknowledges the approval of the individual instrument specialists who are responsible for implementing a request.

The suggestions included a distinction between Theme Group Priority and SOWG priority, but this distinction was viewed as overly complex. As indicated above, a Priority field was retained, but it is not clear whose priority it reflects, or how a difference in theme group and SOWG priority should be represented. The suggestions also included an explicit representation for the status of the observation, acknowledging its role as a boundary object that moves between groups. This suggestion addresses a practical problem known as round trip data tracking, in which the availability of data in response to a request is made explicit. As of this writing, during the extended operations phase of the mission, neither the round-trip data-tracking problem nor the larger status problem has been addressed. Administrative approval fields were originally suggested by JPL personnel, and determined to be unnecessary.

We also recommended a representation for the temporal relationship between activities within an observation, based on Allen (1983). Temporal representations provide shorthand for the qualitative results of prior action, and enable control and sequencing. Initially suspended, Allen's temporal calculus was later incorporated into a different piece of software called the constraint editor, which serves as an interface to a computational scheduling system.

3.2 Field Test Results

The 2002 Field Test provided an opportunity for scientists to use the intent fields while they formulated requests for observation. This field test involved little training on the intent frame fields in order to examine natural responses. A primary conclusion resulting from this analysis was the need for a method field—the field that had been originally deleted—because methodological information was being distributed unsystematically over the rest of the fields. Of the five theme groups, only the Atmospheric scientists regularly utilized the intent fields. Their usage is considered in substantial detail below.

Name—The name field was one of the first to be consistently completed. The intuitive names seemed to correspond to method, instrument or purpose. Specific references to method included "cloud movie", "navcam movie for clouds and dust devils" and "sky-survey".

Purpose—The purpose field was also among the first to be consistently completed. Example purposes included: "Watch the clouds move". Some of the examples appeared to focus on resulting information e.g., "navcam pairs show cloud presence and motion", "navcam movie for clouds and direction". Other entries clearly mixed resulting information and method. For example "aerosol properties anti-Sun" indicates both the results and the method (do this anti-Sun). " Determine water vapor with 2 IPS spectra (4*1000 coadds) and support Navcam" and "Navcam pairs show cloud presence and motion" also suffers from intermingling purpose and method.

Downlink Rationale Field—Scientists began to exploit this field mid-test. The examples reflected an appropriate appreciation for the timely delivery of data (DTE Direct to Earth is faster) and the limited bandwidth available for the most expeditious delivery, e.g., "Important for planning, but 2.4 Mbits". Image data often required timely delivery, to assist in planning the next sol's activities. This is consistent with field practice, in which the spatial layout is always available for planning purposes, but the results from analytic laboratory instruments are not.

Related Observations—Scientists also began to exploit this field mid-test. The interesting results from this field concern the basis of relationship. In some cases the observations are related by time, e.g., "Sky IPS data on other sols to see time variation." In other cases the observation functions as a caveat on the interpretation of another observation, e.g., "Check on rover pointing ability." Related observations may have backward referents e.g., "Sol 14 accidental wind detection" or forward referents, e.g., "Tests for later observations".

Scientific Hypotheses Field—Scientists began using this field mid-test. Some entries suggest they understood a distinction between purpose and hypothesis, e.g., Cumulus brought in by monsoonal flow indicates flow patterns. However, other entries were contaminated with methodological content: 1) There is time variable atmos. Absorption visible to IPS; 2) There is cloud activity and motions will give wind speed; Dust sedimentation can be tracked with target images; dust spectra indicative composition. An explanation for the appearance of methodological information here is that there was no obvious location for such information.

Notes Field—Scientists began to use the notes field early, but usage increased mid-test. One rather unusual note identified an observation as an Education and Public Outreach opportunity, i.e., additional purpose beyond scientific purpose and hypothesis. More typically, notes address pre-and post-conditions, method, and constraints on modifications, e.g., "Seq. Assumes IPS on and uses relative pointing. (do not rearrange)."

These inputs indicate that at least some scientists find methodological information relevant and meaningful. However, in the absence of a "method" field, similar information also appeared under the "purpose" field, and to some extent under the "scientific hypothesis" field. As a result, consumers of intent fields would not know where to look in the frame for methodological information. The prevalence of this information and its unpredictable location supported the reinstatement of a "method" field in the mission version of the framework for documenting intent.

3.3 Mission Results

The 2004 Mission results were mixed. Figure 3 below illustrates the intent fields filled out for one observation, by a theme group that was unusually compliant. However, compliance tendencies established during the field test transferred to the mission, and most theme groups varied from little to some compliance. The recipients of intent information, instrument specialists, accommodated this practice by arriving several hours earlier than scheduled, in order to hear discussions about the observations first hand. Under these circumstances, we suspect that well-known Gricean principles in communication conflict with the documentation of intent (Grice, 1975). People simply do not like to make something explicit that is obvious in current context; being unnecessarily detailed and specific could be considered inefficient and even disrespectful. Indeed, during the mission we observed the science team chiding one theme group about the scope of their detailed intent description of an observation—suggesting that it was an abstract for publication!

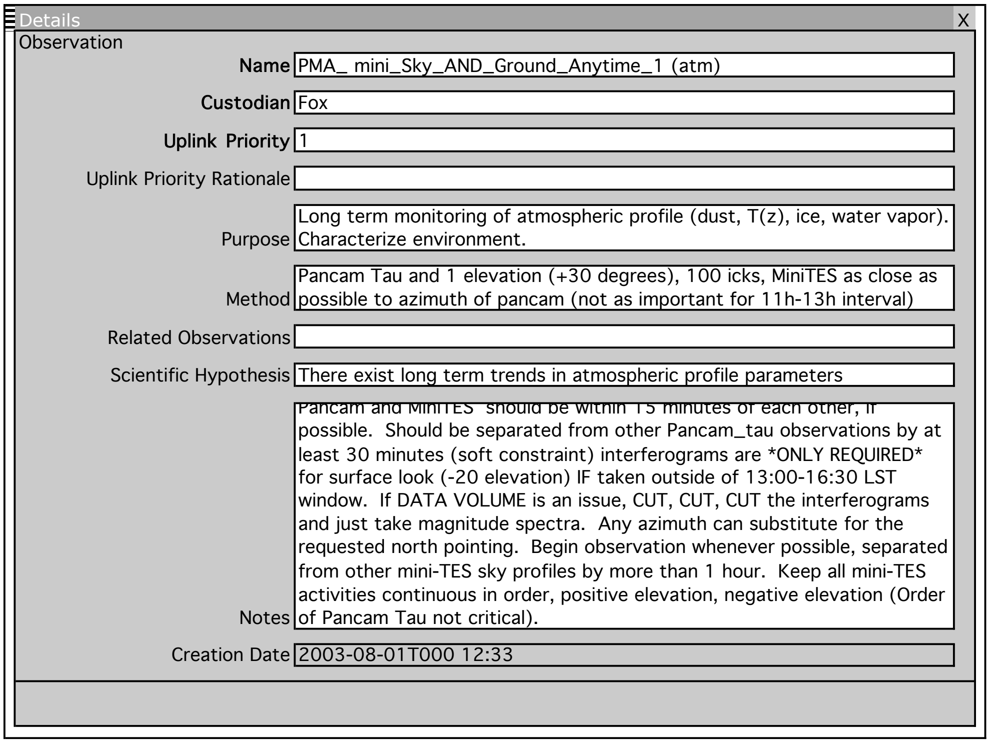
Details X

Observation

Name PMA_ mini_Sky_AND_Ground_Anytime_1 (atm)

Custodian Fox

Uplink Priority 1

Uplink Priority Rationale

Purpose: Long term monitoring of atmospheric profile (dust, T(z), ice, water vapor). Characterize environment.

Method: Pancam Tau and 1 elevation (+30 degrees), 100 icks, MiniTES as close as possible to azimuth of pancam (not as important for 11h-13h interval)

Related Observations

Scientific Hypothesis: There exist long term trends in atmospheric profile parameters

Notes: Pancam and MiniTES should be within 15 minutes of each other, if possible. Should be separated from other Pancam_tau observations by at least 30 minutes (soft constraint) interferograms are *ONLY REQUIRED* for surface look (-20 elevation) IF taken outside of 13:00-16:30 LST window. If DATA VOLUME is an issue, CUT, CUT, CUT the interferograms and just take magnitude spectra. Any azimuth can substitute for the requested north pointing. Begin observation whenever possible, separated from other mini-TES sky profiles by more than 1 hour. Keep all mini-TES activities continuous in order, positive elevation, negative elevation (Order of Pancam Tau not critical).

Creation Date 2003-08-01T000 12:33

Figure 3: An example of a complete intent description.

Figure 4 below illustrates a full candidate plan assembled by the science team. The numbered directory structure at the left indicates that plans are stored inside of sols. The large panel at the right consists of observations, and activities nested inside of them. Each observation has a place for documenting intent, as in Figure 3. The remaining columns illustrate the priority of the observation (and the activities that it includes), the resource requirements of the observations (and activities) and the purpose of an observation, an excerpt from the intent representation.

The behavior we observed provides an opportunity for critical evaluation, so that future representations of intent will reflect an improved understanding. We identify three trends below pertinent to the documentation of intent: the transfer of intent information to observation names; the prevalence of multi-sol science planning; and the impact of missing feature and target semantics.

Transfer of intent information to observation names. We address the intricate issues around observation naming in a different paper (Wales, Shalin & Bass, in preparation). Here we note a qualitative trend to replace textual descriptions within the intent fields with abbreviated shorthand in the names. The initial confusion over method gave way in some cases to the establishment of methodological conventions, and often, names for these methods, such as “dirt-taster”, “mini-Mini”, “dust-movie”. This development of method names and their appearance in observation names was predicted prior to the mission, and it does account for some of the likely reduction in word count associated with intent documentation.

Multi-sol science planning. The sol-based directory did not provide a convenient representation for plans that required several sols to execute. Several different types of situations emerged on the mission to reinforce the need

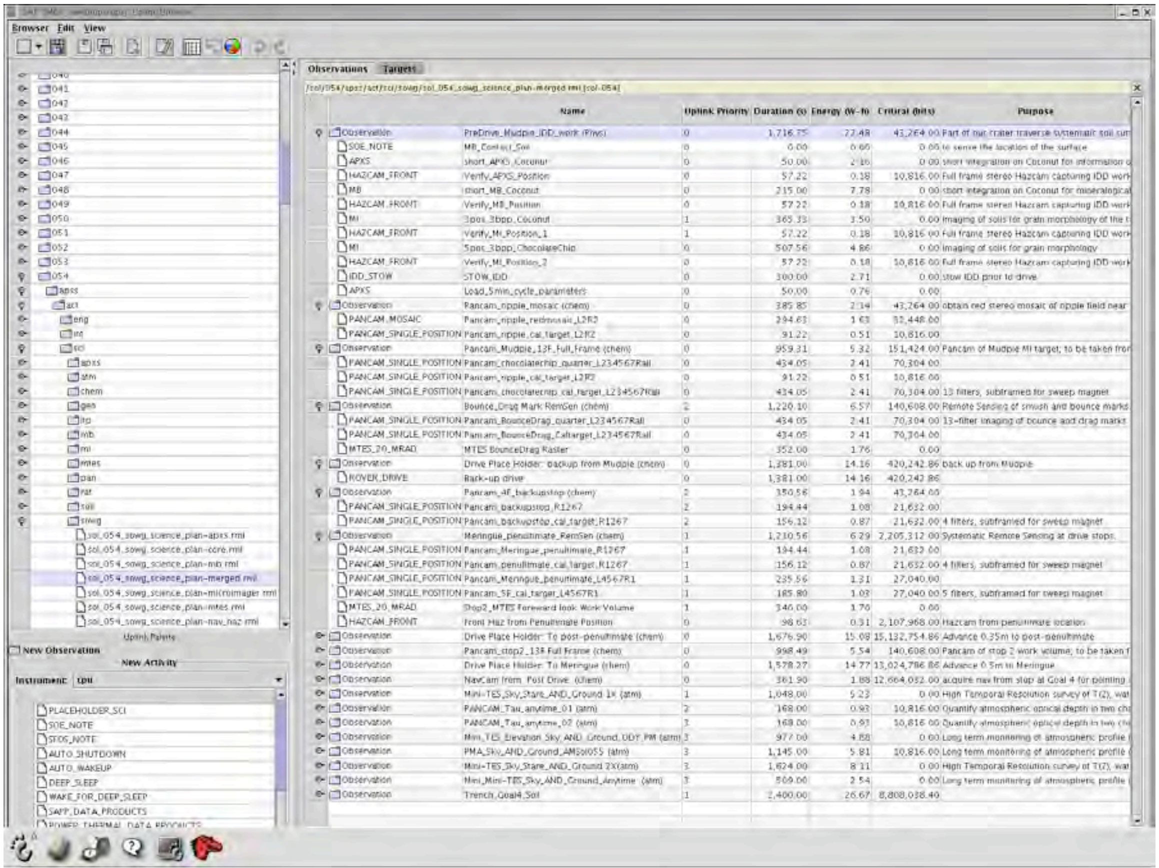

Figure 4: Science Activity Planner (SAP) Tool: screen shot shows a science activity plan as created by the science team with the higher order observations. Open toggles on some observations show subordinate activities that instantiate the observation.

for multi-sol planning tools. In the absence of these tools, the work of each sol appears incoherent. One situation that requires multi-sol planning is the telerobotic campaign. The campaign unfolds over a several days, from the identification of a target of interest, to approaching it, obtaining imagery, abrading it and obtaining further imagery. The sol-by-sol description of plans makes it difficult to see relationships between these observations, particularly if the campaign involves more than one target. A second situation that requires multi-sol planning is the need to record liens. Should the science team delete an observation for the current sol, they may wish to add it to the list of candidates for the next sol. This is cumbersome, and requires navigation across the highly nested sol-based directory structure. A third situation that requires multi-sol planning is the need for repeated measurements, in order to systematically characterize the environment, or to control for artifacts of an individual observation. During extended operations the science team developed a four-sol template, designed in part to address this problem (see Figure 5, right). Though effective, it is a classic work-around that does not solve the general case. It is impossible to specify, within this structure, requests with a frequency less than every four days. As a result, scientists also required representations such as the one shown on the left panel in Figure 5. This finding is completely consistent with the need for multiple levels of abstraction, identified in the motivating representation.

Target and feature semantics. The planning software required scientists to associate named targets for an observation with named features, corresponding to a single pixels in imagery. However, the software did not provide any means for identifying the semantics behind these (e.g., whether the feature is a rock or a soil patch, and the identity of the target, e.g., a dark spot, raised area, etc.) During the nominal mission, scientists acquired the habit of hand-carrying printed imagery to the instrument specialists, or rover drivers in order to point to the target while offering a description, thereby taking advantage of context and non-verbal behavior. This pointing refers to more than just a point---it refers to regions and areas. To help remote participants track targets and features, scientists developed alternative to single pixel pointers, which they included in their own reports, such as the ones shown in Figure 6.

Measurement	Desired Frequency	Last taken	Notes
Pancam of capture and filter magnets	10-14 sols	104	P2113
ISC/VOC measurement	10 sols	99	P2102
Pancam of Mini-TES eternal cal target	10 sols	99	P2104
Mini-TES of external cal target	10 sols	104	P2390. 1.4 Mbits, 224 s Do within 1 day of above pancam
Pancam of sweep magnet	10 – 14 sols	104	P2123
Pancam Survey Sky	14-20 sols	104	P2619 + p2119
Pancam Geofilter Opacity (contained with “Sky Survey”)	4-10 sols	106	P2617 + p2117

Green — OK; Red— Due or Overdue

total min:	277		DRIVE SOL 1	total W-hrs: 104
total data:		88		
position	**min**	**Mbit**	**observations**	
pre-drive	11	1.4	PMA mini sky and ground (atm)	3.5
	1	2.1	clast survey, rover front (geo)	0.4
	2	4.2	nearfield clast survey, side-looking (geo)	0.6
	6	6.4	Pancam, rear-look, science&mapping (ltp)	1.9
drive	120	9.1	Blind (30min) and auto-nav (90min) drives	84
mid-drive	1	2.1	Clast survey, rover front (geo)	0.4
	11	1.4	PMA mini sky and ground (atm)	3.5
	9	8.5	Navcam mapping&science (ltp)	2.7
	8	6.3	Penultimate Navcam&Hazcam; ultimate Hazcam	0.9
post-drive	10	12.9	Pancam drive direction, 4x1 (ltp)	5.3
	6	6.3	Navcam front3X1,drive,targets,science&mapping (ltp)	1.9
	12	12.7	Navcam rear7X1,targets,science&mapping (ltp)	3.7
	6	9.0	mid-field rock survey, forward looking (geo)	2.2
	6	2.5	systematic soil front; Mini-TES (min)	1.7
	20	1.0	Mini-TES recon raster (reduce if necessary) (min)	5.9
	30		Siesta	-20
	19	1.9	PMA sky and ground (do during ODY PM?) (atm)	5.8
Color key:				
Only observation or set of observations of its kind in entire quartet. Don't delete!				
Drive period				
Linked to Sol 3 activity			Linked to Sol 2 activity	29-Apr-08

Figure 5: Two alternative representations used to capture intent across sols. . The panel on the left covers an unlimited time span. The panel on the right is one of four different skeletons in a quartet than spans four sols.

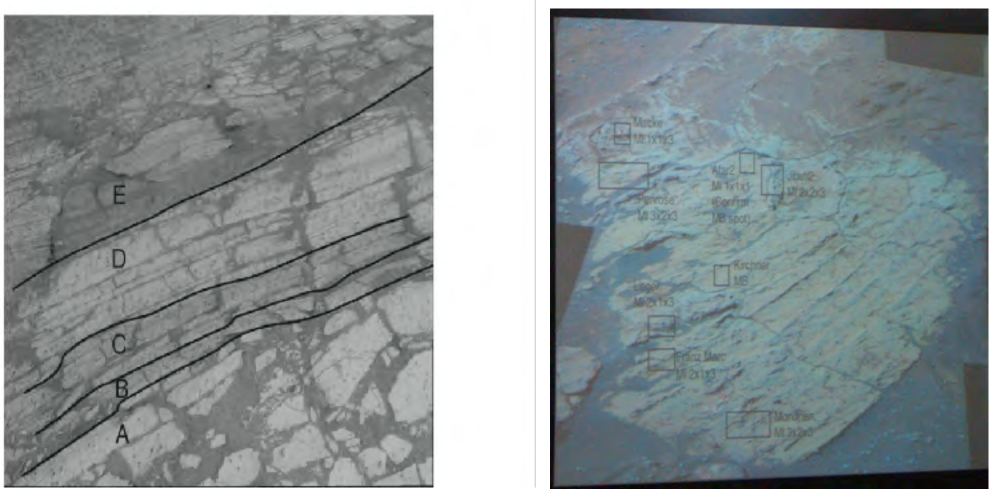

Figure 6: Imagery augmented to indicate areas of interest.

4 Conclusions & Implications

A computational theory of planning inspired the current approach to the documentation of intent behind science observations. The most important distinction in work, between the purposes and methods was reinforced, by the restoration of the method distinction as a field in the intent template, and by the emergence of method names. Purpose, scientific hypotheses, notes and related observations were also used, but to varying degrees.

A number of fields were added that did not reflect the motivating representation, but appear necessary when the intent fields are appreciated as a boundary object circulating among different science and engineering specialties. Said differently, the numerous previous applications of the plan-goal graph did not require these fields, but when the contents were used to guide distributed planning, fields acknowledging individuals and groups became necessary. Further, a priority field acknowledged the lack of good information regarding available resources, and facilitated negotiation among interested parties, revealing two important and unrealistic assumptions prevalent in computational planning: 1) We know the state of the world when we begin planning and 2) Goals do not conflict.

Two broad classes of deficiency emerged in the approach that we used to documenting intent: unit of analysis and limitations of text-based medium for the expression of intent in collocated domains.

Unit of analysis. Characterizing intent as the property of an observation within a sol, without capturing its surrounding context, as part of a campaign, a lien or a repeated cycle is insufficient. We know this because scientists invented their own representations to compensate for these limitations. Similarly, we found that identifying targets and features as individual pixels does not correspond to scientific thinking. Again, we know this because scientists invented their own representations to compensate for these limitations in the expression of context. Both cases exemplify issues that we recognized early on. We were unable to impact the software that imposed these restrictions, however, for two likely reasons: First, the software in question was well under development when we made our suggestions. Second rover code requires a target and field of view, not representations of objects.

Limitations of text-based medium for the expression of intent in a co-located domain. When engaged in an activity, many participants do not see the need to document their intent in text when the intended recipients or their representatives are present. We suspect that this characteristic of human communication will challenge the temporally and geographically distributed work anticipated in NASA's proposed human exploration effort.

References

Allen, J. F. (1983). Maintaining Knowledge about Temporal Intervals. *Communications of the Association for Computing Machinery, 26* (11), 832-843.

Bloomberg, J., Giacomi, J., Mosher, A.& Swenton-Wall, P. (1993) "Ethnographic Field Methods and their Relation to Design", in Schuler and Namoida (eds.) Participatory Design: Perspectives on System Design. Lawrence Erlbaum: Hillsdale, NJ, pp. 123-155.

Forsythe, D. E. (1999) "It's Just a Matter of Common Sense": Ethnography as Invisible Work, Computer Supported Cooperative Work 8: Kluwer: Netherlands, pp. 127-145.

Hayes-Roth, F., Waterman, D. & Lenat, D. (1984). *Building expert systems.* Reading, MA:: Addison-Wesley.

Friedland, P. & Kedes, L.H. (1985). Discovering the secrets of DNA. *Communications of the ACM, 28*, 1164-1186.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41-58). New York: Academic Press.

Jordan, B. (1996) "Ethnographic Workplace Studies and Computer Supported Cooperative Work", in D. Shapiro, M. Tauber and R. Traunmüller (eds.) The Design of Computer-Supported Cooperative Work and Groupware Systems (1996), North Holland/Elsevier Science: Amsterdam, pp. 17-42.

Nardi, B. (1996) "The Use of Ethnographic Methods in Design and Evaluation" In M.G. Helander, T. Landauer, and P. Prabhu (eds). Handbook of Human-Computer Interaction II. Amsterdam: Elsevier

Norman, D.A. (1991). Cognitive artifacts. In J. Carroll (Ed.), *Designing interaction* Cambridge: Cambridge University Press: Cambridge, pp. 17-38.

Sewell, D.R., & Geddes, N.D. (1990). A plan and goal based method for computer-human system design. *Human Computer Interaction: INTERACT 90* (pp. 283-288). New York: North Holland.

Shalin, V.L. (in press). The roles of humans and computers in distributed planning for dynamic domains. To appear in: *Cognition, Technology and Work.*

Shalin, V.L. & McCraw, P.M. (2003). Representations for distributed planning. In E. Hollnagel (ed.) *Handbook of cognitive task design. (pp. 701-725)* New Jersey: Lawrence Erlbaum Associates.

Shalin, V.L., Geddes, N.D., Bertram, D.L., Szczepkowski, M. & DuBois, D. (1997). Expertise in dynamic, physical task domains. In P. Feltovich, K. Ford, & R. Hoffman (eds). *Expertise in context: Human and Machine*. Menlo Park, CA:

Star, S.L. (1989). The structure of ill-structured solutions: Heterogeneous problem solving, boundary objects and distributed artificial intelligence. In M. Hahns and L. Gasser, (Eds.) *Distributed artificial intelligence, 2*, Morgan Kauffman: Menlo Park. pp. 37-54.

Wales, R (in preparation). Earth Space for Martian Work.

Wales, R., O'Neill,, J.& Mirmalek, Z. (2002). Ethnography, Customers and Negotiated Interactions at the Airport" In special issue on Human Centered Computing at NASA. IEEE Intelligent Systems Journal 17(5), pp.15-23.

Wales, R.C., Shalin, V.L. & Bass, D.S. (in review). Requesting distant robotic action: An ontology of work, naming and action identification for planning on the Mars Exploration Rover mission.