AN ONTOLOGY FOR REQUESTING DISTANT ROBOTIC ACTION:

A Case Study in Naming and Action Identification for Planning on the Mars Exploration Rover Mission

Roxana C. Wales
SAIC NASA Ames Research Center
rwales@mail.arc.nasa.gov

Valerie L. Shalin
Wright State University

Deborah S. Bass
NASA Jet Propulsion Laboratory

Abstract

This paper focuses on the development and use of the abbreviated names as well as an emergent ontology associated with making requests for action of a distant robotic rover during the 2003-2004 NASA Mars Exploration Rover (MER) mission, run by the Jet Propulsion Laboratory. The infancy of the domain of Martian tele-robotic science, in which specialists request work from a rover moving through the landscape, as well as the need to consider the interdisciplinary teams involved in the work required an empirical approach. The formulation of this ontology is grounded in human behavior and work practice.

The purpose of this paper is to identify general issues for an ontology of action (specifically for requests for action), while maintaining sensitivity to the users, tools and the work system within a specific technical domain. We found that this ontology of action must take into account a dynamic environment, changing in response to the movement of the rover, changes on the rover itself, as well as be responsive to the purposeful intent of the science requestors. Analysis of MER mission events demonstrates that the work practice and even robotic tool usage changes over time. Therefore, an ontology must adapt and represent both

incremental change and revolutionary change, and the ontology can never be more than a partial agreement on the conceptualizations involved. Although examined in a rather unique technical domain, the general issues pertain to the control of any complex, distributed work system as well as the archival record of its accomplishments.

1. INTRODUCTION

Everyday in work situations, people ask other people to do tasks that appear quite simple, yet, when analyzed, can be quite complex. A colleagues says: “Please make me four copies of this,” and with an impressive success rate, another colleague produces four copies of the appropriate document, because he or she understands the steps that are involved in copy making or will keep adjusting the work (filling the paper tray) until the task is accomplished. The success rate may change if the requests become more complex. Complications to identifying the work to be performed include asking for copies of multiple originals, giving different instructions for different originals, making multiple requests concerning the same original, demanding that some copies be express mailed by a deadline, and asking for copies from originals that are something other than a paper document (such as an electronic file). The fact that the originals are often accessible physical entities, capable of bearing a label and being relocated to the instrument that will perform the work also helps to simplify the problem of copy-making. But, the task of executing the request is invariably far more intricate than the apparently simple one sentence request would imply.

This paper focuses on the abbreviated names (rather than intricate instructions) that are associated with making requests. The work in our description involves multiple specialists collaboratively operating robotic vehicles for the scientific exploration of Mars during the 2003-2004 NASA Mars Exploration Rover (MER) mission, run by the Jet Propulsion Laboratory. Scientists and engineers needed semantics for higher-order actions that developed in an evolving system, and led to the subsequent evolution of an emerging ontology. The infancy of the domain and the interdisciplinary nature of the work required an empirical approach to the formulation of an ontology that was grounded in human behavior.

The purpose of this paper is to identify general issues for an ontology of action (specifically for requests for action), while maintaining sensitivity to the work system within a specific technical domain. Although examined in a rather unique technical domain, the general issues pertain to the control of any complex, distributed work system as well as the archival record of its accomplishments. Further, an ontology of action must address an issue that often can be ignored in an ontology of objects, such as might be found in an internet search application. An ontology of action must take into account a dynamic environment, changing in response to natural physical events as well as intentional actions. The ontology for expressing a request must reflect the influence of context on the meaning of action. For example, photocopying after an express mail deadline has passed may mean something quite different from photocopying before the deadline.

The next sections provide background information that is important to understanding the developing scenarios of the Case Study. First, we describe a description of the current MER mission work system that is the result of three years of design and lengthy training that was used to accomplish the tele-robotic exploration of Mars. This description is followed by a brief review of related literature spanning computer science, and behavioral and social science.

Mars Exploration Rover (MER) Mission: Current Work System for the Robotic Exploration of Mars

This section describes the MER process for developing *science activity plans*, which are requests for robotic action and are made up of higher level *observations* that contain subsets of *activities*. In this section are X subsections, covering preliminary planning, planning tools, science plan integration, rover coding, and the associated planning software in the ground data system. The name of an observation in the science activity plan provides coherence across all of these planning phases, which in turn influences the development of an ontology for action.

MER Planning Process

Each operational day, called a “sol” (a Martian day, which is approx 24 hrs and 39 min) the Science Team convenes in Science Theme Groups to discuss the newly arriving data and decides what to plan for the next sol using the 13 available instruments on the rover. Scientists can use instruments on three different parts of the rover to collect data, the body of the rover, the rover’s mast, and the robotic arm (instrument deploy device), as shown in Table 1.

Table 1.Instrument Locations on the Rover

Type of Instrument	**Name**	**Definition**	**Number**
Engineering instruments on Rover body	Navcam	Navigational cameras	2
	Hazcam	Hazard avoidance cameras	4
Mast Instruments for Remote Sensing	Pancam	High Resolution Panorama Cameras	2 (for stereo imaging)
	MiniTES	Mini Thermal Emission Spectrometer	1
Instrument Deployment Device (IDD) Instruments for Gathering In-situ Data	MI	Microscopic Imager camera	1
	APXS	Alpha Particle Xray Spectrometer	1
	MB	Mössbauer	1
	RAT	Rock Abrasion Tool	1

The five science theme groups are Atmospheric Sciences, Geology, Soils and Physical Properties, Geochemistry/Mineralogy, and Long Term Planning, the last a special group responsible for roughing out a general plan for the sol and ensuring continuity between the previous sol's observations and what the team wants to do with the rover over the next several sols.

All theme groups join in the first meeting of the sol called the Science Context Meeting. During this meeting, the Long Term Planning Lead roughs out the approximate plan that the rover will be expected to accomplish on the next sol based on incoming telemetry. The full science team discusses the plan, and the various Theme Groups suggest and receive assignments for Observation development. For example, the Soils and Physical properties group might develop an observation to examine the detailed morphology of a particular patch of soil. Alternatively, the Atmospheric group might create a request to monitor atmospheric dust.

Midpoint in the planning process, the full Science Team reconvenes at the Science Assessment Meeting to make adjustments to the type of sol being planned, based on a review of the available telemetry from the previous sol and past sols. Each of the Science Theme Groups continues to develop possible science observations as defined by their assignments. Groups and individuals rework observations considering:

- the type of sol being planned (e.g. traverse, approach towards a rock, in-situ instrument use).
- the available resources (e.g. power, data volume available for transmission, operating time)
- possible timing restrictions on when the observation can take place.
- relational events that will influence the observation (e.g. a communication event for data transmission)
- whether there are engineering restrictions imposed on the upcoming sol that will impact this observation.
- options for reducing the resources used by a particular observation (e.g. specifying adjustable parameter values)

Scientists use an emerging ontology to name the observation and convey appropriate information to science team members as well as to other teams in the downstream planning process. For example, the science team might generate an observation to examine changes in the amount of registered sunlight over the course of a Martian sol. To accomplish this, the Pancam cameras will image the sun using the solar filters at various times of day. Acquiring this measurement several times in a row allows the science team to find a trend in how much atmospheric dust loading has occurred. The Atmospheric Theme Group names this observation Pancam_Tau_Anytime. "Pancam" refers to the instrument; "Tau" refers to the method of data collection, and the analysis that will follow. "Anytime" indicates when the observation can be conducted.

The naming convention allows scientists to identify the different work of the instruments. The remote sensing instruments (those that gather information such as images and spectroscopy data from locations distant from the rover) might specify direction relative to the rover rather than a particular object upon which to do work. Post-drive_Navcam_360 is the name for acquiring an observation that requests a navigation camera panorama in a 360 degree circle around the rover at the end of a drive day. An in-situ measurement (data collected by direct or close contact with in place rover instruments) might be named Post-MB_MI_5position_ElCapitan, implying a request to use the Microscopic Imager camera to acquire five pictures on the rock El Capitan after the Mössbauer instrument has completed its measurement.

MER Tools for Creating Science Observations

The planning process involves more than the specification of an observation name. The mission team is responsible for translating their observations into a language for programming the rover. To facilitate the translation of purposeful action into rover language, the science team uses a software tool called the Science Activity Planner (SAP) (Figure 1. and 2.) The SAP software permits the definition of Features and Targets, which serve as the focal points of observations. Features (such as rocks and craters) have unique names that do not change over time due to their prominence in the landscape. Targets specify a particular location or spot on an already-named feature that are the focal point for the work of instruments. Target names might have a whole-part relationship with features. A feature might have one or several targets. For example, a feature Dolphin has a target Fin, and the feature Tamarend Park has targets: Center, Park diamond, Sandbox, Jungle gym. However, the rover does not know about features and targets in its environment. The rover knows how to extend its robotic arm, but not whether it is extending the arm to touch soil or a rock.

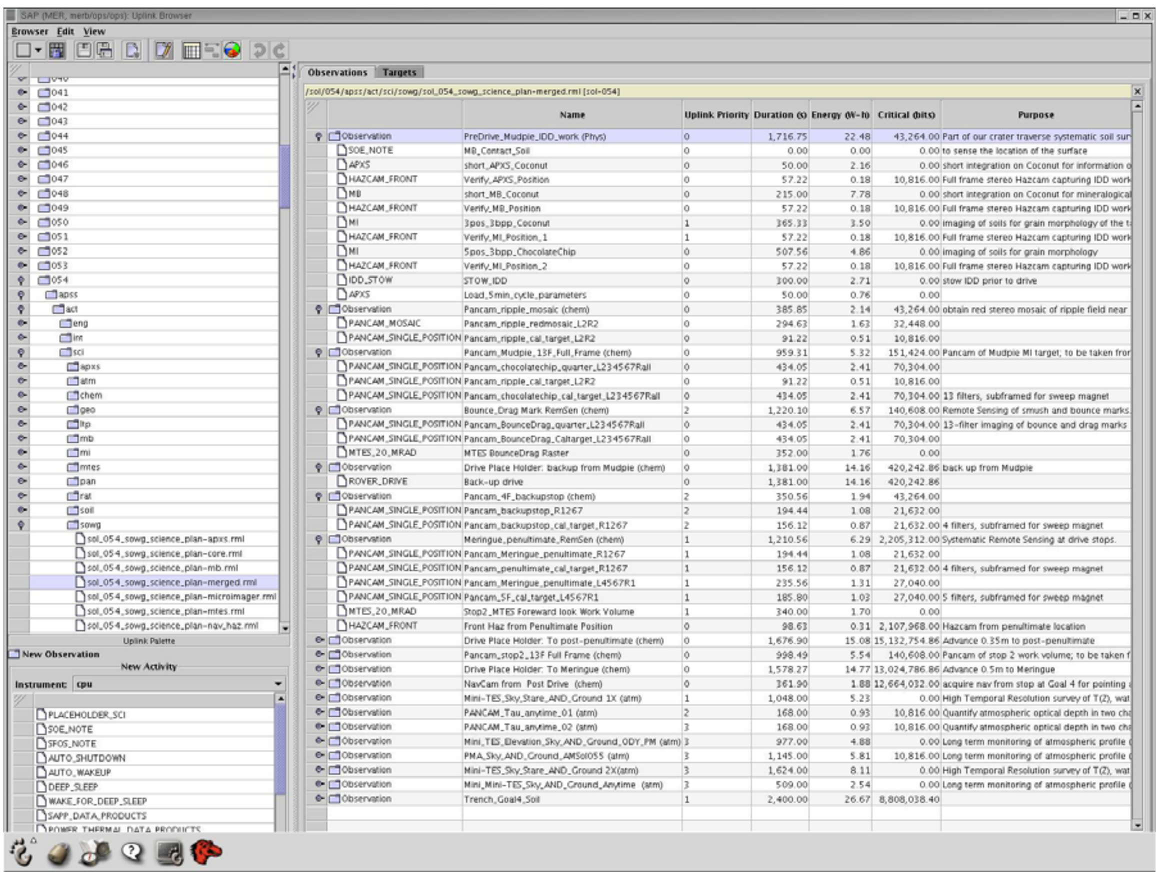

Figure 1. Science Activity Planner (SAP) Tool: screen shot shows a science activity plan as created by the science team with the higher order observations.

Open toggles on some observations show subordinate activities that instantiate the observation**.**

Observation	PreDrive_Mudpie_IDD_work (Phys)	0
SOE_NOTE	MB_Contact_Soil	0
APXS	short_APXS_Coconut	0
HAZCAM_FRONT	Verify_APXS_Position	0
MB	short_MB_Coconut	0
HAZCAM_FRONT	Verify_MB_Position	0
MI	3pos_3bpp_Coconut	1
HAZCAM_FRONT	Verify_MI_Position_1	1
MI	5pos_3bpp_ChocolateChip	0
HAZCAM_FRONT	Verify_MI_Position_2	0
IDD_STOW	STOW_IDD	0
APXS	Load_5min_cycle_parameters	0
Observation	Pancam_ripple_mosaic (chem)	0
PANCAM_MOSAIC	Pancam_ripple_redmosaic_L2R2	0
PANCAM_SINGLE_POSITION	Pancam_ripple_cal_target_L2R2	0
Observation	Pancam_Mudpie_13F_Full_Frame (chem)	0
PANCAM_SINGLE_POSITION	Pancam_chocolatechip_quarter_L234567Rall	0

Figure 2. A partial image of observations and activities in SAP.

Mission engineers use a language to communicate with the rover (command language) that is entirely a function of the rover's own internal states. Humans must construct the relationship between the behavior/actions of the rover, the rover's external, environmental context, and the scientific purposes related to any request. In other words, the rover understands how to turn its own devices on and off, and it knows its own pitch, roll and yaw, but it does not understand its state in relation to the environment nor the scientific experiments it is being asked to execute. As a result the command language of rover action is not explicitly grounded in context and purpose. However, an ontology for human collaboration must relate to and convey information about both the internal world of rover work (instrument use, calibrations) and the external world of rover environment (timing of work feature and target).

Within SAP, scientists identify an observation consisting of a set of actions the rover would accomplish for a specific purpose. The individual actions are called Activities, and a single Observation would be composed of one to many Activities. For example, an observation request might be made for Pre-RAT_MI_Golf. This requests asks that the rover collect a series of microscopic imager (MI) data on the feature, Golf, before executing any commands involving the Rock Abrasion Tool (RAT). Activities are specific building blocks that translate into computer commands the spacecraft will understand. Activities are defined in the Activity Dictionary, which specifies a standard set of parameters with a range of values that must be identified for each instantiation.

MER SOWG: Science Plan Integration, Planning, and Prioritization

After defining their observations in SAP, members of the science team meet again in the Science Operations Working Group (SOWG) meeting to discuss and finalize the science activity plan, which is the complete set of requested observations and rover actions for the next sol. At the beginning of the meeting a representative from the Spacecraft/Rover Engineering Team summarizes the data received from the rover, the resources available for that sol and the necessary engineering activities (e.g., housekeeping or solar panel imaging) that will consume some of these resources. The SOWG group then considers each science observation and its activities as they produce an integrated "Science Activity Plan" appropriate to the available rover resources. They must consider temporal constraints such as whether a certain observation must be completed prior to the execution of a second observation, or whether a given observation might have time of day constraints for temperatures or lighting. The team prioritizes the observations and ensures that the plan achieves the objectives for the day. During this meeting, the observation name improves the efficiency of the work because it highlights important identifiers (such as instrument, method, and constraints) that are necessary to decision making, planning and scheduling. Finally the team identifies a rough planning timeline, and checks the plan against a model that predicts the resources that will be consumed.

MER Sequencing and Ground Data System

After the Science team completes its task, the engineering team translates the observations into the computer code the spacecraft can execute. As a first step, the engineering team works with a spatial representation to schedule the science team Observations in relation to fixed events in the day, such as rover wake up time and opportunities for communication between Mars and Earth. Observation names therefore need to fit within this display and the resolution of time that it provides (Figures 3. and 4.) This scheduling process involves a somewhat higher fidelity resource model, resulting in the removal of low priority observations when they don't fit in the available time and power envelope. Unambiguously and quickly referring to Observations by name is a necessary function of this portion of the process. In fact, the software display that is used for this step shows only the Observation and Activity names, along with colored blocks that represent priorities and durations of Observations and Activities.

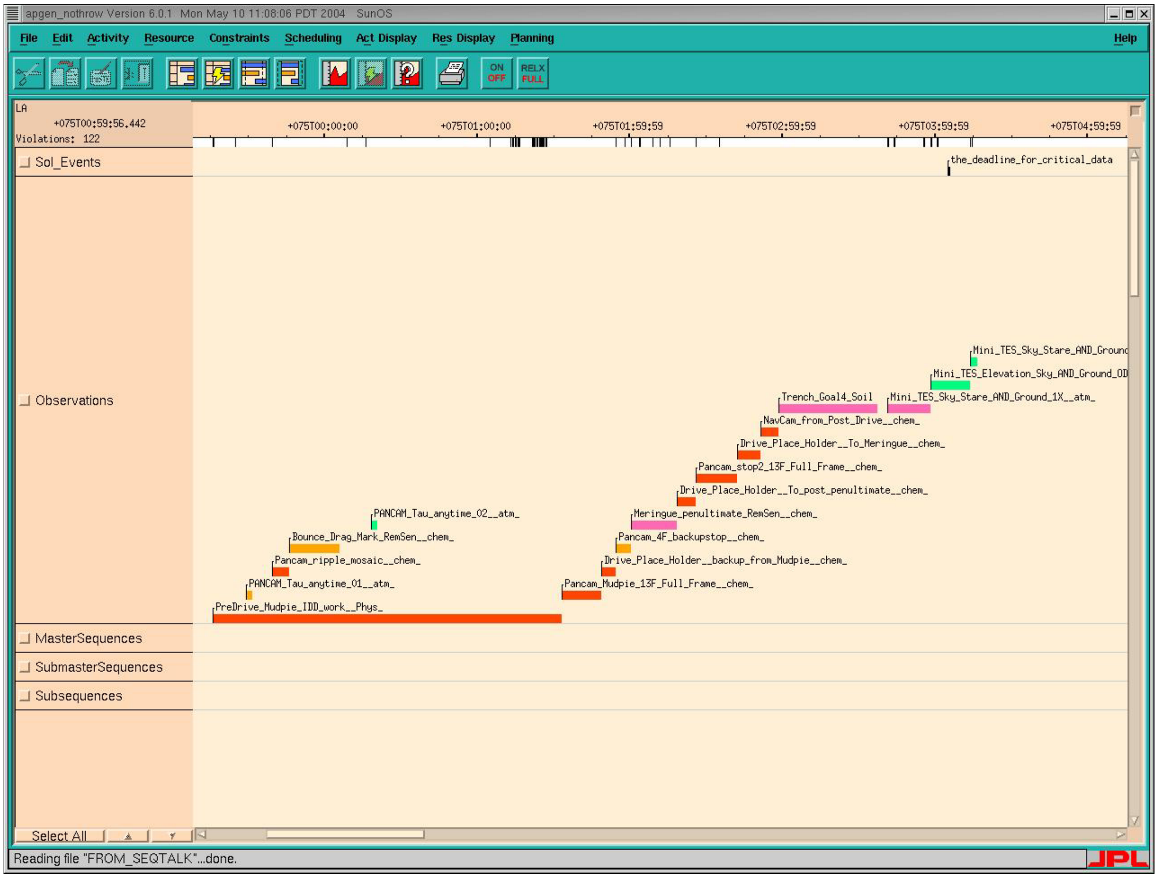

Figure 3. Activity Plan Generator (APGEN): The screen shot of a version of the interface that shows observations as they are "planned" into a timeline for execution. Colors indicate the priority level that scientists have assigned to each observation.

Figure 4. Magnified view of APGEN.

After generating the scheduled plan, the team works to create “sequences”, which are sets of instructions in computer code that the rover can understand. Because the rover only understands its own states, humans need to work as translators to tell the rover what to do and precisely how to do it. A specialized group of scientists works on this translation process, and their involvement has influenced the ontology that has developed. This group of scientist-engineers, called Payload Element Leads, has a very firm grasp on the need to identify rover states and has created language to convey that information. The scientists, as domain experts, have developed and used the language, and it has become increasingly more sophisticated and relevant to their work. For example, they create shorthand names for new methods they have developed and the names become part of the vernacular.

Each set of instructions uplinked to the rover and successfully executed results in sets of data that requires a filename. Engineers, who need to monitor the receipt of data in order to manage on board memory, require unique, representative names. They favor numeric identifiers, which are unique, easy to increment, and easy to manipulate without list processing. Scientists, who need a more human readable and historical record of the observation that leads to the data, favor meaningful and therefore relatively unconstrained names. The science team’s desire for fewer constraints in naming needs to be balanced with the downstream-need for uniqueness and regularity. The science planning software includes large, open, write-in fields. The lack of constraints around the field leaves it up to the scientists to be salient and discriminating to meet the demands of the timeline. As well they need to be concise so information can be more easily honed as it progresses through the system. The names must balance between salience and the need to communicate and convey important information to team members for both discussion and correct execution.

MER Summary of Work System Implications for Naming Observation

Individual and small teams of scientists use observation names in developing their requests. The Science Operations Working Group uses observation names to discuss and develop a plan for the sol, including the merging of similar types of observations. The engineering team use observation names for scheduling, and to trim the schedule in order to meet available resources. Engineers and scientists are concerned with the filenames of data returned after the rover executes observations. The naming problem concerns not only the rover capability but the need for extended planning that requires multidisciplinary communication among human participants. The current work system and Observation names described above evolved over a period of three years—perhaps a relatively long time for an engineering effort, but a rather short time for the evolution of a domain of work. . The case study describes the initial naming attempts, and how the work system evolved and our role in this process. Before turning to this case study, the following section addresses the literature relevant to the naming of observations for this system.

II. Related Literature

The literature related to the problem of how to name requested observations spans computer science (ontology and planning and scheduling), and social and behavioral science. While many of the issues identified in this section were apparent to us at the outset, some were not. Our own awareness of these issues however, was not the decisive influence on the circuitous pathway we encountered in developing observation names. The ultimate influence was the practical need for a naming convention that the scientists recognized as they trained and developed their own tele-robotic expertise which fortunately converges nicely with the following issues.

Computer Science Literature: Ontology

The current work on ontologies reflects burgeoning interest in developing languages and applications to enhance internet search and use. Contemporary researchers acknowledge the need for a explicit specification of conceptualizations across participants in any given domain (Gruber, 1993; Guarino, 1997) thus allowing for search, navigation, inferencing, data integration, pattern completion, and configuration of physical devices (McGuinness, 2001). Ontologies are being developed in domains as diverse as on-line commerce, industry, intelligence analysis and medicine. Key concerns include the criteria for ontologies (e.g., unambiguous interpretation of classes and relationships, (McGuiness, 2001). Focus is on guidance and development of the formalization process and syntax, and relations to the known functional requirements of the target application. However, the emphasis appears to be on optimizing the semantic properties and relationships between constructs within a domain ontology. This inwardly contained focus is generally described in the philosophy of mind literature as intensional semantics.

The ontology literature does recognize that the development of an ontology will require it to evolve (McGuiness, 2001) through an iterative process, with potential for growth and change. Gruber’s definition of ontology in particular acknowledges that the agreement among domain experts may be incomplete (Guarino et al., 1995). Researchers provide advice on the sort of development environments that assist in the collaborative representation of an ontology distributed across various experts (Noy, et al.) and especially evaluation criteria for a developing ontology. However, for the domains in question, contemporary researchers are not so much developing an ontology as they are extracting and representing an ontology from an established domain of practice. An established domain of practice already reflects the important constructs and distinctions, so that as constructs are added to the representation they are rather unlikely to upend the representation in progress. While the task of representing and formalizing an ontology from an established domain of practice is not trivial, the present emergent domain required another order of ontology development. The geologists and engineers as well as the authors of the present paper had to understand the evolving work through a series of tests and training exercises, with the authors reflecting on and feeding back emerging constructs as they developed. This was a matter of practical necessity, due to the demands of the mission timeline and the landings on Mars. Other domains may have an on-going, decades-long evolution of practice.

Contemporary interest in ontologies stands on the shoulders of a longstanding interest in computer models of mind, which also require consideration of the relationship between natural language, its underlying semantics and ontologies (e.g., Schank, 1972; Schank et al., 1973). Critics of the computational models of mind specifically noted that natural language is contextualized (Dreyfus, 1979) and intentional (Searle, 2002), casting doubt on the then prevailing enthusiasm for formalizing natural language as unambiguous, stand-alone constructs for computational manipulation. While contemporary ontologists may have circumvented this concern by working on applications involving search and pattern matching, the role of context and intention is critical for commanding robots that effect action in a dynamic world.

Computer Science Literature: Planning

The planning literature addresses a different sort of task than much of the work on ontology has addressed. A key characteristic is that actions are executed in the world, so that extensional semantics are as important as intensional semantics, which focuses on internal consistency and formality. Actions change the state of the world, so that different sequences of identically parameterized action (target, field of view, elevation angle of the instrument) can have different meanings. For example, moving the rover and then taking a picture focused on a particular target results in a different image than taking a picture before driving. The mere passage of time also results in changes to the state of the world (lighting), so that taking a picture of a particular target at one time of day does not have the same outcome as an identically parameterized image taken at a different time of day.

The computer science literature on planning sometimes assumes that the semantics of action can be adequately expressed as a state change (Georgeff et al, 1986). A rover drive is nothing more than location at one point in time and re-location in another. When action is executed in a real physical environment, the manner in which the action is executed becomes important. The drive description might be enriched to specify the intermediate waypoints and orientations of the vehicle relative to the terrain. Accordingly, objects (the vehicle, the terrain), attributes (e.g., vehicle configuration) and relations (e.g., direction, distance) define the ontology of action. However, single terms might describe a drive method at a different level of analysis, e.g., as a small bump or wiggle.

The early literature on linear planning (Fikes et al., 1971) assumed that the planning of action is suitably described at a single level of analysis. Accordingly, a proper plan consists of a set of actions at one level of analysis arranged in an order that achieves a particular goal. Sacerdoti (1977) demonstrated the limitations of this approach and the need to conceptualize a planning task at multiple levels of abstractions. Consistent with Sacerdoti's insights, the technique of skeletal planning (Friedland et al., 1982; Stefik, 1981) represents possible actions in a hierarchical graph, with the most abstract descriptions at the top and the most specific descriptions at the bottom. The representation is a graph rather than a tree because the very same low level action may play different roles in different plans. For example, we might take a stereo picture of terrain to study an object's texture or to provide distance estimates for navigation. Actions that seem identical at one level of analysis may appear quite different at another. For example, we can take a picture of a rock with NavCam and Hazcam to aid rover navigation. In detailed rover language, the two sets of activities needed to accomplish this request have very little in common, other than the fact that they have the same purpose. These examples illustrate a distinction between the purpose of action, and the method of action (Sewell & Geddes, 1990). For example, in the situation where scientists request moving the rover and then taking a picture, which is focused on a particular target, results in a different image than taking a picture before moving. This illustrates the context sensitivity of the semantics of action and reinforces the distinction between method and purpose; the very same method of picture-taking executed in two different contexts (lighting or location) can achieve two different purposes.

Behavioral & Social Science Literature

On MER scientists used observation names, and some of the general characteristics of cognition and communication had potential influence over their choices of the names. The cognitive science community began to explicitly address the ontology of thought in the late ‘70’s (Gentner & Stevens 1981; Rosch & Lloyd 1978) although the interest persists (Lakoff, 1994). Much of the cognitive science literature concerns adult categorization of objects. An early practical result is that different adults do not consistently use the same language for the same concept (Furnas, 1983). Contemporary researchers would likely attribute this inconsistency to the nuanced meaning that natural language provides. Some researchers investigated the ontologies of children in a developmental or instructional context and documented that the organizing constructs change with experience (Greeno, 1983; McCloskey et al, 1980).

Far less work addresses verbs. The behavioral science literature makes a distinction between (at least) three types of verbs. It suggests that people pay attention to different aspects of verbs. So for example, the MER scientists might focus on action as in “driving” or they might focus on instrument. verbs, such as “RATting”, or they might focus on the resulting state change verbs such as “trenching”, which results in a changed state in the soil or the trench. The literature on adult speech preferences is rife with contradiction, but some researchers suggest that adults prefer instrument verbs because of their efficiency in communicating substantial detail (Behrend, 1995).

Social science research places verb choice in the context of conversation. Well-accepted Gricean (Grice, 1975) principles of communication will influence the level of abstraction and the level of detail utilized in an exchange. Speakers attempt to provide information that is unambiguous but without overwhelming the listener with detail that is irrelevant, either because the detail has no bearing on the purpose of the exchange, or because the listener is already aware of the detail. Context can render detail irrelevant. For example, if there is only one drive, it may not seem necessary in conversation to specify the details of this drive in its name. Or, if participants are examining an image, they may refer to a feature with a pronoun, and point to the image to disambiguate the pronoun. However, a late arriving participant in the exchange may not be privy to the pointing. This is a serious consideration in distributed work. As context changes, what was once obvious (and hence assumed) becomes ambiguous. The communication literature implies that adjusting for the changes that may take place in a future context conflicts with the tendency of humans to simplify in response to an existing context.

III. The Research Method

We used ethnographic methods to both capture and shape the ontology for action. Ethnography has a track record for successful data collection in corporate and organizational settings and is particularly well suited to understanding complex settings. Ethnographic methods provide a number of data collection techniques that 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 us 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 others of us 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) Participant observation requires researchers to be present during on-going events, making field notes of in-situ observations, asking informal real-time questions pertinent to those events and conducting more formal interviews as appropriate. As participant observers we were present during all of the pre-mission tests and all but one of the science team twice-yearly meetings. We listened in on the majority of the science team weekly conference calls. On occasion, we presented information on the developing convention to the team during those meetings. As badged members of the mission team, we were able to move with flexibility within and across the domain, attending meetings and working with software designers. Over time we came to know the science team members and many of the engineers, interacted informally with them, got feedback and worked with them to develop the ontology. We helped train the team in the use of the ontologic convention during the science team training "flight schools."

To understand work practice and the characteristics of the developing ontology, we analyzed field notes and mission design documentation. We made video recordings with hand-held and fixed-mounted cameras and analyzed the data using "interaction analysis" (Jordan, 1995). We identified and categorized information, and analyzed communication exchanges, scientists' work practice and their scientific reasoning. We learned many of the intricacies of the rover instruments. We also assessed the software requirements and interfaces between mission technologies.

We had innumerable teleconferences and email exchanges in which we analyzed and assessed the data and brainstormed new ideas for a convention that would support the work of the science team.

After each test, we analyzed the data in the science planning tool (SAP) as well as from field notes to develop an understanding of the cognitive, linguistic, referential and software needs relative to a naming convention and for the development of an informal ontology. We focused on the learning of the science team and the emergent quality of their work practice and their needs for a naming convention. As the progression of our recommendations indicates, we could not have settled on an adequate ontology at the outset of the project because no one knew what issues would emerge. Our participation enabled us to capture the developing science practice and translate it into a naming convention.

During the mission, we continued our data collection, taking field notes, making video tapes of meetings and collecting copies of mission science activity planning print outs. For this paper, we have done a qualitative analysis of those data.

IV. The Case Study

For almost three years before the landing of the MER rovers on Mars in January of 2004, mission scientists and engineers participated in a series of tests to develop the science planning processes for the mission. In this case study we describe two field tests in 2001 and 2002, pre-mission tests conducted in 2003, and the work practice of the landed mission that began in January of 2004. Our naming convention developed from initial understanding of the problem in the first tests, through the development of a taxonomy of the categories of identifiers that were pertinent to the naming convention, to an emergent ontology with syntax, semantics and a description of the relationships between the categories of identifiers over this period.

We describe in some detail, the first field tests with the Field Integrated Design and Operations (FIDO) rover in 2001 (and precursor “Mars Yard” tests), because the work we observed during this period highlighted the need for a naming convention to support scientific reasoning and a complex collaborative work practice. We also realized that the convention would function as both an identifier of and carrier of information within the ground software. We needed a convention that could work at a natural language level during the collaborative process and then transition into more precise identifiers in the software as the science requests moved from science to engineering teams for eventual translation into commands for the rover.

From our observation and participation in these tests, we first detected and then developed an understanding of many issues relative to referencing, identification and communication that would affect the ability of science team members to efficiently request and share information. We determined that the "names" of science requests would affect communication within and across teams, the collaborative decision processes, the flow of information to engineering teams, and the management of returned data products from the rover. With each successive test, working with the science team and other mission members, we made a series of recommendations that led to the development of a naming convention and influenced the work process described above. The team was learning to do this new, remote Martian science using new geological and robotic tools. Our research, congruent with the science team's research on Mars, was focused on the language of this emergent work as well as on how best to identify or "name" science requests.

In later sections of this paper, we show how the recommendations we made influenced both the work practice of the MER mission and software design.

2001 Mars Yard Tests

In the spring of 2001, MER participants operated the FIDO rover in their first two tele-robotic training tests. The science team worked in a windowless room at the Jet Propulsion Laboratory. During the tests, the rover was located in a nearby sand and rock filled area called the "Mars Yard." The two Mars yard tests lasted two days each.

During the Mars Yard tests as well as the later FIDO 2001 test, the scientists analyzed the images and other data as they were returned from the FIDO rover and used those data to plan new rover movement, drives and the use of the on-board instruments. Despite some differences in the FIDO and MER rovers, such as a different computer language for rover commanding, from the scientist's perspective the FIDO engineering process was a reasonable approximation of the tele-robotic process. The participating engineers commanded the rover and processed the retuned data for the science team. The FIDO test was designed so that meetings, science decision processes and planning events were as realistic a rehearsal as possible for the later Mars work. The tests also included rough simulations of the time-delayed communications between Earth and Mars.

2001 Mars Yard Tests: Work Practice and Naming

"How do we name things?"- Participating scientist on the first day of the first Mars Yard test.

When the Mars Yard tests began, scientists openly acknowledged that they did not have a naming schema and decided to see what worked as they went along.

According to the target-focused attempt, scientists just started naming "targets" by indicating a point of scientific interest on an image in the science planning software. Initially, each group chose a different naming convention for targets. One theme group used letters of the alphabet and added numbers indicating a single target in a number of chosen targets to get names such as "G3." Another group chose flowers, and another just used numbers as in "rock 1." The mineralogy group began by naming rocks according to their appearance ("Geometric" "White"), but a member of the engineering team suggested they call them by letters of the Greek alphabet ("Alpha," "Beta," "Gamma"). The scientists named according to the character or quality of the objects in the terrain, while the engineers named in a way that carried an implicit order between names, guaranteeing uniqueness but reducing meaning.

When the scientists gathered as a group for their first science assessment meeting, they realized the variety of naming schemas the different groups were using. In real time during that meeting, they changed all the existing target names in the software, giving them new, Greek letter names. Almost immediately confusion resulted. Group leads had to remember old and new names in order to make sure they were talking about the right rock in relation to the work that their group wanted to do. “Wait, which rock is Gamma?” “Was that G3?”On the second day, the Greek letter convention started to break down, sometimes a target was named with a Greek letter and sometimes it was identified with a theme group name and a number. The confusions continued into the next team meeting the Science Operations Working Group meeting, in which the team chose their final science activity requests to be commanded for work by the rover.

At the end of the Mars Yard test period, after experiencing several instances of confusion over name choices, the group decided on a naming convention that they would use during the upcoming FIDO 2001 test. The convention called for the following identifiers to be used when scientists first named targets in the theme groups:

Sol #_ science theme group shorthand_target #

An example would be:

sol 4geo5

indicating the fifth target chosen by the geology group in sol 4. To prevent confusion if two groups chose the same target in the theme group meetings, the new procedure called for renaming targets during the SOWG meeting, the last process meeting before sequencing, with an official, unique “SOWG target name”. The SOWG chair would create a list of names to be available for designating the official SOWG target name.

2001 FIDO Rover Test

During FIDO 2001 test, the science team still operated inside a windowless room, but the rover was in an un-disclosed location in California's Mojave Desert. The secrecy of the desert location forced scientists to use only test-derived field data, as they would have to do on Mars, to identify and categorize the terrain and locate the rover on a map. "Cheating" by using aerial photographic data external to the test was discouraged

The FIDO 2001 field test lasted twelve days. In order to get in as many planning cycles as possible, the tests incorporated two, single-sol planning cycles into one Earth day (one sol planned in the morning, one in the afternoon of each test day). The later Mars mission had a single-sol planning cycle per twenty four hour period (approx.). Because of the compression of events during the tests, scientists were able to rely on memory to help them do much of their work and planning, something that would be much more difficult over the longer duration and larger scale of the MER mission. We saw immediately that any request naming confusions in the early tests would almost assuredly be magnified during the MER mission.

Although the naming process became more consistent with the implementation of the new naming convention, there was still confusion The official SOWG target names, chosen by the SOWG chair from the designated list, were most often single words and represented a theme such as baseball players, kinds of guitars, cars and race horses (Ex: [Hank] Aaron, [Mickey] Mantle, [Babe] Ruth, Guild, Corvette or Citation). After a science request was given an official SOWG target name, that name became the most consistent science request identifier: "So then, we will do Guild?" "Where are the results of the MiniTES on Aaron?"

The procedure for renaming targets from the theme group names, such as Sol7geo6 to an official SOWG name was informal one and varied over time.

Targets got their official name when they were accepted as part of the final science activity plan during the SOWG meeting. The SOWG chair verbally assigned the new name and one person entered that name into the software while everyone else looked on. They sometimes renamed targets individually as they went along and sometimes all at once. Participants, including those who were making the official name change in the software, had a hard time keeping up when several changes were made simultaneously, leaving room for mistakes in the electronic record.

Scientists kept individual, precise notes or spread sheets to track these name changes. The original name, such as (sol 7geo6), was over-written and disappeared from the software unless a cumbersome save process was followed within each theme group. Finally, engineers added numbers for rover-centric site and position identifiers as well as the sol number. The site and position numbers helped the engineers track the assumed context for a request, but had little meaning for the science team.

The teams used a variety of the above names in face to face conversation. Theme group members, both in small and large groups, referred to original target names such as “sol 7 geo6” or just “geo 6” and later used the official target name, such as Citation. Interestingly, all test participants often used instrument as a target identifier in conversation, both during the planning process and in the downstream engineering team, for example they would say, “Do we want the Pancam on Citation?” during face to face interactions, but instrument was not entered into the software. The target name was simply Citation.

2001 FIDO Rover Test: Analysis of Field Test Process

In addition to the obvious problem of changing target names during a sol and the unfortunate use of sol and site identifiers that had limited meaning for science team members, initial analysis of the 2001 field test processes suggested several work practice issues related to naming that created problems and confusion for participants. These issues made it difficult to trace the history of work, to refer to work with clarity in conversation and to capture necessary information to support group understanding. In some cases, the problems contributed to the loss of data.

Team members assigned new names to previously designated targets on each new sol, even when they were repeating work on a target/spot where they had already done work.

During the tests there was no consistent procedure or convention that allowed a target to keep its official name from sol to sol. Targets were given new official names when a new request was specified. However, this created confusion in referencing as team members referred to a target with either the old or new name, sometimes meaning the new work but using the old name and vice versa. Toward the end of the test, when members saw the confusion that could result from re-naming a target, they tried keeping the official name of a target from a previous sol, such as Aaron but asking for new work to be done on that target spot. In these cases, however, it became difficult to differentiate between the different instantiations of the data. Was the data from Aaron on sol 13 or sol 15? Misunderstandings multiplied when scientists targeted new, related work on a spot that was close to an already existing target. They gave the target a new name but there was no convention for identifying relationship either between the new and old targets or between the resulting sets of returned data.

A target "name" came to represent more than the targeted point where scientific work had been done.

In their work, participants used the single "target" name to identify several objects: the object (rock or soil patch) in the terrain, the scientific instrument used on that object; and the physical target point or placement of the instrument. Additionally, that target name became the referent for all related information in the software and the entirety of information relative to a science request during discussions in later meetings and with other teams. Because discussions and information were often about parts of the information, which were not explicit, miscommunications and misunderstandings could result.

A rock's target name became part of the vernacular. Team members used the name as a way to refer to developing understandings and make comparisons to earlier findings. The target name, "Ruth", for example became the team's referent for a type of exposed formation known as caliche. Scientists also made statements such as "This is more consistent with a "Mantle" observation" in which "Mantle" was the name of the target. By simply stating the target name, they could refer to the results of previous instrument work as well as a developing body of scientific knowledge. Renaming targets, therefore, blocked scientists' ability to follow a developing trail of scientific understanding and to make comparisons. For example, if Ruth was renamed Gibson for work done on the same caliche formation on later sols, there was now no single referent to the work. However, the practice of keeping a target name across sols had its own problems. It was hard to differentiate between the different sets of data or instrument measurements on the same target spot if they were all named "Guild."

FIDO 2001: Implications for Developing a Naming Convention

In later analysis, we further de-constructed the work of the science team and their use of "target" names and identified several principles important to a naming convention and basic to supporting their work.

- *Names have to contain more information. They represent **complex science team requests**. Single target names are not capable of conveying and carrying that complexity.* A naming convention needs to make all important information explicit to be shared across team members and over time.

- *A naming convention must identify both the **objects** that are located in the terrain **and** the **activity** that is being requested on those objects.* In retrospect, the need for differentiation between object and activity may seem obvious, but initially it was difficult to parse and understand this blurring and the underlying assumptions that caused it. One of the issues involved, we believe, is the cognitive/linguistic meaning that is implicit within a word such as “target,” especially when it is being used in a face to face setting. Target implies that *action* will be done on a point, or that the target point is the *goal of some action*. However, without making that action explicit in the name as well as traceable in the software, the associated action becomes lost over time both in the software and to members who were not present. The target name alone does not make the action explicit.

- *Once named, Target/Objects must keep that name to minimize confusion when future or related work is done on the same object. However, the name must identify the different instantiations of such work.* Consistent names support developing understanding and make reference to objects useful and meaningful. Necessary name changes must be carefully documented and made available to participants.

- *Target/Objects are types of things (soil or rock for example) and they are nested, sitting on or within larger objects. Larger objects sit within areas on the terrain. These relationships must be made explicit.* A target is point of interest for collecting data with an instrument. It sits on a larger object, a rock, for example. A single rock may contain several points of interest or “targets”.

Drawing on the above principles, we defined categories that were important and relevant for identification and reference in doing tele-robotic science with a rover. Initially, we identified two categories, each with sub categories. Working with the science team and mission members over time, we defined these categories more closely and made some statements about them.

- *The objects in the terrain*

  Features: an object of scientific interest or reference in the terrain, such as a rock, cliff face, crater, hill, soil patch.

  Target: A specific location on a feature that is a point where an instrument will be pointed or placed in order to gather data.

- *The actions or activities of the instruments on the rover*.

  Activity: Actions of the instruments and rover stated as groups of commands that the rover understands

  Observation: A set of actions that are done in relationship in order to gather data for a specific, common scientific purpose. The concept of observation came from a past mission, and serves to unify the set of actions.

2002 FIDO Rover Test

In preparation for the second major field test, FIDO 2002, when scientists were again located at JPL, but this time working with a rover placed remotely in the Arizona desert, we continued to hone the relationships between the object and action categories and sub categories. We created the following taxonomy and began to establish relationships between the parts of the taxonomy.

Observations and Activities

- Observations in the software are the high level goals that are a container for any activity or set of activities required of the rover. An observation must contain at least one activity to be defined as an observation.
- Activities are identified by type in an activity dictionary. For example, one type of Pancam activity is a single-frame. Another type is a mosaic.
- Activities are parts of observations.
- Observations must have a feature as well as a target as the location for activity.

- Activities have to be associated with a higher level category of observation in the software.
- Rover instruments, when used in a name, indicate the request for activity or action of a certain type. They are important identifiers. Additionally, instruments help identify the desired state of the rover when that activity will be done. For example, the rover arm must be extended when using arm instruments to do in-situ work on a rock or soil in front of the rover. Using instruments on the rover's mast indicates remote sensing or data collection that is distant from the rover, as in taking a picture with the Panorama camera.

Features and Targets

- Features would be identified in the planning software. Features do not have to be associated with other objects or actions, such as activity or target, in the software. They are objects of reference and interest in the terrain that can remain independent in the naming convention.
- Targets are parts of features.
- Targets have to be associated with the higher level category of feature in the software.
- Activities require a target. Targets indicate the point at which an activity will take place in order to collect data.

Based on our continuing analysis, we suggested a convention for the next field test and made recommendations for additional functionality in the ground software. It allowed participants to identify features (unlike FIDO 2001), associate targets with features (providing the ability to differentiate between features and targets), and create observations. The FIDO 2001 software had only allowed scientists to indicate targets and to choose from a set of rover activities.

Our FIDO 2002 naming recommendations stated:

- Observation names should identify the instrument as well as the feature name and a target.

- The observation name should have a consistent syntax. Instrument should be identified first as the most consistent reference, then feature and then target. Instrument_Feature_Target. Example: Pancam_Shoe_Heel, identifying first the Panorama camera instrument, pointing at the feature Shoe, with the center image point on the target Heel.
- Observation names should be descriptive of relationships between the work of the instruments and features and targets if possible. Observation names could indicate basic relationships between objects and actions, such as which instrument was used on which feature. Observation names could also point to remote objects and point to more than one object. For the FIDO 2002, we recommended *Survey* as an appropriate concept for indicating such relationships.
- Target names should identify, if possible, a whole-part relationship with the feature such that the Feature represents the whole, and the Target represents the part. Example: Feature=Shoe; Target=Heel. Whole-part names establish relationship between target points as well as identify the relationship between targets and a feature and facilitate face to face discussion.

2002 FIDO Rover Test: Work Practice and Naming

As we observed and analyzed the work of the science team in the second test, we understood that our recommendations had changed work practice, regarding the function of environmental features, reference to rover instrumentation and the development of systematic methods for tele-robotic exploration.

Features

Once the scientists could identify features as well as targets in the software, they began to name and track distant features that did not have targets for specific work. The features became reference points, established context and aided in indicating longer range goals. As they would on Mars, where small objects in the landscape would not have names, scientists began naming distant buttes and rocks in the desert. The names became shared references to the landscape, direction and terrain. Scientists also identified closer rocks and soil patches as features of interest. These features became markers for planning the work and traverse of the rover. The team also printed images, marked feature names on them and hung them on the walls as points of reference and context. Now that features and targets had separate names, scientists could also refer to the entire body of data and findings on a single feature as well as the data that might be associated with one or more particular target points

Instrument Names

Scientists indicated their willingness to include the instrument in the observation name. However, the compressed timeline limited their ability to fill names in completely in the software. When participants did fill in instrument names, it improved team understanding and enhanced the understanding of downstream teams. They did get more proficient at filling in the fields over the period of the tests. Example: On sol 8 a name was “sky spectral spot”. On sol 11, the name was expanded to include an instrument identifier (IPS, an instrument on the FIDO rover) and then abbreviated removing the word spectral, as in Sky_survey_IPS_2ELl_ sol 11. This request was for a spectral survey of 2 distinct elevations in the sky on sol 11. In face to face situations, using natural language, the team continued to use the instrument name as they had done in FIDO 2001. Based on this data, and our analysis that showed the instrument as central to activity description, we decided that instrument name was an important part of the naming convention and kept it for later iterations.

Method

Based on analysis of the tests data we decided that method was also an important part of a naming convention, because it allowed for efficiency in the naming process. While we had suggested Survey as a method for identifying camera pointing and data collection, we anticipated that scientists would begin to

use and develop other methods as they became more proficient at doing tele-robotic science. We expected to see methods that combined the actions of certain sets of instruments for repeated and consistent data collection strategies. We also expected to see methods that would standardize individual instrument activities and parameter settings making them consistently re-useable.
". During the 2002 test, methods did emerge. The most prominent example was the science team's use of a method for measuring atmospheric opacity or atmospheric dust content, also called "Tau

2002 FIDO Rover Test: Problems in the Convention

Team members had to work quickly because of the short timeline of each FIDO test sol. On some occasions, with as little as ten minutes to fill in the numerous software fields for each observation, scientists used descriptions that were quick and easy such as only identifying the instrument or a purpose designator (such as arm_plan or imagery_after_drive) and often left partially blank fields. Scientists did enter feature and target information in other fields in the software, because it was necessary for determining the placement of instruments and getting data. However, they often named a feature and used numbers to identify targets e. g. Feature: Rio Grande and Targets: Rio_1, Rio_2, Rio_3. Another theme group used a shorthand name for the instrument they were using (min for the MiniTES instrument) and a number as a target identifier, e.g. Feature: Kiptopeke and Targets: Min_1, Min_2, Min_3.

While numbers could indicate relationship between target points, which target was to be done first in a series for instance, or the relationship between returned sets of data in one observation, this procedure was problematic for other reasons. Numbers are less distinctive and memorable. When members talked about a target, team members found it harder to keep track of information that was identified with a number rather than a name, especially if there were a large number of targets being discussed. If a group decided to do future work on a feature, they had to remember or else check notes to find out what target numbers had already been used on that feature. During one sol, we observed the sequencing team dealing with two different identified Min_8 targets. This caused initial confusion during the commanding phase, but after some discussion, the engineering team was able to differentiate the targets. They were in distinguishably different locations, and differentially appropriate for the different science requests. There are other problems with numbers. The scientists themselves noted that using the same name and number across sols (Min_1 for example) resulted in confusion when searching for and retrieving data. We also knew that numbers would not scale over ninety or more sols of operation on the mission itself.

2002 FIDO Rover Test: Implications for a Naming Convention

FIDO 2 resulted in several findings important to the development of a naming convention.

- The separation of objects from actions in the name, as we had theorized, helped support referencing during discussions.
- The official identification of the instrument helped in knowledge and information management tasks.
- The connection between instrument and method became more apparent.
- Identifiers for certain abstractions, such as method, were appropriate and useful in the observation name. They established relationship between instrument work and features and added dimensionality to the identification of the science work.
- The ability to identify features separately from targets increased the team's planning perspective and their ability to talk about data at varying levels of granularity
- Under time pressure, scientists used purpose as an identifier

2003 Pre-Mission Tests and Trainings

The MER project began a series of tests in the fall of 2002 and throughout 2003 to integrate the science process, which had been developing in the FIDO tests, with the developing engineering processes. We continued to assess and refine the naming convention though out this period. We made recommendations to the science planning software that called for a hierarchy of observations with subordinate activities (see Figure 1.) and for the identification of both features and targets. Over this test cycle we also made recommendations for the sets of identifiers that would together become the ontology for naming observations and activities.

2003 Pre-Mission Tests and Trainings: Observation Name Development

We defined the taxonomy for the instrument package on the rover, and included the rover as an instrument (at the suggestion of the science team), to be used especially when the rover wheels were used to do trenching or displacement of the Martian soil. We defined shorthands for some of the instruments. We then differentiated between use of single or multiple instruments in an observation, and expanded the official taxonomy and the relationships between observation, activity, feature and target. For instance, from the FIDO 2002 test, we learned that *in-situ* observations, observations that place the rover's instrument arm on a rock or soil patch, are easier to reference than *remote sensing* observations, those that require pointing to objects in the distance. This is true in every day situations as well, where the distance to a target and the accuracy of both correctly identifying and hitting a target are correlated. For instance, the instruction to "place a penny on that flat rock in front of you" is a lot easier to follow than "take a picture of the dark area on the middle ridge on the second hill to the left."

We defined these two different observations types as in-situ and remote sensing, and determined that scientists should not ask for the two types in one observation request. The cognitive differences in these types of observations, planning difficulties and the configuration of the rover made them separate types of science requests. We established that "remote sensing" and "in-situ", terms already used by the science team, should be the observation name identifiers for work using more than one instrument while doing these kinds of work. The science team asked to use “IDD” and “PMA” as shorthand identifiers in the software for these two types of work. Those names referencing instruments located on the robotic arm Instrument Deployment Device (IDD) or on the Pancam Mast Assembly (PMA). Table 2.shows the relationships between single and multiple instrument use and instrument name identifiers.

Table 2.Single and Multiple Instrument Observations

Observation Type	Name (shorthand)
Single Instrument	APXS, Hazcam (Haz), Mössbauer (MB), Microscopic Imager (MI), MiniTES, Navcam (Nav), Pancam, RAT, Rover
Multiple Instrument	Instrument Deploy Device (IDD) Pancam Mast Assembly (PMA)

With each test, we continued to check the convention for internal consistency. Did the convention and its relationships hold for all types of observations that might be made? Tables 3 and 4 indicate the relationships we considered. The convention held across all observation types.

2003 Pre-Mission Tests and Trainings: Activity Name Development

In any given sol, the science team had to decide on observations that they wished to make and the associated features for those observations and enter the information in the software. Additionally, they had to instantiate observations with the activities that would fulfill the observation. It was the activity information that would be expanded by downstream teams into sequences and commands for the rover.

Initially, the mission design called for the scientists to simply choose activity types from a dictionary in the science planning software and populate the observation.[1] However, as the uplink process was tested in practice, the payload uplink leads (PULs), who expanded the activities into sequences, found that they needed more information from the science team in order to efficiently differentiate between the numbers of activities that appeared in their downstream software tool. For instance, a Pancam PUL might see multiple Pancam_mosaic_stereo activity types. To aid the work of the PULs, we recommended that scientists give unique activity level names to the activity types in the software. Science team members were to indicate important distinguishing parameters, such as filter color or the camera mosaic size (2x2), making the PULs' jobs easier. Once activity names were in place, we realized that the target name should be entered at the activity level not at the feature level. The appropriate place to identifying the actual target for instrument work was at the activity level. In addition, adding the target to the activity name made the name more unique.

[1] The activity dictionary was developed by JPL colleagues and was not part of our research. We simply incorporated the use of the existing dictionary into the work practice and into the developing ontology.

Table 3: In-situ Instruments and relationship between Instruments, Features and Targets

# of Instruments	**One Feature**	**Multiple Features**
One Instrument *(Use instrument or shorthand name)*	Include feature name MB Boulder	(Not Possible)
	Include feature name and relation or method for grouping unmentioned targets with feature MB_Sniff_Boulder	Include one or two feature names and relation or method for grouping unmentioned targets with features APXS_comparison_Boulder_ShipsProw
Multiple Instruments *(Use IDD instrument class name)*	Include feature name IDD_Boulder_	(Not Possible)
	Include feature name and relation or method for grouping unmentioned targets with feature IDD_Survey_Boulder	(Separate Observations)

Table 4: Remote Sensing Instruments and relationships between Instruments, Features and Targets

# of Instruments	One Feature	Multiple Features
One Instrument *(Use instrument or shorthand name)*	Include feature name Pancam_ShipsProw	Include one or two features and relation identifier or method for grouping features Pancam_Surveyaround_ShipsProw
	Include feature name and relation identifier or method for grouping associated targets with feature MiniTES_ShipsProw	Include one or two features and relation identifier or method for grouping associated targets with features MiniTES_Surveyaround_ShipsProw
Multiple Instruments *(Use PMA instrument class name)*	Include feature name PMA_postScratchSniff_ShipsProw	Include one or two features and relation identifier or method for grouping features with target name PMA_Surveyaround_ShipsProw
	Include feature name and relation identifier or method for grouping associated targets with feature PMA_Surveyon_ShipsProw	Include one or two features and relation identifier or method for grouping associated targets with features PMA_Surveyfrom_ShipsProw to Boulder

2003 Pre-Mission Tests and Trainings: Syntax for Observation and Activity Naming

We determined a syntax for both observations and activities. Based on our analysis of the science teams' use of different identifiers while working in face to face collaboration, we structured the relationships between the identifiers. We based the syntax on relative importance for knowledge management and referencing, a desire to create consistency within the software and the need to frame an observation thus delimiting it at both ends of its representation. Features were consistent with the higher level of abstraction found at the observation level, while target was more consistent with the activity level of abstraction.

At the observation level, the syntax represented the relationships between instrument, method, feature and other identifiers. The first identifier was instrument and the last was feature.

- The syntax for naming an observation became

Instrument_Method_Other Identifier_Feature

- Multiple instrument observations had the following syntax

PMA_Method_Other Identifiers_Feature
IDD_Method_Other Identifiers_Feature

At the activity level, the software already contained a set of activity types within an activity dictionary. These were types of activities that had associated sequences and resource models. The purpose of entering an activity name was to uniquely identify an already existing activity type to convey important information to downstream teams. Scientists could choose the activity type from a menu in the software. The scientists did not have to re-enter the activity type name, such as pancam_single position, but did have to enter the identifiers in a name that made it unique. The identifiers were important distinguishing parameters, such as filter color and the target.

- The syntax for naming an activity became

Distinguishing parameter_Target.

2003 Pre-Mission Tests and Trainings: Method Name Identification and Development

Method as an overall category became an increasingly important identifier for the emergent ontology as the science team developed more and more methods in their work. Some examples of methods follow.

- Accordion (an observation that can expand or compress as the time or data volume allows)
- Approach (driving close enough to a Feature to use the robotic arm)
- Blind (for MiniTES activity without a supporting image or Pancam activity without a target)
- Comparison (association of two Features)
- Drive (movement of the rover)
- Camera variations while driving ”quick look”, “rubber neck”, “systematic”, “sashay”
- Movie (repeated camera images looking at the same feature over time)
- Rat (use of the Rock abrasion tool RAT)
- Scratch (use of the RAT to grind a rock)
- Sniff (use of one of the other in-situ instruments)
- Surveys: Survey around, between, covering, from. . to, including (conducting a broad overview that might cover more than one feature)
- Brush (the use of the RAT instrument not to grind into the rock but to clear debris from the target area)
- Tau (identification of atmospheric opacity or amount of dust in the atmosphere)
- Trench (using the robotic wheels of the rover to dig a hole)

2003 Pre-Mission Tests and Trainings: Other Identifiers and Constraints

As we watched the scientists work and analyzed that work, we realized that the observation names had to be consistent yet flexible. No single or fixed convention would ever be appropriate to support the multiple types of observations, the complexity of the work and accommodate learning and changes that would take place as the team discovered new ways to use the rover and new methods. Finding a syntax and structure that would frame the problem while accommodating change was a challenge. We knew that instrument was an important identifier, as was method and feature. Then we realized that the missing piece in the convention was an open invitation to enter what we came to call “other identifiers” as they were deemed appropriate. We had already recognized a number of these identifiers from our analysis of the scientists’ work. Most of them came under the classification of temporal or spatial constraints.

Temporal constraints:

- Words like: afternoon, morning, morning after, AM, PM, overnight, and before and after through the use of the terms pre and post.

Spatial constraints:

- Direction: North, South, East, West and elevation (for sky observations)
- Distance: Long and Short (often associated with rover approach moves toward a feature or with drives)
- Location: around, between or the identification and name of an area surrounding a feature

Other identifiers: helped situate the context of the environment.

- Soil: because soil features were so ambiguous by their nature, identifying an observation as being on a soil feature was often helpful

Table 5. Examples of Observation and Activity Naming and their meaning:

Observation Names	Meaning
MI post rat Buffalo	Take a Microscopic Image of Buffalo, after using the Rock Abrasion tool
Mini-TES_Movie_30deg_Sky	Take several consecutive MiniTES measurements of the sky at a 30 degree elevation
IDD_Post Scratch_Plymouth Rock	Take several different kinds of in-situ measurements of Plymouth Rock, after scratching the rock with the RAT
Activity Names	
Red single Pilgrim	Take a single frame image of the target pilgrim, using the red filter of the Pancam
5 filter vent_center	Take a Pancam image of the target vent_center using five filters.
4x1 Knob	Take a 4 by 1 MiniTES raster measurement of the target Knob

2003 Pre-Mission Tests and Trainings: Issues

Pre-mission, we were especially interested in seeing how two parts of the convention and ontology would work out. The first related to the difficult problem of pointing at objects in the distance. Because the rover did not "know" about pointing, we attempted to work within the early functionality of the software by creating a shared conceptualization about pointing for capturing distant information, such as for camera work and remote sensing Mini-TES instrumentation. We had scientists identify remote features and targets and then indicate activity on them, with the idea that instruments would focus on the target point. We knew, however, that difficulties might arise because pointing to the distance from a rover that is constantly changing its position is challenging. Reference to distant objects can be determined either from the perspective of the object being pointed at, as in a focal target point in the distance, or from the perspective of the pointer, as in a field of view (FOV). If the rover always stayed in a fixed position, one could identify targets by establishing a fixed grid over the field of view and then setting targets. However, that was not possible for the MER rovers and our ontology had to facilitate shared conceptualizations for remote pointing.

After we had been working on the ontology for some time, some new functionality was added to the software that allowed for designation of more footprints and a FOV designation. Because of this new functionality, scientists now had a choice between naming a focal point with targets and features and naming from a FOV perspective. For example, a navcam request could ask for two 45degree images to cover a slightly less than 90 degree FOV (because of the two stereo images overlapping). So a scientist could ask for the two images to be taken using a directional indicator, such as azimuth 30 degrees. Or the scientist could ask for navcam image pointing toward a target on a feature in the distance. We were curious as to how the convention would support both types of requests.

Our second issue was a concern with how closely the scientists would follow the convention for observation and activity naming. We knew that observations and activities as well as features and targets were high level classifications that were now supported by the software design. However, the field for writing in a name was an open space and scientists could theoretically fill in anything they wanted. The team had gone through training and mission procedures required they follow the convention. However, as work practice researchers, we knew that people would learn to do their work as efficiently as possible. They would use the ontology and our procedures to the degree that they supported their work. We expected there would be variations to the ontology in actual use. Shared conceptualizations might change. The work might call for new parts to the ontology. We did not know what impact events would have on information and knowledge sharing within the group or with downstream teams of engineers.

2004 Mars Exploration Rover Mission

Once the two rovers, Spirit and Opportunity, had successfully landed on opposite sides of the Martian surface, our work focused on assessing how well the ontology supported the work of the science teams, now split into two groups each working at the different landing sites. Generally, based on a initial, random selections of the mission data, we found the constructs we developed held over time. The high level representations of observation, activity, feature and target became a part of everyday work practice. Scientists used these not just within the prescribed functionality of the software but also, with some exceptions, in report writing and face to face discussion. Features were the most obvious referent. Targets were referred to when mentioning specific data or work of the in-situ (IDD) instruments, but otherwise became a much more implicit part of the work, at least during the science discussions. Our findings from the FIDO 2001 were confirmed and our recommendations supported.

2004 Mars Exploration Rover Mission: Work Practice and Naming

The basic ontology held through the mission. The science team did not follow the syntax precisely. Scientists adapted the naming convention as they gained experience with the operational environment and developed methods of tele-robotic exploration. As time went on, they began to situate important information, not feature, at the end of the name. Our idea was to frame the name so that feature always identified an end to the name. They placed important constraints and other identifiers at the end of the name. We suggest that this may have been a way to make important information more accessible.

Specific observations, examples, and deviations appear in Table 6 below and will be discussed below in the order of their appearance in the table.

Table 6. Observation Name Highlights

	Mid Mission	End of Nominal Mission
Temporal Constraints	13:30 LST Midday Anytime Post MB Prebrush Sol 46 PreMGS Ultimate/penultimate/ antepenultimate	Before 14:30 Post backup Plan A, IF Dist GT .085m Overnight science Pre or Post ODY
Methods	Traverse clast survey Mini-MiniTES Stutter step	Super clast survey Ground Stare 3x1x255 Stares
Purposes	Recon Transient Temperature Doc	Dust Devil Finder Phobos Set
Features	Trex cheek Soil Ejecta blanket IDD work volume	Crater floor Heatshield

2004 Mars Exploration Rover Mission: Temporal Constraints and Other Identifiers

The category of temporal constraints expanded dramatically, with a number of different subcategories. Specific (numeric) and general timing constraints appeared as proxies for changing temperature and lighting, and with these specific proxies, the need also arose to indicate the absence of temporal constraint (anytime). Temporal constraints also expressed synchronization with rover events, for example to ensure that the data from two observations reflected the same underlying conditions (e.g., Post MB). Some observations (destructive) changed the conditions of the world (e.g., driving, trenching, RATting), requiring scientists to identify the states they were assuming when they parameterized an observation (e.g., prebrush). Late in the mission, when the planning cycle diverged from Mars time, the scientists began to make conditional plans that depended upon unknown data values (e.g., If traveled distance were greater than .85m). Atmospheric scientists found a rarely claimed time slot after the next sol's wake-up but before the next Uplink, so they often prefixed an observation with an incremented sol number. Several mission events influenced the timing of an observation, including the availability of a downlink (data relayed from the Odyssey Orbiter) so that data could arrive before the next cycle of planning. This resulted in names that specified the downlink pass that scientists would like. Spirit's rock-laden landscape broke up a traverse into small steps (ultimate and penultimate), and scientists identified observations with each of these steps. These numerous examples confirmed our expectations that other identifiers and constraints were needed to support unanticipated needs.

As the mission progressed, scientists began to re-arrange the order of these identifiers as we have mentioned. For example, scientists doing remote sensing (PMA) work placed temporal constraints at the end of the name in the early part of the mission. Example_Pancam_Tau_Anytime. When the mission moved into an extended operations phase, the planning process became more standardized and engineers began to use templates for pre-planning activity requests. Because the temporal constraints were key in this template planning, engineers requested that the science team place temporal constraints first in the observation name. So instead of 'Pancam AM' they wrote 'AM Pancam'.

2004 Mars Exploration Rover Mission: Method Development

Experience with the specific tool suite lead to the development of numerous specific methods such as a scuff and go, brushing, mini-Mini-TES, and stutter step. The science team also named different ways to drive as the table indicates, new method names were still appearing after 45 days of operations. The ability to name clusters of activities with a single label lends support to the idea of observations as containers that render the work coherent.

As we had started to see during the field tests, purpose began to appear in the observation names. Some of the purposes were primarily operations-relevant such as reconnaissance or turning for communication. However, some of the purposes were scientific, such as documenting transient temperature,

2004 Mars Exploration Rover Mission: Feature Names

The use of features in the observation name also evolved with the mission. We had recommended that scientists use a whole-part relationship when identifying features and targets to help with information and knowledge management during Uplink discussions and in finding information in returned data. Whole part relationships were used more consistently with the in-situ IDD instruments. We discuss this further when we turn to activity names and the use of targets. Here we note that target name, which was supposed to be an activity identifier was elevated to the observation name to create specificity and distinctiveness. For repeated work on a feature, the science team wanted to make sure the all team members and down stream participants understood the exact spot to do the new work. A number of generic feature names appeared, starting with sky which we had seen in field tests, and introducing other regions such as soil, drive direction. The mission also resulted in some feature names that implied geological origin (e.g., ejecta blanket). Finally, scientists used rover or space craft parts as a feature, such as IDD work volume, RAT magnet, cal target, heat shield, magnets.

2004 Mars Exploration Rover Mission: Activity Name Development

Table 7 indicates some of the additions to activity names that appeared during the mission. Some activity names included temporal constraints, but certainly not with the regularity we observe in observation names. Activity names also acquired some method names, generally referring to parameter settings. Purpose also crept in to activity names, to capture both operational and scientific rationale. The most prevalent descriptor on an activity functioned as both a method and a target. When doing remote sensing pointing, as we had anticipated might happen, scientists increasingly used the FOV perspective, relying on azimuth and elevation numbers and footprints to define the FOV. Remote sensing data collection was most accurate for pointing when done before a rover move, because the feature locations did not change after movement. Because the rover itself did not “know” about pointing, the command used azimuth and elevation co-ordinates. Since there was no way to know the exact set of the rover coordinates after it had moved without analyzing returned telemetry data, the most accurate azimuth and elevation pointing was done before a move. Further, features were less important in remote sensing, because the product of remote sensing is typically a region rather than a particular spot. The data gathered from remote sensing observations provided context and information for making decisions about next steps in the process. These kinds of observations were done frequently and many activity names reference azimuth and elevation rather than target.

Table 7 Activity Name Highlights

	Mid Mission	**End of Nominal Mission**
Temporal Constraints	16:10 Nighttime MI preMB Post Drive Ultimate/penultimate	Daytime Postgrind
Methods	Cal target filters Triple Play Color stereo	Cal plus sweep magnet 1x1x50 Block
Purposes	For MTES overlay Verify placement Mineralogy Layer Study	Document placement Verify position
Features & Targets	Cherry center Below Sun Rear view tracks	Target 1 Placement 1 Drive direction Filter magnets

We also saw the occasional use of numbers to identify targets. Sometimes scientists simply preferred to use numbers to identify a target. In the context of a particular static situation it seems an acceptable practice and so team members sometimes used it. There was another variation in target naming work practice that was of particular interest, however. Theme group members sometimes had to identify several target points in the software before finding the exact spot for the placement of the RAT on a rock because the placement on the rock had to be optimal for surface abrasion and yet within the reach of the rover arm. Because of this, the team sometimes used numbers to identify various candidate targets. Sometimes they would use the feature name with the number to help keep the number in context, such as McKittrick_1, MicKittrick_2. We suggest here that the science team found it the most expedient way to target a number of points, knowing that they would use only one in the end. Cumulative knowledge management of these variations was not as serious an issue as it would have been if every target in the mission had been identified only by a number.

2004 Mars Exploration Rover Mission: Observation Name Development

As the mission went on, observation names got longer. We believe this tendency correlated with the increased use and a standardization of methods (Shalin et al., in prep) and the indicated desire of the science team to make sure that important relevant information was obvious in the software at both the observation and activity levels. For example, in-situations where scientists were requesting re-work on the same feature, they sometimes elevated the new target name to the observation level to make sure that others understood that this was a request for new work. Important parameters were also elevated to the observation level on occasion. As the work approached an extended mission phase, constraints (Odyssey communication pass, or ODY, and later in day, or PM) went to the front of the name

Examples of a longer name from later mission work are:

MTES Elevation Sky AND Ground ODY PM
Pancam Midway 1 4Fs (Four Filters on Soil)
PM ODY mini TES Elevation Sky AND Ground Beta Pancam Photometry Photometric Equator[2]

Extreme variations of use outside the parameters of the ontology and the procedures were rare. However there was one instance where a scientist entered a variety of non-standard characters, such as “&” and “%” into an

[2] This is a multispectral Pancam along the photometric equator. The Beta Pancam Photometry was an addition to the name to group four coordinated observations together.

observation name. The downstream software did not recognize the information and the observation was subsequently lost in the automatic, digital file transfer to the engineering team members. The record of the original observations was in the previous file format, and scientists who worked in the downstream process realized that the observation, an important one for that sol's activity, was missing. They re-created the observation, but it took a significant amount of time, Not only did a failure to work within the agreed specifications result in inefficiencies, it posed a risk to the success of the sol. Time was a precious resource in the mission. Failure to approve and command an activity plan in time to meet the Deep Space Network's transmission window meant the loss of a day's science. Variations such as this highlight the importance of agreed upon conceptualizations and the need for specification as information flows through different software applications.

V. Discussion

The need to share structured information that represents shared conceptualizations (Musen,1992; Gruber 1994) is at the basis of ontology development. As we have shown in this paper, the need to define, frame and standardize shared conceptualizations (abstract models) of the work of Martian tele-robotic science into consistent representations was the main driver for our ontology development. Generally, in established domains, ontology development takes place when an expert ontologist works with one or more domain experts to translate existing domain knowledge into an ontology (Noy et al). In the present study, there was very little pre-existing knowledge about the need for an ontology, nor about the demands of the field of tele-robotic geologic science. When their training began, the MER scientists were neither expert nor proficient in the method and process of doing time-delayed tele-robotic science with a rover on Mars. Over the period of their training and throughout the mission, the scientists learned constantly and continued to refine their work practice. We identified and framed parts of the developing knowledge and the shared conceptualizations as they emerged from the on-going work. Over a three year period before the mission, we worked to identify the components of that work and create a frame for consistent representations that would support natural language discussions at one end of the work process, translate through a variety of software formats and provide references translatable into rover commands at the other end.

The resulting ontology for tele-robotic work meets the definitions for an ontology referred to by Gruber (1994), Guarino and Giaretta (1995) and Guarino (1997). That is, it is not so much a complete *specification* of shared conceptualizations as it is an incomplete or *partial agreement* about those conceptualizations. The later point is important to our Case Study. One of our central findings was that in the dynamic, environment of Martian tele-robotic science, terms require semantics that are not just internally consistent within the software. Rather, they must also support action with an on-going, dynamic external environment. Additionally, those terms must adjust to users who are themselves changing, that is, learning over time. An internally consistent, intensional semantics would not be sufficient to constrain meaning in such a dynamic environment. This finding is consistent with past work (Greeno, 1983) that states that as leaning takes place, new conceptualizations will develop. It is consistent with the understanding in current ontology development (McGuiness, 2001) that ontologies require extensibility, or the ability to adapt to the needs of users and projects. We believe that that this Case Study demonstrates how an ontology for action in a dynamic environment demands the greatest flexibility, and that when constructed for an emerging domain, the ontologist should expect dramatic, frequent revisions and have the capability to capture and support those revisions.

Our ontology had to represent not only the *objects* in the domain and on the rover (features, targets, instruments), but also *events* (communication passes, ends of drives), *constraints* (time, before and after), as well as the *action* in the domain (methods, rover movement and activity, rover interaction with differing terrains, interaction with moving robotic satellites), and user learning. Because of this dynamic, we determined that the ontology must adapt and represent both incremental change and revolutionary change over time. The addition of the open concept of “other identifier” into our naming convention allowed this to happen. We found scientists included temporal and spatial constraints more than any other kind of identifier in this conceptualization. However, other kinds were used, for example indicators of relationship to other observations and between instruments.

Change was a constant in this domain, and so we suggest that our ontology can never be more than a partial agreement on the conceptualizations involved. We believe, however, that the constructs in our ontology are at a consistent and sufficient level of abstraction to capture the basic work of the domain. They were adequate to at least explicate most new conceptualizations, as evidenced by their continued presence and elaboration throughout the mission. We believe also that the basic categories within the ontology will be re-useable in other tele-robotic domains.

VI. Conclusion

An Ontology of Tele-robotic Work

Observations and activities represent basic constructs of the work of tele-robotic science as they support the construction of a science activity plan that makes requests of the rover. Each observation and activity must have an associated location relative to the request, e. g. a feature and a target. These four constructs, together with method, instrument and a space for other identifiers form a part of an emergent ontology of tele-robotic work (Figures 5).

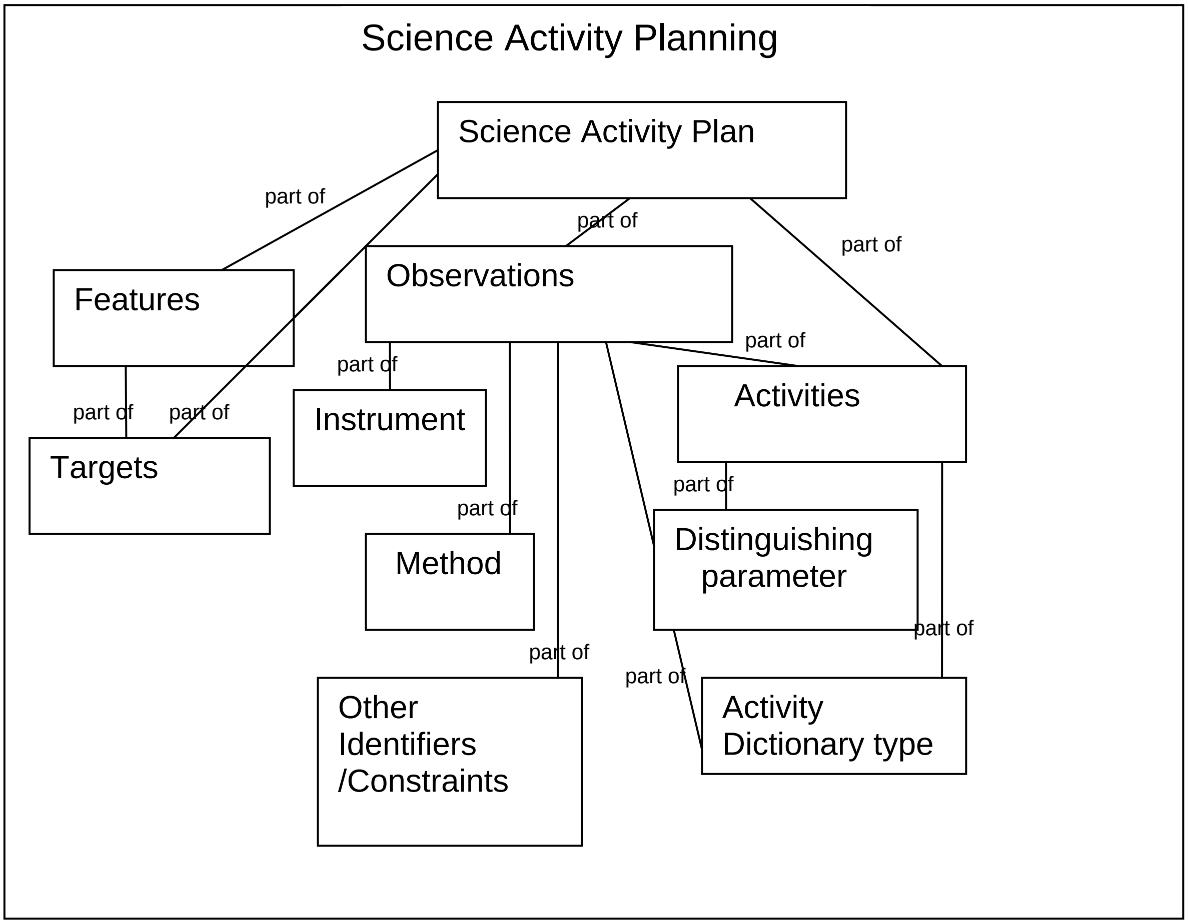

Figure 5. The parts of the science activity plan

In this emergent ontology, as can be seen from the figures, there is a structure of shared conceptualizations, a classification of things and defined relationships between all of these conceptualizations. At the observation level, the ontology allows for the definition of natural language identifiers that can be used for decision making, planning and then honed for commanding as they move through the work system and the software system. (Figure 6.)

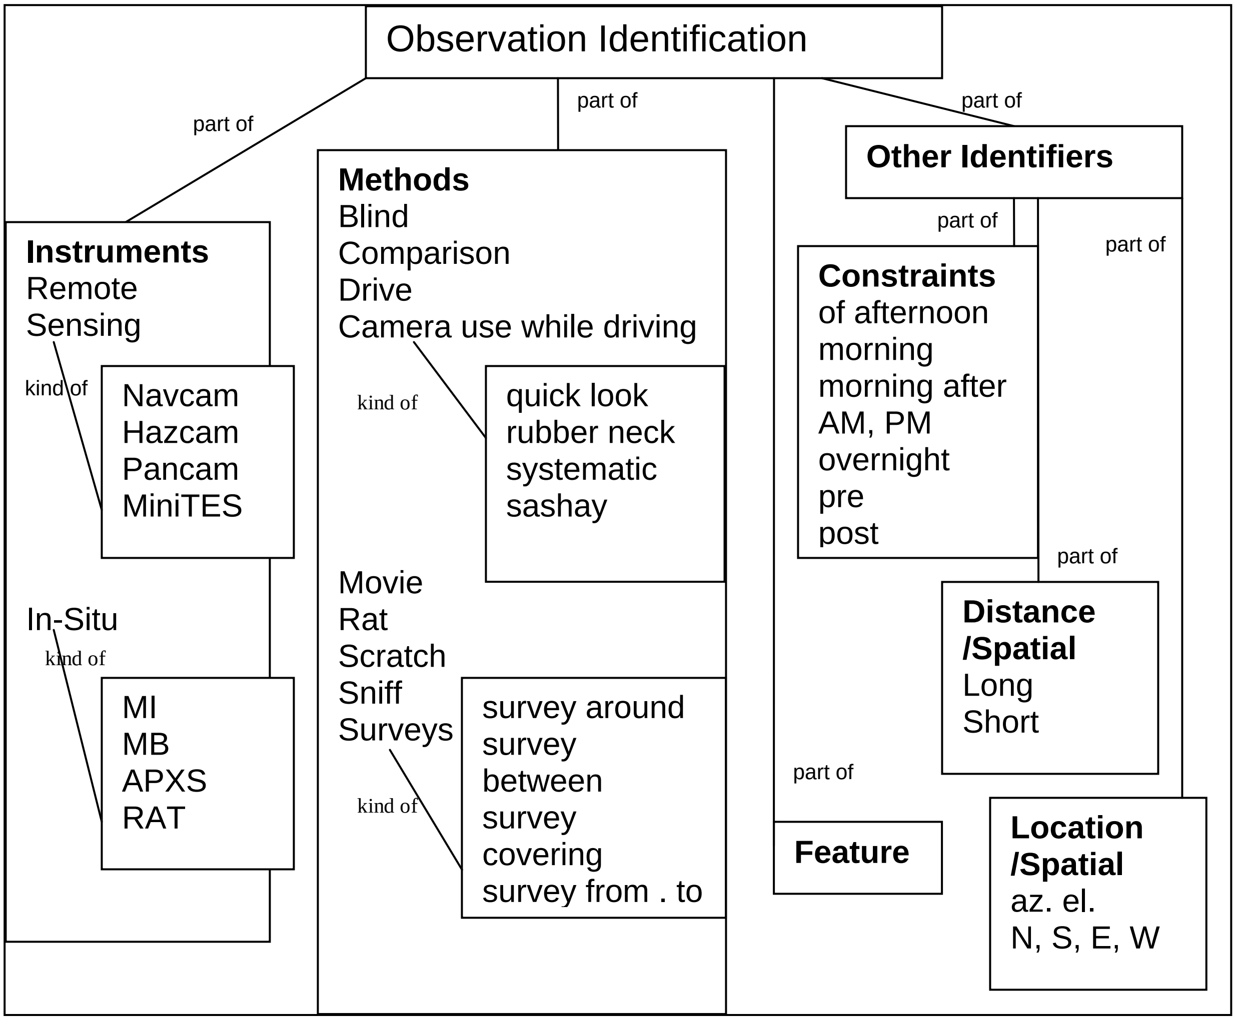

Figure 6: An ontology of Observation Identification

Analysis of MER mission events demonstrates that learning in an expert domain continues and as a result work process and even tool use can change over extended periods of time. The parts of the observation and activity name, their semantics and their syntax all were important constructs in this mission. However, we found that because scientists could further develop existing classifications such as method as well as enlarge the relationship of their work through new shared conceptualizations under “other identifier” helped the ontology adapt and support learning and change. This flexibility is essential in any expert domain that interacts with a dynamic environment.

Originally, the science planning software provided scientists with a blank field to enter the observation name and parts of the activity name. Since no one knew exactly what the work would look like, constraining it to some hypothetical convention would have been problematic, as we saw demonstrated in the very first field tests. While the open field limited the formality of requests, it offered flexibility and supported learning as scientists became increasingly sophisticated in their work and created new shared conceptualizations, which we identified and formalized into the emerging ontology. As the domain has developed during the mission, our analysis shows that the ontology held certain consistencies, and we believe that fields congruent with our ontology can be implemented.

Fields that contain pre-set taxonomies of instrument, constraints, and some methods will offer all scientists and all engineers the capability of viewing the information that is most salient to their work, represented in a way that can be flexibly re-configured for each task. The increased formality will capture and present information consistently across the various tools within the system, but allow for unique identification and discern-ability as teams carry out their tasks.

We believe that the definition of the emerging ontology in this domain served several purposes. It provided procedural and work practice constraints and an agreed upon language for use. It established consistency for software representations. It assured that information important to the problem was incorporated in the ontology, and it allowed for the unique specification of requests for the science team and downstream engineering teams. And finally, it captured shared conceptualizations and representations for the historical record. While we were working to support the work of the MER mission, a longer term goal is to provide the basics for a more developed ontology for tele-robotic science. Future software applications can contain fields for some of this work, such as instrument and method, feature and target and the work can begin with a certain taxonomy and ontology in place. We believe, also, that many of the constructs for requesting distant action are applicable to other complex, distributed work systems and to the archival record of their accomplishments.

VI. REFERENCES

Behrend, D. A. (1995) "Processes Involved in the Initial Mapping of Verb Meanings", in M. Tomasello & W. E. Merriman (Eds.), Beyond Names for Things. Hillsdale, NJ: Erlbaum, pp. 251-273.

Bloomberg, J., and J. Giacomi, A. Mosher, and P. Swenton-Wall (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.

Dreyfus, H.L. (1999) *What Computers Can't Do: A Critique of Artificial Reason,* New York: Harper and Row.

Fikes, R.E. and N.J. Nilsson (1971) "STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving", Artificial Intelligence, 2, pp 189-208.

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.

Friedland, P. et al. (1982) “Genesis, a Knowledge-based Genetic-Engineering Simulation System for Representation of Genetic Data and Experiment Planning”, *Nucleic Acids Research,* 10(1), 323-340.

Furnas, G.W. et al. (1983) “Statistical Semantics: Analysis of the Potential Performance of Key Word Information Systems”, Bell *System Technical Journal.* 62(6), 1753-1806.

Gentner, D. and A.L. Stevens (eds.) (1983) Mental *Models,* Hillsdale, NJ: LEA.

Georgeff, M.P. and A.L. Lansky (1986) “Procedural Knowledge”, Proceedings *of the IEEE.* V 74 (10) 1383-1398.

Greeno, J.G. (1983) “Conceptual Entities,” in D. Genter and A. Stevens (eds.) *Mental Models,* Hillsdale, NJ: LEA.

Grice, H.P. (1975) Syntax *and Semantics,* New York: Seminar Press.

Gruber, T.R. (1993) “A Translation Approach to Portable Ontology Specification”, *Knowledge Acquisition,* 5, 199-220.

Gruber, T. (1994) email communication, SRKB Mailing list, in Guarino, N. "Understanding, building and using ontologies," *International Journal of Human-Computer Studies*, vol. 46, pp. 293-310, 1997.

Guarino, N. (1995) “Formal ontology, conceptual analysis and Information Systems”, in *Proceedings of Formal Ontology in Information Systems*, Washington DC: IOS-Press.

Guarino, N. and P. Giaretta, (1995) “Ontologies and Knowledge Bases: towards a terminological clarification” (1995) In N. Mars, (ed) Towards Very Large Knowledge Bases; Knowledge Building and Knowledge Sharing, Amsterdam: IOS Press, pp.25-32

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.

Lakoff, G. (1994) “What is a Conceptual System”, in W.G. Overton and D.S. Palermo (eds.) The *Nature and Ontogenesis of Meaning.* Hillsdale, NJ: LEA.

McCloskey, M. and A. Caramazza, and A. B. Green (1980) “Curvilinear Motion in the Absence of External Forces - Naive Beliefs about the Motion of Objects”, *Science.* 210 (4474): 1139-1141.

McGuinness, D.L. (2001) “Ontologies Come of Age”, in D. Fensel, J. Hendler, H. Lieberman, and W. Wahlster, (eds.) *The Semantic Web: Why, What, and How*, MIT Press.

McGuinness, D.L. (2000) “Conceptual Modeling for Distributed Ontology Environments,” in *Proceedings of The Eighth International Conference on Conceptual Structures Logical, Linguistic, and Computational Issues* (ICCS 2000), Darmstadt, Germany , August 14-18.

Musen, M. (1992) “Dimensions of Knowledge Sharing and Reuse,” *Computers and Biomedical Research,* 25, 435-467.

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

Noy, N.F. and D.L. McGuinness (unpublished manuscript) “Ontology Development 101: A Guide to Creating your First Ontology.

Rosch, E. and B.B. Lloyd (1978) *Cognition and Categorization,* Hillsdale, NJ: LEA.

Sacerdoti, E.D. (1977) *A Structure for Plans and Behavior,* New York: Elsevier.

Schank, R.C. (1972) “Conceptual dependence: A theory of natural language understanding”, *Cognitive Psychology,* 3(4), 552-631.

Schank, R.C. and K.M. Colby (1973) *Computer Models of Thought and Language,* San Francisco, W.H. Freeman.

Searle, J.R. (2002) *Consciousness and Language,* New York: Cambridge University Press.

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

Shalin, V.L. et al. (in preparation) Development of work methods in novel tasks.

Stefik, M. (1981) Planning with constraints. *Artificial Intelligence* 6(2), 111-139.

Wales, R., J. O’Neill, and Z. Mirmalek (2002) Ethnography, customers and negotiated interactions at the airport. In special issue on Human Centered Computing at NASA. *IEEE Intelligent Systems Journal, September/October*. Vol 17 No 5 pp 15-23.

Appendix A: LIST OF ACRONYMS

APXS	Alpha Particle X-ray Spectrometer
FIDO	Field Integrated Design and Operations
Hazcam	Hazard Avoidance Camera
IDD	Instrument Deployment Device
JPL	Jet Propulsion Laboratory
MB	Mössbauer Spectrometer
MER	Mars Exploration Rover
MI	Microscopic Imager
Mini-TES	Miniature Thermal Emission Spectrometer
MTES	Miniature Thermal Emission Spectrometer
Navcam	Navigational Camera
Pancam	Panoramic Camera
PMA	Pancam Mast Assembly
PUL	Payload Uplink Lead
RAT	Rock Abrasion Tool
SAP	Science Activity Planner
SOWG	Science Operations Working Group

About the Authors

Roxana C. Wales, Ph.D. is a human-centered computing research scientist at the NASA Ames Research Center. As an ethnographer focused on work system design and evaluation, she is especially interested in revealing explicit as well as implicit information in technological settings that is essential for creating better work systems and supporting technology use within expert domains. At NASA, she has worked on projects in Mission Control at Johnson Space Center for both the International Space Station and the Space Shuttle. She also conducted a two and a half year study of airline operations and airline delays at a major airline for NASA.

Valerie L. Shalin Ph.D. is currently an Associate Professor in the Department of Psychology at Wright State University, and has held positions at SUNY Buffalo Department of Industrial Engineering and Honeywell Systems and Research Center. She has thirteen journal articles and book chapters, and numerous proceedings papers in the area of human workplace expertise and associated empirical and analytic methods to support the design and evaluation of workplace aiding and training technology.

Deborah S. Bass, Ph.D. currently works at the NASA Jet Propulsion Laboratory, where she was the Deputy Science Team Chief for the Mars Exploration Rover (MER) Project. Dr. Bass is now the Project Science Systems Engineer for the 2007 Mars Lander, Phoenix. She received her PhD from the University of California at Los Angeles, and her science expertise focuses on Mars polar geology, with an emphasis on water transport in and out of the polar regions