REQUESTING DISTANT ROBOTIC ACTION:

An Ontology of Work, 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 an ontology and the associated abbreviated names that comprise requests for action of a distant robotic rover during the 2003-2004 NASA Mars Exploration Rover (MER) mission, run by the Jet Propulsion Laboratory. We demonstrate how an ontology can support action in work and create structures for organizing that work *(an ontology of work)* as well as support the identification of objects and their semantic relationships for archiving, sharing and identifying information within and across software tools. An ontology of action and work must take into account a dynamic environment, changing in response to natural physical events as well as intentional actions and reflect the influence of context on the meaning of action. The nascent domain of Martian tele-robotic science, in which specialists request work from a rover moving through a distant landscape, as well as the need to consider the interdisciplinary teams involved in completing that work, required an empirical approach to understanding the issues involved. The formulation of this ontology uses ethnographic methods and is grounded in a study of human behavior and work practice.

Analysis of MER mission events demonstrates that work practice and robotic tool usage change over time. Therefore, we suggest that an ontology must adapt and represent both incremental change and revolutionary change, and that the ontology can never be more than a partial agreement on the conceptualizations involved. We offer recommendations for understanding and eliciting knowledge within a domain and creating ontologies in expert domains.

I. 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: "We need 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 and will keep adjusting the work (e.g., filling the paper tray) until the task is accomplished. The success rate may change if the requests become more complex, asking for copies of multiple originals, giving different instructions for different originals, making multiple requests concerning the same original, and asking for copies from originals that are something other than a paper document (such as an electronic file). However, an abbreviated name will most often be used to identify and give meaning to the task, even though the execution is invariably far more intricate than the apparently simple one sentence request would imply.

This paper focuses on the development of abbreviated names and the related ontology that supported the collaborative work of multiple types of specialists making requests of a robotic rover in the scientific exploration of Mars. The ontology was used to frame and instantiate plans for rover action and carry that information across different tools and software in the mission uplink process. We suggest that the abbreviated names and related ontology defined the work in such a way that together they are an *ontology of work.* That is, their meaning, format and usage is integral to the referencing and identification of information, and that it frames the work process and carries the information across tools and

teams as it is executed in the work system.

This research took place over a three and a half year period during the design, development and operations phases of the 2003-2004 NASA Mars Exploration Rover (MER) mission, run by the Jet Propulsion Laboratory (JPL) in Pasadena, CA. In established domains, ontology development generally occurs via a knowledge acquisition expert, who works with one or more domain experts to translate existing knowledge into an ontology (Noy et al, Forsythe et al, 1989; Meyer, 1992) In the present study, the work of doing tele-robotic, geologic science with a rover moving large distances across the Martian landscape was new, and domain knowledge regarding how the work should be done was limited. When their training began, the MER scientists were neither expert nor proficient in the process of doing time-delayed, tele-robotic science of this type. Over the period of their training and into the mission, they became increasingly proficient and expert, while continuing to refine their work practice.

While the task of extracting and formalizing the important constructs and distinctions from established domains of practice is not trivial, our research required another order of ontology development in many respects. First, the domain was emerging as work practice developed, and the ontology had to evolve with that work. Second, the interdisciplinary nature of the work required an ontology that was suitable for both engineering work and science work in both software and in natural language collaborations. Third, scientists and engineers needed semantics and labels not just for objects but also for higher-order work actions, consistent with the view that state changes alone are insufficient to communicate the semantics of action (Georgeff et al, 1986) and the idea that action is represented at multiple levels of analysis (Sacerdoti, 1977) Fourth, the success of the ontology required more than internal consistency. Critics of computational linguistics (and computational models of mind) specifically note that natural language is contextualized (Dreyfus, 1979) and intentional (Searle, 2002), damping enthusiasm for being able to formalize natural language as

unambiguous, stand-alone constructs for computational manipulation. While much contemporary work on ontologies can focus on applications involving search and pattern matching (McGuiness, 2001), the role of context and intention is critical for commanding robots that effect action in a dynamic world. 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 requesting 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 a request for taking a picture of a particular target at one time of day does not have the same outcome as an identically parameterized request for an image taken at a different time of day.

The nascent domain and the interdisciplinary nature of the work in this study required an empirical approach to the formulation of an ontology that was grounded in human behavior. As researchers, we had to understand the work as it evolved from earliest tests and training exercises through surface operations on Mars. As mission participants, we had to identify, reflect on, and feed back emerging constructs to the domain in order to support the developing work and meet the demands of the mission timeline for the landings on Mars.

Drawing on the Case Study, this paper describes how an ontology can:

1. Create structures for organizing work (an ontology of work).
2. Identify action in a changing environment. An ontology of action must take into account a dynamic environment, changing in response to natural physical events as well as intentional actions and reflect the influence of context on the meaning of action
3. Identify the relationship between action and objects in a domain, and define semantic relationships for archiving, sharing and identifying information.
4. Translate the work of one set of experts for use by other experts.

The next two sections describe our Research Methodology and the Case Study. The Case Study includes descriptions of both the design and operational phases of the mission. In it, we first offer a description of the current MER mission work system for the tele-robotic exploration of Mars. The work system is the result of three years of iterative design and training. Then we offer a review of the development of the ontology that helped structure parts of that work system and frame shared conceptualizations. Finally we assess the work during the mission itself. Our goal during design and development was to identify consistent cognitive and software representations that would support the work of mission personnel as they created activity plans for tele-robotic work on Mars.

II. THE RESEARCH METHOD

We used ethnographic methods to both capture and shape the ontology as it emerged from the work. 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 the current authors was the science operations systems engineer and later the deputy science team chief for MER and had daily access to on-going mission design work. The others authors were tasked by NASA to provide human-centered computing work systems design recommendations to the mission, spending extended periods of time over three years at the Jet Propulsion Laboratory (JPL). Our team brought a cross-disciplinary perspective to the research, drawing on backgrounds in geology, cognitive psychology and cultural psychology/anthropology.

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. 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. We presented and iterated on the developing ontology and naming convention with the science team during those and other meetings. As badged members of the mission, 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. As researchers, 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 from 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.

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 the

issues that would emerge. The extended nature of our participation observation allowed us to identify and become completely familiar with the problem, consider what it would mean to have it solved, capture the developing expert 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. The mission planning and commanding process required the development of an inherent electronic record of work that we could access and analyze. For this paper, we have done a qualitative analysis of those data.

The use of ethnographic methods allowed us to do iteration testing, make supporting process and procedure recommendations, evaluate usefulness and appropriateness, and do a functional assessment of the ontology over time within the work of the domain. We were concerned with the development of a naming convention and how an ontology behaves in practice. It is impossible to say what naming convention would have developed had we not made our recommendations but rather simply observed and gathered data as the work practice developed on its own. The naming convention may have developed with longer or shorter names and with less information and a greater reliance on in-situ understanding and context resulting in problems developing down stream when the context was missing. Given the other work and training of the science and mission engineering teams, it is unlikely that so formalized a convention would have developed in time for the mission. We suggest that our involvement enabled the mission to shape and formalize the work, the naming convention and the ontology and integrate them into software and mission tools within the mission development timeline, while creating information flows that fed across the mission planning tools.

III. THE CASE STUDY

In this section are three major subsections:

- Mars Exploration Rover (MER) Mission: Work System for the Robotic Exploration of Mars.
- Pre-Mission Field Tests and Operation Readiness Tests.
- 2004 Mars Exploration Rover Mission: Work Practice, Naming and the use of the Ontology.

MARS EXPLORATION ROVER (MER) MISSION: WORK SYSTEMS FOR THE TELE-ROBOTIC EXPLORATION OF MARS

This section covers the work of the mission during surface operations on Mars, including preliminary planning, planning tools, science plan integration, and the associated planning software in the ground data system. It also describes the MER mission 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* .The name of an observation in the science activity plan provides coherence across all of these planning phases. In turn it frames the work across the teams and influences the development of ontology for action.

MER Planning Process

Each operational day, called a “sol” (a Martian day, which is approximately 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 thirteen 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
		Hazard	

	Hazcam	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 together 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 independent 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 a naming convention and 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 requesting rover action to generate 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 a rover arm instrument) might be named Post-MB_MI_5position_ElCapitan, describing 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 science and engineering 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 mission team uses a series of software tools that translates the requests into different representations. At the front end of the process, 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 are prominent objects within the landscape (such as rocks and craters) and have unique names that do not change over time. Targets specify a particular location or spot on an already-named feature. They 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.

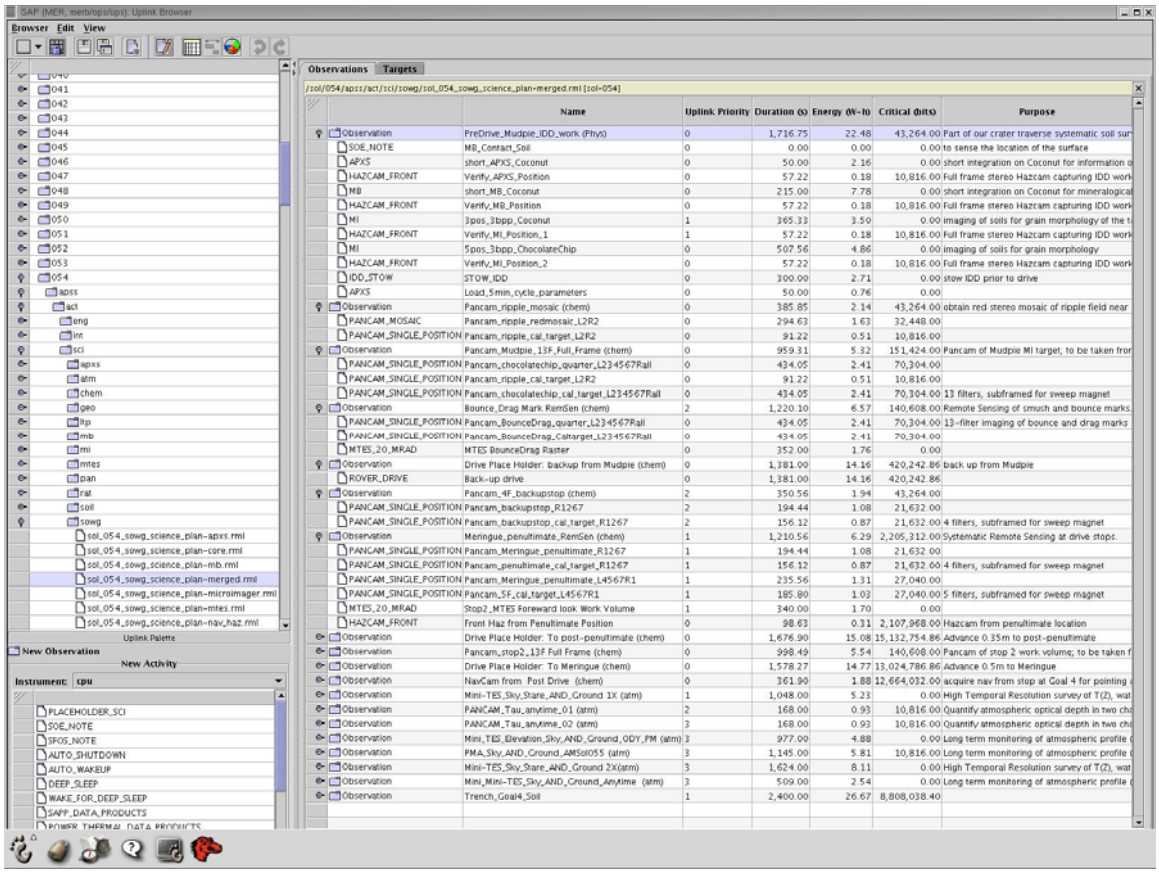

Figure 1. Science Activity Planner (SAP) Tool: screen shot shows a science activity plan as created by the science team with the higher order observations. Open toggles on some observations show subordinate activities that instantiate the observation.

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 request 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 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 software that uses 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.

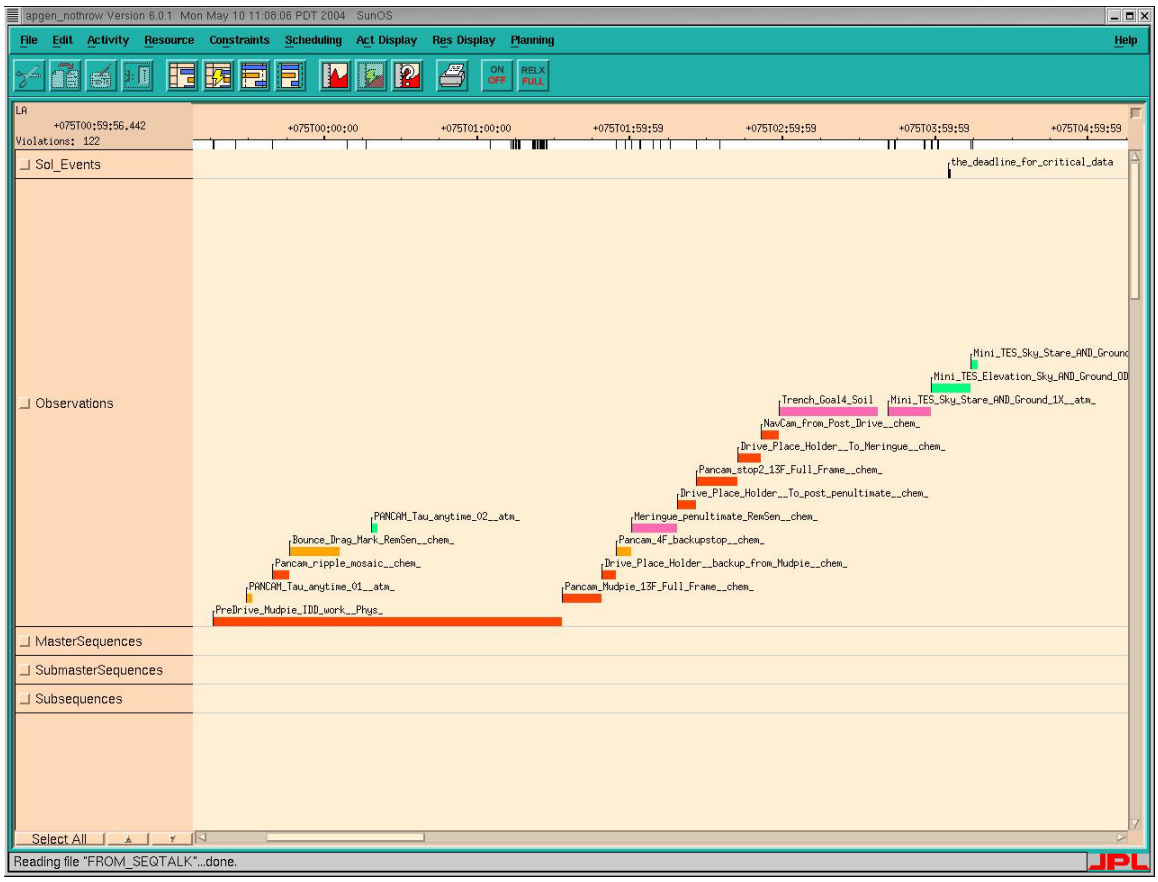

Figure 3. Activity Plan Generator (APGEN): The screen shot of a version of the interface that shows observations as they are "planned" into a timeline for

execution. Colors indicate the priority level that scientists have assigned to each observation.

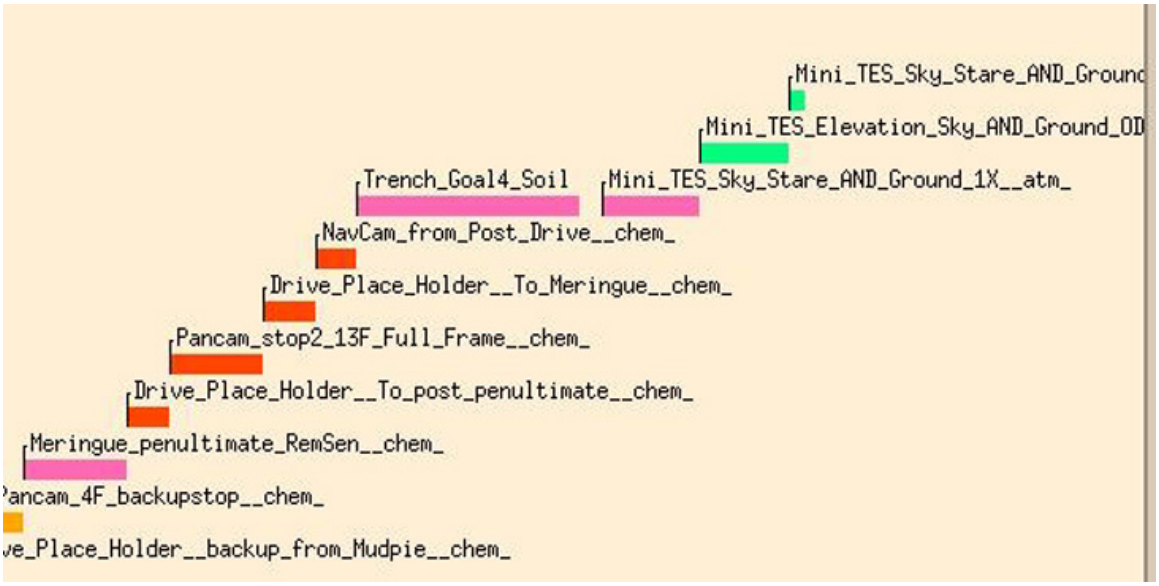

Figure 4. Magnified view of APGEN.

After generating the scheduled plan, the combined engineering and science 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 Uplink 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.

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 returned 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 to identify 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.

Summary of Work System Implications for Naming, Ontology Use and an Ontology of Work

Individual and small teams of scientists use observation names in developing their requests. The whole science team 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 the use of 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.

PRE-MISSION FIELD TESTS AND OPERATION READINESS TESTS

We describe below two field tests in 2001 and 2002, as well as pre-mission tests conducted in 2003. This section describes the initial naming attempts by the science team beginning three years before the mission, how the work system

evolved over time, the development of an ontology with syntax and semantics, and our role in this process.

During this period we determined that we needed a naming convention that could work at a natural language level during the collaborative process and then transition into more precise identifiers in the software as science requests moved from science to engineering teams for eventual translation into commands for the rover. We started with an 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 a flexible syntax, semantics and a description of the relationships between the categories of identifiers.

With each successive test, working with the science team and other mission members, we made a series of recommendations that led to the 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.

2001 Mars Yard Tests

During the spring 2001 tests, the rover was located at JPL in a nearby sand and rock filled area called the "Mars Yard," while the scientists and engineers sent commands from a near-by windowless room. The two Mars yard tests lasted two days each. During the Mars Yard tests as well as the later FIDO (Field Integrated Design and Operations) rover 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. The participating engineers commanded the rover and processed the retuned data for the science team.

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. Table 2. presents the various attempts. Initially, each group chose a different naming convention for targets, a point of scientific interest on an image in the

Table 2

Identification of Early Name Issues

Event	Ontology Issue	Example of Names Used	Originating Team	Consequence & Result
2001 Mars Yard	Observations Require Names	Geometric Rock, White Rock	Science	Meaningful, but multiple schemes developed; lacked systematicity—Abandoned
		Alpha, Beta, Gamma rock numbers	Engineering	Carries an implicit order, but difficult to remember—Abandoned
		Sol#_science theme group_target#	Science	Required renaming if different groups identified the same target; names must be unique— Everything renamed during group meeting
2001 FIDO Rover Test	Observations require Unique Names	Aaron, Ruth, Mantle (baseball theme names)	Science	Difficulty tracking and renaming. Objects must retain names when possible New names assigned to same target for new sol or for different targets on same feature within a sol. Related observations require related names; Some observations are related due to a common geological feature Names used in informal conversation included instrument; Formal target names acquired implicit instrument meaning; observation names must identify both objects and activities Names used to classify pattern of findings Formal target names acquired implicit geological meaning

science planning software. The scientists at first named targets according to the character or quality of the objects in the terrain, while engineers' numeric schemes carried an implicit order between names, guaranteeing uniqueness but reducing meaning. When the scientists gathered as a group, they realized the variety of naming schemas that the different groups were using. In real time during that group meeting, they decided to change all of the existing target names in the software, giving them a coherent scheme of Greek letter names. Group leads then 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 [what we called] G3?" After experiencing several instances of confusion over name choices, the group decided on a unified naming convention that they would use to name targets in the theme groups, during the upcoming FIDO 2001 test:

Sol #_ science theme group shorthand_target #

An example would be:
sol4geo5

It indicates the fifth target chosen by the geology group in sol 4. Should two groups chose the same target in the theme group meetings, the new procedure called for renaming targets during the SOWG meeting with an official, unique "SOWG target name". The SOWG was the last process meeting before handing off to the engineering team for sequencing,

2001 FIDO Rover Test

The FIDO rover 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.). Any observation naming confusions in the

early tests, where it was relatively easy to remember what happened across sols, would surely be magnified over the longer duration and larger scale of the MER mission.

Based on the Mars Yard experience, initial planning and naming proceeded within the theme groups using informal target names. Targets acquired their official name upon acceptance as part of the final science activity plan during the SOWG meeting. As shown in Table 2, the SOWG chair chose an official SOWG target names from a designated list, typically single words that represented a theme of names to be used in the sol, such as baseball players. Scientists kept individual, precise notes or spread sheets to link the official target name to the original theme group target name, which was over-written and disappeared from the software. Finally, engineers added numbers for rover-centric site and position identifiers as well as the sol number, to help the engineers track the assumed context for a request, but they had little meaning for the science team.

The teams used a variety of names in face to face conversation. Theme group members, both in small and large groups, referred to original target names such as “sol7 geo6” or just “geo 6”. After a science request received 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?” Interestingly, during face to face interactions all test participants often used instrument as an identifier, e.g., “Do we want the Pancam on Citation?” but the instrument designator was not entered into the software as part of the name. Scientists simply entered the target name “Citation.”

The 2001 FIDO Rover Test provided the authors with a foundation for making a number of recommendations that would be exercised in the 2002 field test. 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, analysis of the 2001 field test processes suggested several work

practice issues related to naming. 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.

2001 Naming Issues

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 convention that allowed a target to keep its official name from sol to sol. Targets received new official names when a new request was specified. However, 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 work and resulting 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. We concluded:

- *Once named, Target/Objects must keep that name to minimize confusion when future or related work is done on the same object. However, the request name must identify the different, albeit related, instantiations of work on the same target.* Consistent names support developing understanding and make reference to objects useful and meaningful.

A target "name" came to represent more than the targeted point where scientific

work had been done.

Participants used the single "target" name to identify several referents: 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. 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. The target name came to refer to the results of previous instrument work as well as a developing body of scientific knowledge. Renaming targets blocked scientists' ability to follow a developing trail of scientific understanding and to make comparisons. However, the retention of 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." We concluded:

- *A naming convention must identify both the **objects** that are located in the terrain **and** the **activity** that is being requested on those objects.* The concept of *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.

2002 FIDO Rover Test

In preparation for the second major field test, FIDO 2002, we created the taxonomy in Table 3, provided definitions for the constructs and began to establish relationships between the parts of the taxonomy. We suggested a

convention for the 2002 field test, allowing participants to identify *features* (unlike FIDO 2001), associate targets with features (providing the ability to differentiate between features and targets), and create *observations* (that grouped individual activities with a common purpose). All of these required modifications to the

Table 3

Progression of Naming Issues Following Systematic Analysis & the Authors' Recommendations

Event	Ontology Issue	Definition	Example	Consequence & Result
2002 FIDO Rover Test	Features contain targets	Features are objects of interest in the terrain (cliff face, crater, hill) and do not require association with actions. Targets exist only as parts of identified features, and serve as focal points for instrument activity. A single feature may contain several targets.	Feature = shoe, Target = heel	Scientists tracked distant features as directional reference points, indicated in both software and large panorama images. Reference to features in talk subsumed individual targets. 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. With separate names for features, 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
	Observations contain activities	Observations are higher level containers for individual activities directed towards a scientific purpose. Activities only exist as parts of observations.		
	Activities have types	Activities are actions of the instruments and rover state as groups of rover commands. Activity types are based on instruments and are enumerated in an activity dictionary	Single frame Pancam	
	Activities have targets	Activities require a target		Time limits encouraged scientists to enumerate target names, based on feature names, e.g., RIO_1, RIO_2 resulting in less distinctive names that required coordination during assignment to avoid repetitions. If named after instruments, e.g., Min_1, Min_2 the same target name referred to different targets across sols.

Table 3

Authors' Recommendations (cont'd.)

Event	Ontology Issue	Definition	Example(s)	Consequence/Result
2002 FIDO Rover Test (contd.)	Observations have types	Observation names reflect different basic relationships between instrument and feature, e.g., survey	Survey, e.g., Sky_survey _IPS_2ELl_ sol 11 (a spectral survey of 2 distinct elevations in the sky on sol 11)	Time limits affected the inclusion of instrument in software names, but talk continued to reflect instrument names. Observation types evolved into methods, e.g., Tau.
	Activities have targets	Activities require a target		Time limits encouraged scientists to enumerate target names, based on feature names, e.g., RIO_1, RIO_2 resulting in less distinctive names that required coordination during assignment to avoid repetitions. If named after instruments, e.g., Min_1, Min_2 the same target name referred to different targets across sols.
	Observations have types	Observation names reflect different basic relationships between instrument and feature, e.g., survey	Survey, e.g., Sky_survey _IPS_2ELl_ sol 11 (a spectral survey of 2 distinct elevations in the sky on sol 11)	Time limits affected the inclusion of instrument in software names, but talk continued to reflect instrument names. Observation types evolved into methods, e.g., Tau.

Table 3

Authors' Recommendations (cont'd.)

Event	Ontology Issue	Definition	Example(s)	Consequence/Result
2003 Pre-Mission Tests & Trainings	Instruments have classes	Classes reflect fundamental differences in instrument execution.	Remote Sensing, In Situ	Observations could not contain activities that cross classes. In-situ observations should not include more than one feature.
	Activities require unique identifiers	Activity identifiers distinguish between multiple occurrences in same sol.	Pancam_mosaic_stereo_2x2	Each instance of an activity type had its own explicit semantics.
	Observation types reflect methods	Methods correspond to stereotypical combinations of instruments, or special ways of using an instrument.	Comparison, movie, quick look, blind (for MiniTES activity without a supporting image)	Provided semantics for sets of activities. The identification of method names became apparent in these training opportunities.
	Other identifiers and constraints	An open invitation to highlight unanticipated descriptors	Temporal: Afternoon Spatial: North, Long, Around	Provided additional opportunity to provide explicit semantics.
	FOV naming	Requests based on azimuth and elevation	Azimuth 30 degrees	Provided alternative approach to targeting.

FIDO 2001 software, which up to that time had only allowed scientists to indicate targets and to choose from a set of rover activities. We recommended that:

- Target names should reflect if possible, a whole-part relationship with a 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 multiple targets and a common feature.
- 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 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.

As we observed and analyzed the work of the science team in the second test, we understood that our recommendations had changed the developing work practice of the scientists with regard to the use of features in the environment, reference to rover instrumentation and the development of systematic methods for tele-robotic exploration. These appear in Table 3 under Consequences & Results. After the 2002 test, we also recommended that the software include a “location” identifier that would change after a yet to be determined length of rover traverse and a designator of a larger “area” to help scientists keep track of developing data sets. These recommendations were not included in the software.

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 parallel development of the engineering processes. We made recommendations to the science planning software and other tools that called for a hierarchy of observations with subordinate activities (see Figure 1 and 2.) Over this test cycle we expanded and refined the set of identifiers that would together become the ontology for working with the activity plan and for naming observations and activities.

We identified a convention that differentiated between the use of single and multiple instruments in a single observation, and expanded the official taxonomy and the relationships between observation, activity, feature and target. *In-situ* observations, those 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, when the distance to an object and the accuracy of both correctly identifying and carrying out an activity on an object 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."

The cognitive differences in these types of observations, the 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 4.shows the relationships between single and multiple instruments use and instrument name identifiers.

Table 4. 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)

Initially, the mission design called for the scientists to simply choose activity types from a dictionary in the science planning software and populate their observation with formalized sets of information that could be used by downstream teams in planning and commanding.[1] Scientists could choose the activity type (such as pancam_single position) from a menu in the software. However, as the uplink process was tested in practice, the Payload Uplink Leads (PULs), who expanded the activities into sequences, needed to distinguish between the different activities that appeared in their downstream software tool. For instance, a Pancam PUL might see multiple Pancam_mosaic_stereo activities. We recommended that scientists give unique activity names to the activity types in the software, indicating important distinguishing parameters, such as filter color or the camera mosaic size (2x2), making the PULs' jobs easier. We had previously recommended placing the target name at the

[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.

observation level, but with the development of activity names, we recommended the activity level as appropriate to identify the spot for instrument action. In addition, adding the target to the activity name made the name more unique.

- The syntax for naming an activity became

Distinguishing parameter_Target.

Table 5: 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 6: 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

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

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 of which are indicated in Table 7.

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 scientific methods. Finding a syntax and structure that would frame the problem while accommodating change was a challenge. The missing piece in the convention was an open invitation to enter what we came to call “other identifiers”, typically temporal or spatial constraints and pointing (30 degrees).

Table 7. 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

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 from the 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.

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 45-degree images to cover a slightly less than 90-degree FOV at an azimuth of 30-degrees (because of the two stereo images overlapping). Or the scientist could ask for navcam image centered on and pointing toward a target on a feature in the distance.

By the end of the test period, the ontology was in place but as work practice researchers, we knew that the experts in the domain would learn to do their work as efficiently as possible. They would use the ontology and our procedures only to the degree that they supported the work. We expected there would be variations to the ontology in actual use, and the work might call for new parts to the ontology.

2004 MARS EXPLORATION ROVER MISSION: WORK PRACTICE, NAMING AND THE USE OF THE ONTOLOGY

This section describes some of our findings from a qualitative analysis of the science plan data and the work practice during the mission that began in January of 2004. It describes the use of constraints, the development of method, observations and activities, and the use of feature and targets within the actual mission context of use.

Work Practice and Naming

The basic ontology held through the mission, but some of the elements continued to develop. Scientists adapted the naming convention as they gained experience with the operational environment and developed new methods of tele-robotic exploration. The science team adapted the syntax to the context. Our formal

ontology called for framing the name so that feature always identified an end to the name. As time went on, however, they began to situate other information, not feature, at the end of the name. We suggest that this may have been a way to make important information more obvious to each other and downstream teams.

Specific observations, examples, and deviations appear in Table 8 below and will be discussed below in the order of their appearance in the table.

Table 8. 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

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 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, brushing), 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. The availability of a downlink time (data relayed from the Odyssey Orbiter) was identified (pre-ODY) so that data could arrive before the next cycle of planning. Spirit's rock-laden landscape broke up a traverse into small steps (penultimate and ultimate), and scientists identified observations with each of these steps. These numerous examples confirmed our expectations that other identifiers, mostly as constraints, were needed to support unanticipated needs.

As the mission progressed, scientists began to rearrange 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'.

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 sols of operations. The ability to name clusters of activities with a single label lends support to the idea that observations were containers that rendered the work coherent and that the ontology helped frame the tele-robotic work of the science team. As we saw 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.

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 sometimes 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 (as undifferentiated and ubiquitous) which we had seen in field tests, and introducing others, such as soil and 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.

Activity Name Development

Table 9 indicates some of the additions to activity names that appeared during the mission. While activity names included temporal constraints, it was not with the regularity we observed 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 rover itself did not “know” about pointing and since there was no way to know the exact set of the rover coordinates after the rover had moved without analyzing returned telemetry data, the most accurate pointing was done using azimuth and elevation coordinates 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 9. 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. Over time, numbers are not unique, but within the context of a particular static situation, they can seem an acceptable practice and so team members sometimes used them. 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 two 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.

Observation Name Development

As the mission went on, observation names got longer. We believe this tendency correlated with the increased use and 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 this was a request for new target 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 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.

[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.

Variations such as this highlight the importance of agreed upon conceptualizations and the need for specification as information flows through different software applications.

IV. CONCLUSIONS

An Ontology of Work

The need to share structured information (Musen, 1992, Gruber 1993) that represent agreements about *shared* conceptualizations (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 behind our research and ontology development.

As the research progressed, however, we found that as we identified and formalized shared conceptualizations, we were also structuring and organizing not only the work of the rover on Mars but also the planning work of the mission scientists and engineers. We moved beyond the elicitation of knowledge and the definition of the internal relationships of objects to a concern for how the ontology would structure work and behave in practice.

Compatible with Clancey (1993), we focused on observed behavior in the environment, but found that we could not only draw on that behavior at a conceptual level but also feed back into it at a conceptual level, influencing the work that was being done. Observations and activities represent the basic constructs of the work of tele-robotic science. 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 the emergent *ontology of work* (Figures 5 and Figure 6).

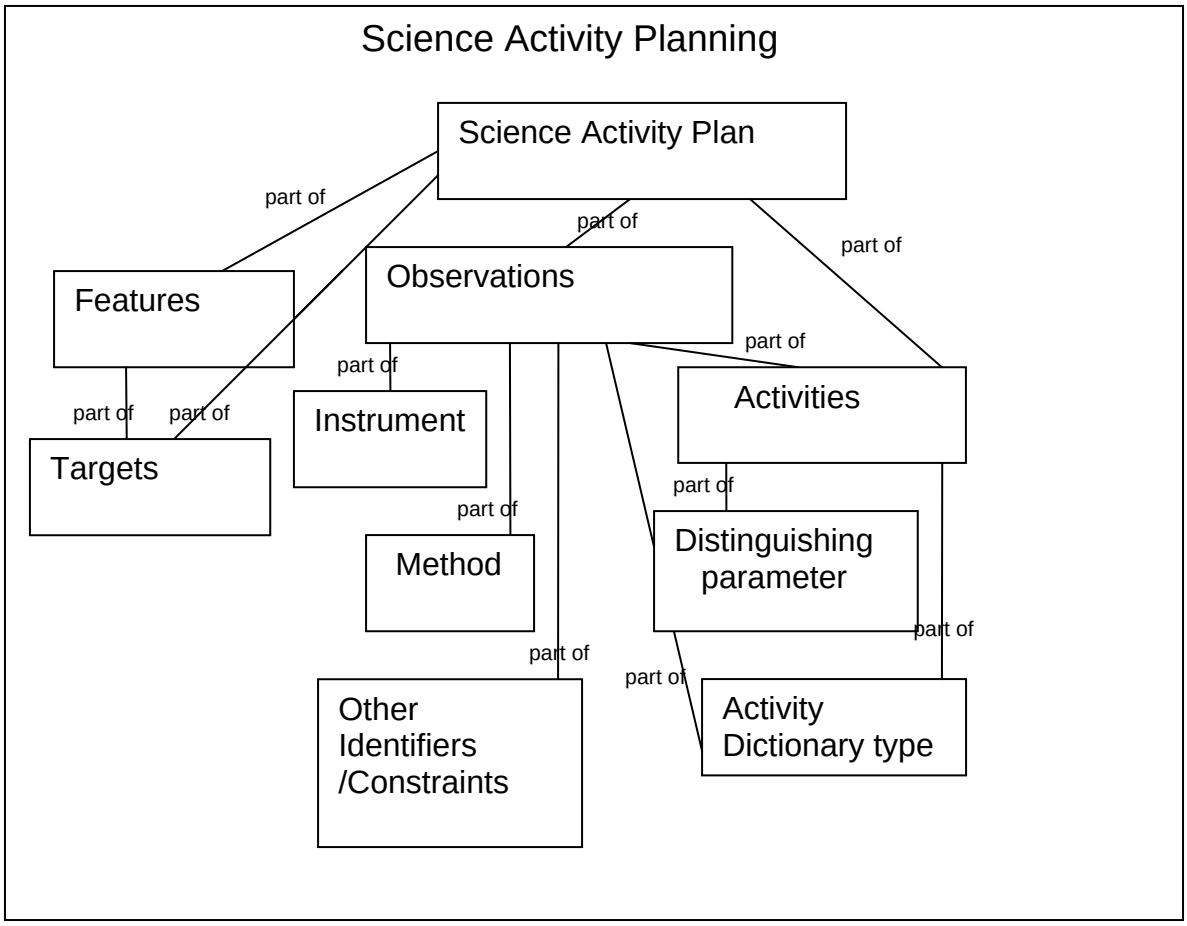

Figure 5. The parts of the science activity plan also serve as way to structure the work of the rover, frame the requests of tele-robotic science and move requests through the work system.

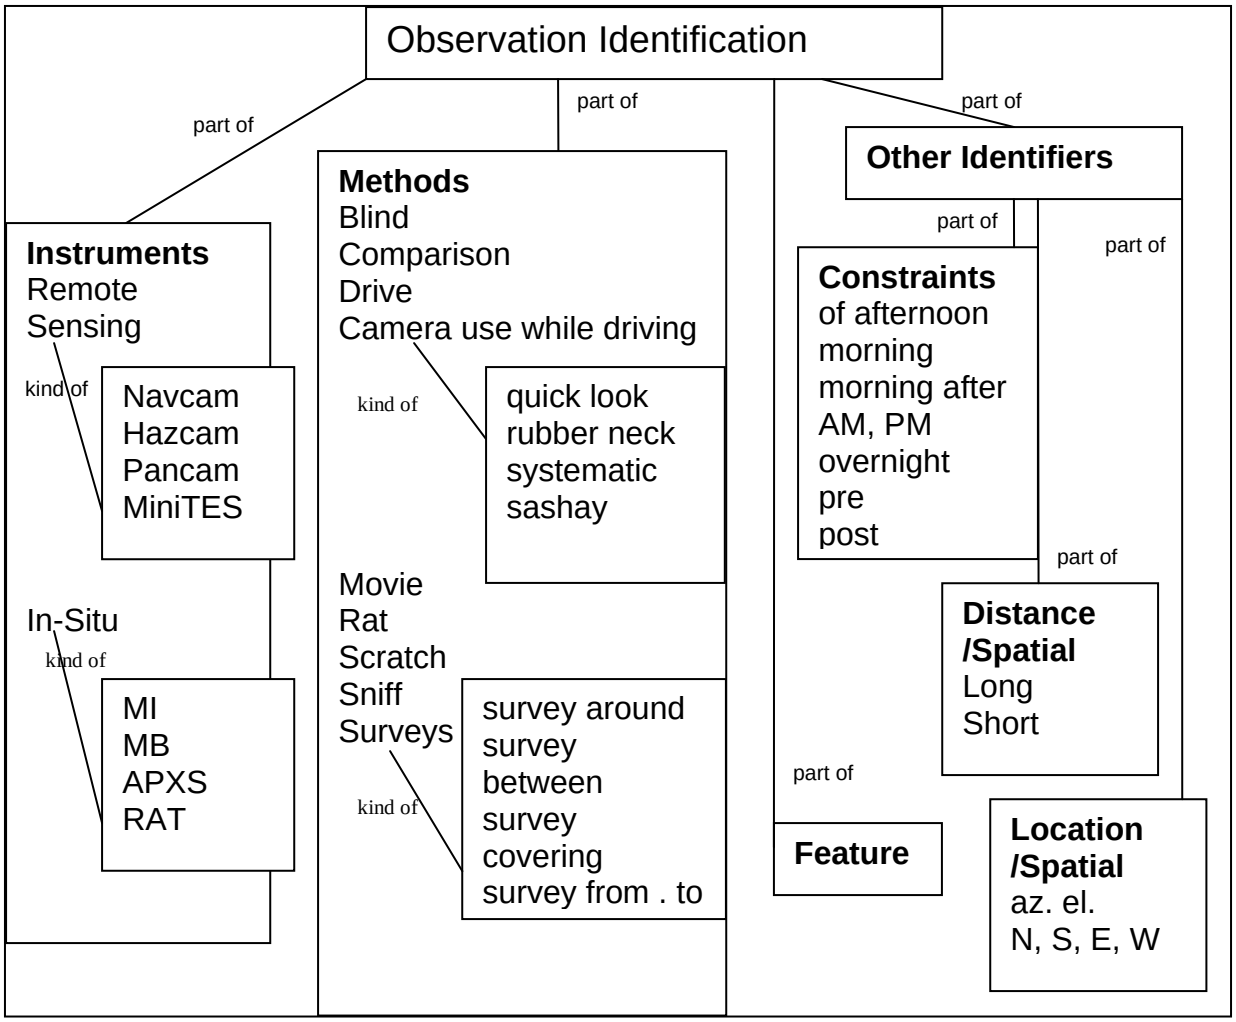

Figure 6: Requests are changed from natural language discussions into observations that can be used for decision making and planning as they move through the planning system, and then honed into more explicit identifiers used to command the rover.

The MER ontology of work and the associated abbreviated names meets the definitions for an ontology referred to by Gruber (1994), and Guarino and Giaretta (1995). That is, it is not so much a complete *specification* of shared conceptualizations as it is an incomplete or *partial agreement or account* 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 a dynamic external environment. Additionally, those terms must be adjustable by 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 user needs 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 of work 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), *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. 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 in the beginning, 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, the work has

become more streamlined.

In the future, 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.

Flexibility is essential in any expert domain that interacts with a dynamic environment and learning will continue over time. As a result, work process and even tool use can change. MER scientists developed existing classifications, such as method, and created new shared conceptualizations. That work was facilitated by the use of the “other identifier” field. This helped the ontology adapt and support learning and change.

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, such an ontology could support the work of other tele-robotic science, including tele-medicine.

Implications for Research in Ontology Development and Knowledge Elicitation

Past ethnographic research in acquiring expert knowledge has focused on the

elicitation of knowledge from existing situations of use (Meyer, 1992, Forsythe et al, 1989). To this discussion we offer the finding that as we captured and refined the ontology for tele-robotic science, we also framed the work that was being done, creating an *ontology of work*. The science team work practice was to first define observations they wanted to do and then instantiate them with activities and methods that would complete that work. Teams and meetings were dedicated to making these particular decisions within the larger work of creating a science activity plan for each Martian sol. We suggest that other researchers may find clues to understanding a domain and to defining the appropriate ontology by looking at the structure of the work practice that exists in the domain

Drawing from the Case Study and the above analysis, we also identify six findings that we believe have implications beyond the current research and make a contribution to the effort of knowledge elicitation and ontology development in expert domains.

- An ontology can structure and organize the **work** o**f** the domain (an ontology of work), as well as identify, name and give semantic relationships to the **knowledge** within the domain.
- Ontologies of work in information systems must contain both **actions** and **objects** to identify and represent all aspects of the work involved. The names must represent the basic units of the action in the work system (e. g. method, instrument) in relation to the objects (e. g. features and targets) on which the work will be done.
- Abbreviated names can be successful identifiers in an ontology of work as long as they contain consistent and systematic representations of the work to be done and draw on pre-identified parts of the ontology to create descriptions of the work being done.
- Both the ontology and the abbreviated names must be capable of representing and incorporating change and the learning that will take place in an expert domain as it responds to new needs.

- Abbreviated names allow for natural language referencing and collaboration in a domain.
- Abbreviated names and an ontology of work allow for the transfer of knowledge and information across disciplines and tools within a work system.

V. REFERENCES

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.

Clancey, W. J. (1993) "The Knowledge Level Reinterpreted: Modeling Socio-technical Systems*", International Journal of Intelligent Systems*, 8, pp. 33-49.

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

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.

Forsythe, D.E and B. G. Buchanan (1989) "Knowledge Acquisition for Expert Systems: Some Pitfalls and Suggestions", *IEEE Transactions on Systems, Man and Cybernetics* 19(3), pp. 435-442

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

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

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 Uschold, M & M. Gruninger (1996) "Ontologies: Principles, Methods and Applications", *The Knowledge Engineering Review,* 11, pp. 93-136.

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. (1995) and A. Henderson (1995) “Interaction Analysis: Foundations and Practice”, *The Journal of the Learning Sciences* 4(1) pp. 39-103.

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.

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.

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.

Meyer, M. A. (1992) “How to Apply the anthropological Technique of Participant Observation to Knowledge Acquisition for Expert Systems*”, IEEE Transactions on Systems, Man and Cybernetics* 22 (5) pp.983-991.

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.

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

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

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

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* 17(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 senior human-centered computing research scientist with SAIC 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