Optimizing Information Value: Improving Rover Sensor Data Collection

Justin M. Glasgow, Geb Thomas, *Member, IEEE*, Erin Pudenz, Nathalie Cabrol, David Wettergreen, and Peter Coppin

***Abstract*—Robotic exploration is an excellent method for obtaining information about sites too dangerous for people to explore. The operator's understanding of the environment depends on the rover returning useful information. Robotic mission bandwidth is frequently constrained, limiting the amount of information the rover can return. This paper explores the tradeoff between information and bandwidth based on two years of observations during a robotic astrobiology field study. The developed theory begins by analyzing the search task conducted by robot operators. This analysis leads to an information optimization model. Important parameters in the model include the value associated with detecting a target, the probability of locating a target, and the bandwidth required to collect the information from the environment. Optimizing the information return between regions creates an image and provides the necessary information while reducing bandwidth. Application of the model to the analyzed field study results in an optimized image that requires 48.3% less bandwidth to collect. The model also predicts several data collection patterns that could serve as the basis of data collection templates for improving mission effectiveness. The developed optimization model reduces the bandwidth necessary to collect information, thus aiding missions in collecting more information from the environment.**

***Index Terms*—Human computer interaction, images, robot, search.**

I. Introduction

RECENT technological advances in robotics currently allow for the safe deployment of teleoperated robots into extreme environments too hostile for safe human exploration. Example sites of teleoperated robot deployments include the Chernobyl nuclear reactors [1], the World Trade Center towers as an Urban Search and Rescue (USAR) effort [2], and the surface of Mars [3]. In these situations and many other teleoperation scenarios, the robot operator must navigate an unknown environment primarily using information provided by the rover. Mobile robots typically deliver this information as streaming video or still images, collected by one or more mounted cameras. Some robots utilize sonar and laser range finders to provide information useful in avoiding obstacles and determining distances. Even with this extra information, image quality strongly influences whether an operator can safely navigate and achieve the mission goals.

The USAR experience at the World Trade Center clearly demonstrated the importance of image information to mission success. During the USAR operations, robots assisted with the following four different tasks: confined space search, semistructure search, area monitoring, and medical payload transportation [4]. All four tasks involve the use of images; however, the space and semistructure searches are of particular interest, as these have the clearest analogs across multiple robot platforms. These tasks required operators to navigate a robot into areas within the rubble too small or dangerous for a person. Using returned images, the operator attempted to navigate the robot through the structure, working to identify survivors, search for signs of victims, identify structural instability, and avoid damaging the robot. The search-and-rescue effort deployed robots into the World Trade Center in eight situations [4]. Each of the robotic deployments had a number of errors occur during the mission. On each deployment into the structure, the robot became stuck an average of 2.1 times. During the sixth and eighth deployment, Casper and Murphy [4] calculated that operators spent 54% of the actual drop time diagnosing problems with the robot and not performing mission-critical tasks. On one drop, the operators could not solve the problem with the provided imagery and, instead, had to pull the robot out of the structure and start the search anew.

One of the conclusions from the World Trade Center response was that operators needed better information to aid their navigation of the structure. Some observers might suggest adding additional sensors, such as sonar, that aid in obstacle detection. These systems could provide the additional information the operators need to navigate the structure. Others think that images with higher resolution, field of view (FOV), or update rate would have aided navigation and reduced operator stress by reducing situations where an operator detects ambiguous objects that may indicate a victim [4]. While these and other suggestions are valid solutions, implementing them on robotic systems is not always a straightforward task. Robotic systems have a number of constraints, such as size, power, weight, and bandwidth that limit implementation of these suggestions.

Focused engineering efforts should allow the addition of these tools to many platforms; however, an important additional consideration is how to improve and optimize data collection tools. These tools must assist operators when other systems degrade or fail. Optimization of visual information within the

Manuscript received April 20, 2006; revised December 1, 2006. This paper was recommended by Associate Editor R. Hess.

J. M. Glasgow, G. Thomas, and E. Pudenz are with the Department of Mechanical and Industrial Engineering, The University of Iowa, Iowa City, IA 52242-1595 USA (e-mail: geb-thomas@uiowa.edu).

N. Cabrol is with the SETI Institute, Mountain View, CA 94043 USA and also with NASA Ames Research Center, Moffett Field, CA 94035 USA (e-mail: ncabrol@mail.arc.nasa.gov).

D. Wettergreen and P. Coppin are with Carnegie Mellon University, Pittsburgh, PA 15213 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMCA.2008.918614

constraints of robotic systems requires accounting for the associated bandwidth. The objective of this paper is to develop a methodology for controlling bandwidth consumption while increasing the value of information returned by an exploration rover. First, this paper reviews literature relevant to designing robotic visual systems. Second, it reviews findings from a two-year field experiment that applies visual systems to the study of geology and biology in desert environments [5]–[7]. These findings lead to the development of a theory that optimizes data collection. This paper concludes by validating the model and proposing ideas on how to implement the theory on future robotic missions.

II. Background

A. Image Collection Constraints

Visual information plays a number of roles for operators during a robotic exploration mission. Often, the critical role for operators is successful navigation through the remote environment. If the operator cannot safely navigate the robot, it is impossible to achieve any other mission goals. Images provide operators with the ability to avoid obstacles, localize landmarks, and detect targets. The domain of image collection faces a number of constraints to balance when developing visual imagery collection systems.

In an experiment looking at how drivers navigate vehicles in indirect viewing situations, van Erp and Padmos [8] identified the following six parameters that affect operator control: FOV, magnification, camera viewpoint, presence of reference or orientation viewpoints, image quality (contrast and resolution), and frame rate. For robotic systems that operate under fixed and limited bandwidth conditions, the tradeoff becomes a matter of balancing bits per pixel, pixels per image, and images per second [9]. Choosing the correct balance depends on the mission and is often a decision made early in the process by the robot designers with little input by potential operators.

In a paper focused on developing an onboard vision system for an unmanned aerial vehicle (UAV), Garcia *et al.* [10] use a commercially available FireWire camera to capture images as the UAV flies over terrain. Their primary focus is to create a system that detects large objects while minimizing battery power and weight. The authors focus on traditional design constraints but do briefly consider the role of resolution and FOV in meeting the mission requirements.

Several other papers have documented the optical parameters for mobile-robot cameras. The Pandora robot for urban military actions uses a stereo pair of forward-looking charge-coupled device (CCD) chips [11]. Similarly, the designers of the Rocky 7 Mars prototype at the Jet Propulsion Laboratory used an off-the-shelf CCD camera to collect images [12]. Another paper, developing the Griffon prototype, utilized a 320 $\times$ 240 pixel color CCD camera [13]. These papers documenting the development of the camera systems for Mars Exploration Rover (MER) are good examples of the careful attention that designers can put into developing a camera [14]. These papers all develop design requirements and use them to develop useful cameras; however, all of them do not consider a design requirement essential to a successful mission.

The constraint these other papers do not consider is information flow. Operator success depends on the amount and update rate of the information received from the robot. By considering the information flow, designers must make choices between resolution and FOV. Consider a commercially available five-megapixel digital camera. A single JPEG compressed image is 4.25 MB. File sizes this large quickly overwhelm system bandwidths, and this limits the information a rover can capture and return to the operator. Special lenses that capture a larger FOV, such as fisheye lenses, can potentially capture all the information needed by the operator with a single image. Large FOV lenses have the tradeoff that they reduce resolution (in terms of degrees per pixel) to a point insufficient for operator needs. In systems that maintain resolution and provide a large FOV, the camera captures multiple tiles and uses these to create a mosaic of the region. This solution requires capturing a number of images, risking an overload of the available bandwidth. No single solution provides an optimal balance across all the optical parameters. Instead, camera system designers must understand how operators view and use images during the mission to determine the optimal tradeoff among the optical parameters.

B. Human Vision

It is also important to consider how the human visual system works and how people view and extract information from images. Only with this understanding is it possible to create a robotic visual system that matches or outperforms the human visual system. An important concept to consider is how the brain selectively attends to important cues. Without this capability, the incoming information would easily overload the brain, in a manner similar to how returned images can overload the bandwidth of many robotic systems.

A breakdown of the human visual system shows that the high-resolution area of the eye, the fovea, only covers a 2° spot [15]. It is only through image integration and multiple scans that we perceive our surroundings at a high level of resolution. Even with the limitation in what a person can see at full resolution and process in working memory capacity, visual information acquisition is highly efficient in humans [16]. This efficiency is, in part, due to the ability to parallel process information, attend to salient objects, and use schemas to identify areas of interest. Essentially, people pay attention to objects that stand out in the scene or fit spatial schemes they have for target placement. Background knowledge and past experience is a significant contributor to how people collect and interpret information from the environment. Understanding the role of these schemas and experiences aids in improving the human–robot interaction.

One notable viewing pattern is the tendency to examine the borders of images less. Dubbed the "edge effect," Parasuraman [17] showed that a natural viewing pattern focuses on the center of image. A system that increases the saliency of objects along the border will help ensure that operators can identify targets of interest [18]. A partial explanation for the edge effect is people's tendency to pay attention to areas possessing high levels of visual detail [19]. When looking at a portrait, the tendency is to pay attention to the details in the face before looking at the typically less descript background. Observers choose to look at the areas containing the most information as a timesaving schema. By emphasizing or deemphasizing the detail of certain areas, a designer can help direct fixations to points of interest.

Saliency effects, such as the edge effect, show the importance on understanding how operators view images to extract information. In robotic exploration, the operators are using images to detect targets. Unfortunately, target search is a poorly structured process and is difficult to model [18]. The difficulty arises from the fact that search is primarily a cognitive task. The searcher operates off a mental set of target location probability maps about the area based on experiences in similar situations. Depending on the level of experience, the accuracy of such probability maps highly varies. However, an understanding of the searcher's probability maps can aid in designing camera systems that take advantage of these probability maps. The dependence on probability maps suggests that there is no standard search pattern. Some people may search from left to right as in reading a book; others may bounce from one high-probability area to another [18]. Yarbus [19] showed that people use different scan paths over the same picture when asked to look for different targets.

Despite the variability inherent to operator search, robot designers do have some control and, possibly, obligation to assist operators in quickly determining the presence or absence of interesting targets. Morawski *et al.* [20], [21] have developed models of the optimal search times associated with random and systematic visual inspection for a single target. These models determine an optimal stopping time for a search based on the values and costs associated with the search. When evaluated using the same parameters, the models suggest that systematic search outperforms random search, and designers should ensure that interfaces support a systematic search of the environment. Designers should also plan on training searchers on the interface and its optimal use to improve performance [22]. These models are limited in their application to robotics as Morawski *et al.* [20], [21] developed the models to only account for the identification of a single target, which is unlikely in robotic search. Hong [23] expanded the systematic-search model to account for searches that require the identification of an unknown number of targets.

Human vision and image information extraction is a complicated and difficult-to-model process. Designers can utilize saliency effects and search models to improve image acquisition, but often, overall performance depends on the skill of the operators. If the operator does not have appropriate experience, they will not efficiently search images and often miss targets of interest. In addition to training operators to improve information extraction, designers can develop robots with autonomous abilities to perform some image-search tasks, alleviating some of the operator burden.

C. Automation

Increasing robot autonomy allows the robot to complete tasks typically relegated to the operator. Increased autonomy can serve two purposes. First, it reduces the number of targets the operator has to consider, thus reducing overall cognitive effort required of the operator. Second, onboard processing of images can reduce the amount of information needed by the operator, thus reducing the bandwidth of information returned to the operator. Increased autonomy algorithms usually exist to serve one of two purposes, either target search or obstacle avoidance.

Basic automated search algorithms have existed since 1996 [24]. These early algorithms modeled animal search and forage patterns to improve robotic search [24]. The concern with algorithms for fully autonomous search is that they require some general knowledge about the target. Many clever methods, such as using reference images and identification based on image parameters [25], result in useful search algorithms that have use in structured environments. They, however, have little applicability in uncontrolled situations common in robotic exploration.

More useful to robotic exploration are collision avoidance algorithms that aid the operator in safely navigating the robot [26]–[30]. Everett *et al.* [28] developed a "telereflexive teleoperation" scheme that alters operator commands based on readings from the robot sensors. Implementation of some of these collision-avoidance algorithms may help reduce the bandwidth needed by the operator; however, this does risk reducing the operator's situation awareness.

The need for operators to understand the situation they observe is paramount to successful missions. Operators often compare driving a mobile robot to the experience of trying to navigate by looking through a straw [31]. The ability to view objects within the context of the environment helps the operator determine if a particular object is a target of interest. A full awareness of the situation is necessary for error-free operation [32]–[36]. Situation awareness is a three-stage process covering correct perception of the environment, comprehension of the interaction between individual parts, and projection of decisions into the future [34]. In a robotic-search task, this means that the operator cannot only perceive potential targets but must also understand the importance of the target within the greater context of the environment. Only then can the operator make appropriate decisions that satisfy the mission goals. Interface designers often measure an operator's overall situation awareness to determine the efficiency of an interface, which is a metric that designers of camera systems must consider [37], [38].

D. Summary

Developing camera systems to aid operators in completing mission goals is a complicated process. The system must balance the tradeoffs between the different optical parameters of resolution, frame rate, and FOV. There is no way to prescribe a single solution for every robot mission; instead, each individual mission has a combination of these parameters that works best. To understand which tradeoff between the parameters works best, it helps to understand how operators search images for mission-specific targets. Human search is an unstructured process that is highly dependent upon cognitive probabilistic maps and saliency effects. Even with its unstructured nature, it is possible to model target search and understand the optimal time needed by an operator to search an image and detect targets.

It is possible to remove some of the burden from the operator, and thus reduce the amount of information to return, by automating some of the operator tasks. Automated search algorithms and obstacle-avoidance algorithms can improve mission performance but, sometimes, at a cost to operator's situation awareness. If an operator is still ultimately in charge of the robot, then the operator must receive the appropriate information to understand what is happening in the remote environment.

Understanding general characteristics of human search and operator needs is important to developing an efficient information acquisition system. However, the system must also consider mission-specific parameters in determining what information to collect from the environment. This paper proceeds by reviewing findings from a two-year case study of an astrobiology field test. The findings from this case study build the framework for a quantitative methodology that is useful in designing camera systems that minimize bandwidth needs while maintaining information quality.

III. Life in the Atacama (LITA) Field Experiment

A. Background

The 2004 and 2005 LITA field experiment constitute a joint effort between Carnegie Mellon University and the NASA Ames Research Center to develop and test an autonomous rover capable of detecting microorganisms and organic compounds. The general objective is to characterize the habitats and distribution of microbial life in the Atacama Desert in Chile, an extremely arid desert that supports little life. The rover, Zoë, executed daily plans that are uploaded by a team of scientists (geologists and biologists) that sent it traversing up to 10 km a day across the desert. Along the way, the rover would stop and use its instruments and cameras to collect information from the environment. The 2004 field season consisted of two separate site investigations, in September and October of 2004, referred to as sites B and C, respectively. The 2005 field season covered three sites, D, E, and F, explored during September and October of 2005. Previous publications cover the full details of the mission design, objectives, and data collection [5]–[7]; this paper addresses the use of scientific information and reviews findings that are essential to the understanding of the scientists' behavior.

A table reviewing all the instruments in Zoë's payload appears in [5]; of these, two specific instruments constituted 94% of the bandwidth used in returning data during the investigation. The science team had up to 150 MB of bandwidth available each day. The primary visual information source in this investigation was the stereoscopic panoramic imager panoramas. The panoramas provide a 360° view of the environment from the ground in front of the rover up to the sky. Collection of the panorama consists of capturing a number of high-resolution tiles and combining them into a mosaic. As the name suggests, this camera system design collects tiles in triplicate to create stereo images. Based on analysis of science team usage during the 2004 investigation, the 2005 investigation only returned a single tile. It is the bandwidth for the collection and return of a single tile that the following analyses utilize.

The other information source is the fluorescence imager (FI) [39]. This instrument images a 10 cm $\times$ 10 cm square underneath the rover. The imager sprays a mixture of four dyes, onto the area under the rover, that react with proteins, amino acids, lipids, and carbohydrates to cause fluorescence. The camera collects images of this fluoresce, thus providing the science team with the primary method of detecting signs of life.

B. Previous Findings

Much of the previous analysis of the LITA investigation has focused on how the science team views the panoramas. Of interest is the fact that the panorama represents on average 70% of the utilized bandwidth [5]. However, because of its large file size, the scientists do not view the panorama at full size. During the 2004 season, the science team viewed the panorama initially at 22% of its full resolution and, then, click on areas of interest to view an individual tile at full resolution. Part of this team's observations included maintaining an access log of which tiles the science team viewed at full resolution. Final analysis showed that the science team viewed 52% of the available tiles at full resolution [5]. This raised the question of whether there were any systematic patterns as to how the science team viewed tiles in the panorama.

To determine these viewing patterns, the process examined the search task the scientists go through. This led to the hypothesis that the science team uses the panorama to safely navigate Zoë through the environment, determine its position within the environment, and look for large-scale features suggesting the past presence of water in an area [7]. That paper evaluated a two-part hypothesis that theorized that the science team would preferentially view tiles just below and along the horizon for water features and topographic highs useful in determining rover position and tiles directly in front of the rover for navigation obstacles. Analysis of the viewing patterns showed that the science team does preferentially view tiles below and along the horizon. It did not show any difference in preferring to view tiles in front of the rover as opposed to either side of the image.

That viewing-pattern analysis considered the tiles examined by the science team during the 2004 investigation. Reference [7] then institutes a task analysis that analyzes how the scientists used the panorama during the 2005 investigation. The objective was to determine the tasks the science team completed while viewing the panorama to provide insight on the importance of targets to the science team. The science team spent 54% of the panorama viewing time trying to estimate Zoë's position in the environment [7]. Another 28% of the time went to data-analysis activities. These findings further supported the findings from the 2004 analysis by showing that local-navigation-related tasks were of little importance when viewing the panorama due to the rover's autonomous hazard-avoidance system. The final product of that paper was a list of targets (sun, topographic highs, clouds, drainages, channels, slopes, and drop-offs) that the science team searches for in the panorama.

C. Selective Information Extraction

The previous analysis clearly shows that the science team uses the panorama for specific purposes and that certain areas are more likely to contain needed information. The next step is to determine a method that takes advantage of this knowledge to acquire necessary information from the environment while committing less bandwidth to the panoramas. The design concept underlying the research presented here is to engineer a camera system that selectively harvests only the needed information from the environment. The premise is that an optimized camera system need not collect an entire scene at a single level of resolution. Instead, the camera can photograph areas at a resolution that correlates to the amount of information likely to exist in that region. The resolution of regions in an image will vary based on the size and detail of the targets the operator searches for in that area. This process will save bandwidth for robotic missions by reducing extraneous bits of information. This engineered extraction method provides a greater wealth of information at a lower overall bandwidth cost.

IV. Theory Development

A. Image Optimization

The results of the 2004 and 2005 LITA analysis show that the science team searches collected images with the specific objective of trying to identify targets of interest. The next step is to translate these search patterns into a model that is useful in designing an information acquisition system. The process of developing this model starts by returning to the models that address search in the context of industrial inspection.

The general basis of the models developed by Morawski *et al.* [20], [21], [40] was to account for the value of detecting an error and the cost associated with missing an error or the cost of the inspector's search time. A similar approach is possible in robotic search. Operators gain value from detecting a target but have an associated cost of the bandwidth to acquire the image. The model should provide a way to optimize information value while controlling or reducing the bandwidth.

As search begins with defining the targets [7], the model's derivation begins with the concept of a target. The model defines the set of all desired targets T consisting of individual targets, $\{t_i|t_1, t_2, \ldots, t_N\}$, such that $t_i \in T$. In robotic exploration, the information encoded by targets is not the same for each target. Therefore, it is important to understand what value $V(t_i)$ the mission context assigns to the target. This value represents a relative importance ranking of the information encoded by that target to the successful completion of the mission objectives.

The targets are located in the environment R, which divides into finite regions r_j for search. The individual regions consist of a set of adjacent pixels captured at a similar resolution and image encoding (bits per pixel). Each region, based on its size, resolution, and encoding, has a defined number of bits $B(r_j)$.

Associated with each target is the probability of locating that target within the region. This probability depends upon the likelihood that the target exists in the image and the resolution of the image provides the appropriate detail for the operator to detect the target. This gives the conditional probability $p(t_i|B(r_j))$ signifying the probability identifying t_i, given the bits used to acquire the region $B(r_j)$. Multiplication of the target value by the probability of locating the target within a given region provides the associated information value, I

$$I(t_i, r_j) = V(t_i)p\left(t_i|B(r_j)\right). \tag{1}$$

The objective of the model is to increase available information as in (2). In the context of mobile robotics, any attempt to maximize information must balance the need to control the number of bits transmitted, or bandwidth as in (3)

$$\max \sum_{j=0}^{m} \sum_{i=0}^{n} I(t_i, r_j) \tag{2}$$

$$\min \sum_{j=0}^{m} B(r_j). \tag{3}$$

The solution to the tradeoff between maximizing information value and minimizing bandwidth is to maximize the quotient of information over bandwidth. The optimized solution of this equation will maximize the numerator, information, and minimize the denominator, bandwidth. Equation (4) shows this division with (1) substituted in for (2)

$$\max \frac{\sum\limits_{j=0}^{m} \sum\limits_{i=0}^{n} V(t_i)p\left(t_i|B(r_j)\right)}{\sum\limits_{j=0}^{m} B(r_j)}. \tag{4}$$

Equation (4) has one potential problem that may prohibit its application to an actual mission. In a robotic geology task, and many other robotic exploration missions, the number of potential targets is simply too large to feasibly optimize. Consider a geologist examining the foreground for rocks that fall in a single class, such as those that are less than 20 cm in diameter. In the field, geologists classify these rocks using a 6 × 6 matrix that rates the rock angularity and sphericity. The color of the rock also often relays information critical to the mission. Human absolute-judgment ability limits the number of subcategories in which humans can judge color to approximately seven categories [15]. Simplifying to six colors and the 6 × 6 matrix, geologists could consider 216 different targets for the class of rocks under 20 cm in diameter. A complete target list for a mission could number in the thousands.

Instead of trying to optimize for every possible target, a different option is to group targets based on their probability of detection. This necessary simplification provides a solvable situation. This new entity, a target class, represents all targets that operators have a similar probability of detecting in an image. Grouping into classes in this manner allows for the partitioning of all the targets into pairwise disjoint sets $A_1, A_2, \ldots, A_m$, where all the t_i's in a set share the same probability density function $p_{ti}(B(r_j))$

$$A_k = \{t_i : p\left(t_i|B(r_j)\right) \equiv p\left(t_d|B(r_j)\right) \forall t_d \in A_k\} \tag{5}$$

$$A_k \cap A_m = \phi \forall k \neq m. \tag{6}$$

Given this set of conditions and calculating a new value for each A, (4) now becomes

$$\max \frac{\sum_{j=0}^{m} \sum_{k=0}^{s} V(A_k)p\left(A_k|B(r_j)\right)}{\sum_{j=0}^{m} B(r_j)}. \tag{7}$$

TABLE I
LIST OF TARGET CLASSES AND THEIR APPROXIMATED VALUES BASED ON OBSERVATIONS DURING THE LITA FIELD TESTS

	Target Class (A_k) Value	(0-1)
0	Sun	1
1	Clouds	0.8
2	Local Topographic Highs *(i.e. Mountains/Hills)*	0.9
3	Slopes/Drop-offs *(Difficult to traverse regions)*	0.5
4	Drainages / Channels	0.75
5	Rocks >1m	0.3
6	Rocks <1m	0.3
7	Sediment	0.3

B. Model Evaluation

Evaluation of the model involves in determining whether the model can account for the science team's differential viewing patterns when looking at the panorama. The analysis begins with the identification of target classes and the assignment of relative values for each target class. Table I shows the eight target classes and their associated values. The target classes are those identified in [7], along with some additional classes identified when reviewing daily science team's summaries from the mission. The objective is to determine the value of each of these targets to the mission. The values discussed as follows represent results of a post-hoc analysis of scientist's behavior during the LITA investigations. Actual robotic exploration values may differ from those based on the LITA observations.

The 2005 analysis showed that the scientists spent the most time viewing the panorama to localize the rover [7]. This task requires locating the sun to determine direction and, then, matching topographic highs viewed in the panorama to mountains on the orbital maps. Therefore, these two target classes have the highest values.

The two targets with the next highest values are clouds and drainages or channels. These two targets fit well with the science team's objective to follow the water present in the environment. Additionally, clouds can signify the presence of fog, which is the main source of atmospheric water in the region. The team wanted to run special morning operations to try to detect life in the presence of fog. However, the team could not regularly run these morning operations, as morning operations consume battery power that the team had to conserve.

After these four target classes, the science team observed the remaining targets less during the investigation. Detecting slopes and drop-offs helps in choosing paths to navigate the rover. Zoë has the capability to determine safe paths around obstacles, reducing the effort the science team must put into determining navigation paths. The other remaining target classes consider targets the science team occasionally discussed in their daily science summaries. Due to environmental conditions and difficulty of detecting rocks in the panorama, the scientists only held limited discussions about these targets.

Since the science team differentially views regions of the panorama based on camera elevation angle, these divisions serve as the regions for the analysis. Table II shows the listing of the regions and their corresponding elevations.

TABLE II
LIST OF REGIONS AND THEIR ASSOCIATED CAMERA ELEVATION ANGLES

Region (r_j)	Elevation
0	-90
1	-76
2	-62
3	-48
4	-34
5	-20
6	-6
7	8
8	22

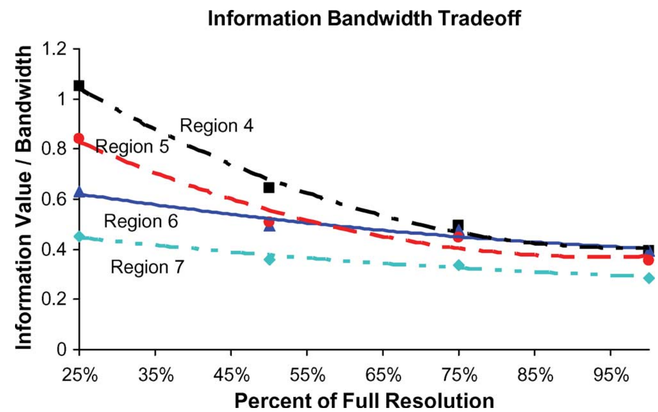

Fig. 1. Solution to (6) for regions 4–7 given the estimated value, probability, and bandwidth for LITA investigation targets.

Next, the analysis determines a probability for detecting each of the target classes in each of the regions given the capture resolution. The full probability tables appear in the Appendix. Capture resolution was 25%, 50%, 75%, or 100% of full resolution available during the LITA investigation. The probabilities represent estimations based on the reviewed visual-search literature and research on how well geologists detect targets in images [41]. The findings, as shown in Fig. 1, result from evaluating (8) with the generated values, probabilities, and bandwidths. Note that Fig. 1 only shows the results of regions 4–7 for clarity purposes. The other regions show a similar trend.

The unexpected result is that the model predicts a low acquisition resolution for all images. Equation (7) is sensitive to increases in bandwidth and, as such, predicts an unfeasible resolution. The science team could not complete the LITA investigation with only 25% resolution images. Fig. 2 shows the interaction between information value and bandwidth. The hope is that the figure will provide insight into how to refine the model to give it some realistic predictive value.

An unexpected feature of these graphs is the presence of a logarithmic best fit curve. Having a logarithmic fit to the data suggests that regions have a high initial rate of information return, but as the bandwidth increases, this rate of information return decreases. The rate of information return for each region depends on the targets present in that region, the value of those targets, and the probability of detecting the targets. Therefore, optimization of data return during a robotic mission requires in determining the rate of information return for each region.

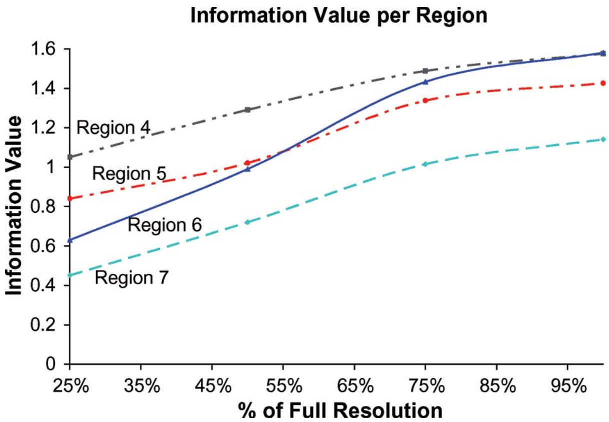

Fig. 2. Comparison of the information value to the collection bandwidth for regions 4–7.

TABLE III
LOGARITHMIC BEST FIT EQUATIONS FOR THE VARIOUS REGIONS IN A PANORAMIC IMAGE

Region	Logarithmic Fit	R^2 Value
0	0.1614Ln(x) + 0.9861	0.9676
1	0.1614Ln(x) + 0.9861	0.9676
2	0.1929Ln(x) + 1.1773	0.9847
3	0.1614Ln(x) + 1.4361	0.9676
4	0.3868Ln(x) + 1.5795	0.9945
5	0.4420Ln(x) + 1.4172	0.9471
6	0.7100Ln(x) + 1.5783	0.9751
7	0.5098Ln(x) + 1.1330	0.9826
8	0.1937Ln(x) + 2.1733	0.9676
FI	1.0118Ln(x) + 2.3350	0.9959

Calculation of the rate of information return $\dot{I}$ requires taking the derivative of (7) to produce

$$\dot{I} = \frac{dV}{dB} \frac{\sum_{j=0}^{m} \sum_{k=0}^{s} V(A_k) p_{A_k} (B(r_j))}{\sum_{j=0}^{m} B(r_j)}. \tag{8}$$

C. Theoretical Summary

During a robotic mission, when choosing between data products, operators consider how much bandwidth they have to invest to get a certain amount of information. This model, as developed in (8), measures the rate of information return, thus measuring the salient mission aspect. Measuring the absolute information value is not relevant and difficult to accomplish within the context of a robotic mission. The rover design and capabilities plays a large role in generating a frame of reference for the operator. This frame of reference affects the relative value of data products, and thus, operators understand rates of information return, which are not absolute information rates. Whether optimizing between existing systems on a robot or using the model to design a new system, measuring the rate of information return provides a measure that operators can understand.

This paper proceeds by showing different methods for using the model to optimize data return. The model can optimize panorama collection to generate bandwidth savings. Additionally, it predicts the relative importance of different data products. This prediction can help generate optimal data collection templates, a feature that improves operator efficiency [6].

V. DISCUSSION

A. Optimized Panorama Collection

Based on Fig. 2, a best fit logarithmic curve for each data series provided a smoothing of the data and generated an equation for each region. Table III shows the equation for each line along with the corresponding R^2 value.

The derivative of these derived logarithmic equations provides the rate of information return for each region at different resolution levels. The question becomes as follows: what is the appropriate rate of information return to require from the instrument. Determining this value will depend on knowing the mission characteristics and system constraints. In a situation where the rover only returns one data product, the selected rate of information return will represent that value that best utilizes the available bandwidth. In the case of the LITA mission, there are other data products, particularly the FI, to consider. In this situation, the system should strive for similar rates of information return across data products. Included in Table III is the best fit equation for the FI. The equation for the FI was determined using a similar process to the one used in determining the regions for the panorama.

The science team did not have the option of collecting the FI at a lower resolution, so they always received FI images at a rate of information return of 1.0118. This rate of information return is much higher than the full-resolution rate of information return for the panorama regions. This discrepancy is in part due to an increased number of targets in the FI over the panorama. Normalizing the rate of information return across image sources so that it becomes an average rate of information return per target allows easier comparison between information sources. This normalization process only takes into account targets with a recorded probability of occurring in the region.

The normalized values permit a comparison between the rates of information return between different imaging sources. Assuming all else is equal, the scientists will want at least the same rate of information return per bit of returned image data from each region of the panorama as they receive from the FI. Using the normalized FI return rate of 0.1690 and acquiring at least that level from the panorama leads to Fig. 3, which shows the various acquisition resolutions for the different regions.

Shown in Fig. 4 is the observed viewing pattern of the science team during the LITA 2004 investigation [7]. The capture resolution as predicted by the model has a high correlation with the actual observed viewing pattern. A best fit line to a scatter plot graph with the actual viewing percentage on the x-axis and the predicted capture resolution on the y-axis has an $R^2 = 0.7895$ value. This finding provides considerable support as to the predictive ability of the model and its actual utility to a robotic mission.

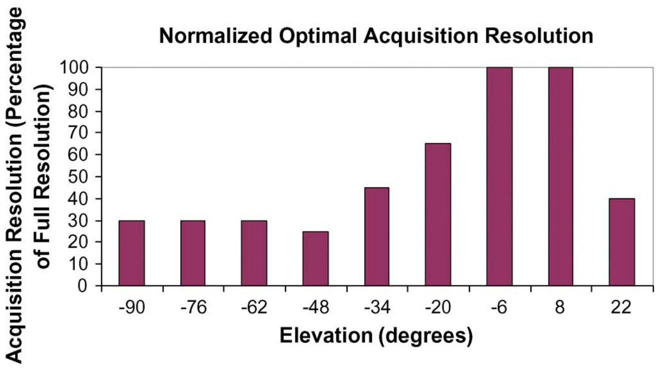

Fig. 3. Optimal resolution acquisition for each region in the panorama that ensures a rate of information return around 0.1690.

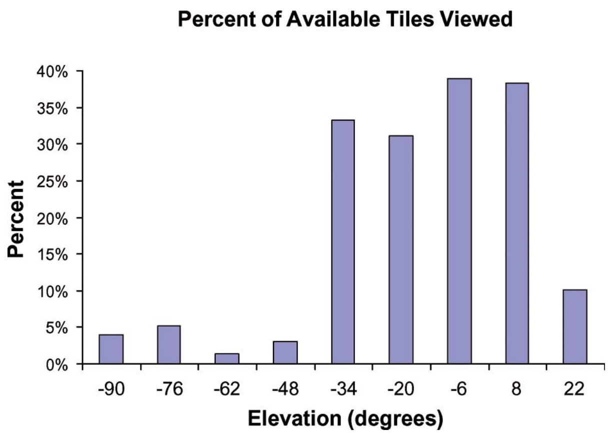

Fig. 4. Percentage of tiles viewed by the science team during the critical search period split by collection elevation.

Another way to look at the model utility is to understand what bandwidth savings the model generates for the LITA investigation. Capturing a full panorama during the LITA investigation required approximately 33 or 3.67 MB per region. The final bandwidth for the panorama collected based on the values shown in Fig. 3 is 17 MB. Implementation of this optimization process on the LITA investigation could save 16 MB, or 48.5%, of each panorama's bandwidth. In the context of LITA, a bandwidth saving of 15 MB provides enough bandwidth for the team to collect three full FI series. Since the FI is the main method of detecting life, it is beneficial to allow the science team to increase their ability to collect FI images.

During the 2005 investigation, the science team developed a standard periodic sampling unit (SPSU) which consisted of taking three full FIs separated by 90 m to characterize a geologic unit. Assuming enough available time, the team would typically prefer to collect an additional SPSU, given that they did not have to lose information gained by collecting a panorama. The optimized panorama provides the science team with the opportunity to collect the important data from the panorama while setting aside enough bandwidth for an extra SPSU.

During the investigation, the team often only collected the top six regions of the panorama. Using the model, the science team could collect all nine rows and still save 4.95 MB of bandwidth in comparison to a six-row panorama. The model's six-region panorama saves 8.35 MB (37.9%) over the six-region panorama collected during the investigation.

TABLE IV
NORMALIZED RATE OF INFORMATION RETURN FOR EACH REGION DURING A MIDDAY CAPTURE PANORAMA

Region	Normalized Rate Return	Number of Targets
SPI-0	0.0538	3
SPI-1	0.0538	3
SPI-2	0.0456	4
SPI-3	0.0404	4
SPI-4	0.0620	5
SPI-5	0.1066	4
SPI-6	0.1547	4
SPI-7	0.1255	3
SPI-8	0.0760	3

B. Template Design

In a complex mission with multiple data products, operators often have initial difficulty in understanding when to utilize the different data products. Using the model to determine the rate of information return for different data products can help the operators understand under what circumstances to use the different data products.

An example from the LITA investigation involves how many panoramas to collect throughout the day. As analysis has shown, the team spent 54% of the panorama viewing time attempting to localize the rover's position [7]. Localization most often only occurs for the last site of the day. If the science team does not use the panorama to localize, the panorama may have a lower overall mission utility. It will still serve a purpose in developing a scientific context of other data collected near the panorama but many of the regions of the panorama may have a suboptimal rate of information return. An observation from the 2005 investigation was that the science team collected fewer midday panoramas as the investigation progressed. During the investigation of site D, the team initially requested five intermediate panoramas. The next day, they requested four, and on subsequent days, only requested two. When the team returned for site E, they began by only requesting one panorama a day at the final stop. This pattern persisted through site F. A number of factors change throughout the investigation and correlate with the science team's choice to collect fewer panoramas. These possibilities include different investigation lengths, hardware issues during site F, and the wrong triangulation of the rover's position during site F.

As discussed, the important data source for the science team is the FI. If the midday panoramas do not return information at the same rate as the FI, the science team must make the difficult choice to drop panoramas in favor of collecting FI data. Recalculating the rate of information return for midday panoramas, which includes reducing the value of detecting targets for localization or navigation obstacles, shows a marked decrease in the rate of information return, as shown in Table IV. The finding from this analysis is that, in general, the rate of information return for a midday panorama is less than that of the FI. Regions 5, 6, and 7 have a reasonable enough rate of information return to suggest a benefit of collecting these regions at a resolution varying between 50% and 75%. Collecting a three-region panorama can provide the science team with the context they need of the region without committing a sizable percentage of the team's bandwidth. Presenting this information to the science team at the beginning of site D would have provided the team with a template to work with from day one. This would save the science team the time spent in testing different data collection protocols and provide an optimal solution that they did not consider.

C. Applying the Model

Based on observations during the 2004 and 2005 LITA investigations, target search is a complex visual process. Studying the tasks and underlying visual mechanisms led to the development of a model that provides a method for determining the rate of information return for a collected image. Then, a post-hoc evaluation provided values for each of the target classes that reflect the observed science team's actions during the investigation. These values may not entirely represent the true values for planetary robotic exploration. They do provide an excellent opportunity to observe that the model provides predictions that match typical scientist behavior. The key point to extract is that optimizing the rate of information return provides an excellent method for improving the efficiency of data collection during a robotic exploration task.

To improve data collection during a robotic mission, the process must start with a clear understanding of the objective task. The designer of the system must work with potential operators to learn what information they extract from images. This information provides a list of target classes. From there, the next steps involve in determining the ability to identify the targets at different resolutions and the value of the targets. Obtaining value estimates from operators is a difficult task. Many of the targets associated with robotic exploration do not have an intrinsic value such as in industrial inspection but, instead, only have a relative value. Presenting operators with statements where they have to choose the ability to detect one target over another is one potential method to develop an understanding of the relative importance of targets within the operator's minds.

Once the mission team decides on the appropriate values for their mission targets, they must decide on what is the optimal rate of information return for the mission. As occurred in this paper's analysis, comparing information rate of return between different image sources is a useful method to optimize across data products. This however assumes the availability of multiple image sources. The original development of the model worked under the presumption of a situation where the rover only collects images from a single source. In these instances, the controlling parameter is bandwidth. The system places some constraint on the absolute amount of imagery the rover can return. The model operates inside this system by providing a relationship between different regions in the image to let the camera system optimize its data collection to provide the maximal information for the available bandwidth.

VI. CONCLUSION

A thorough analysis of the complicated search task associated with target identification led to a predictive model that evaluates the rate of information return provided by different regions of an image. Robotic exploration missions must constantly deal with bandwidth limits that degrade the overall level of information available to operators. By analyzing the information needed by rover operators, it is possible to reduce the bandwidth utilized by low information sources. It is critical to analyze thoroughly the operator's information needs; otherwise, the collected data may lead to the operators reaching erroneous conclusions and missing the mission objectives. This provides additional bandwidth for other information sources. The differentiation between low and high information depends on the probability of locating a target of interest within the region and the relative value of detecting the target in relationship to other potential targets in the environment.

Validation of the model includes showing its ability to predict observed behaviors during the 2004 and 2005 LITA investigations. Typical use of the method should involve analyzing operator needs before a mission and then collecting optimized data throughout the mission. The post-hoc analysis method does show that the model's predictions provide optimized data collection methods that closely model normal operator optimization behavior. During the LITA investigation, the scientists viewed high-resolution tiles in the panorama based on the information contained in those tiles. The model can predict the areas needed at high resolution based on the targets the science team detected in those regions. The optimization covered in this paper only applies to the specific set of circumstances experienced during the LITA investigation. The specific optimization for a mission will vary based on its objective and the expected environment.

Utilizing the model for robotic exploration provides operators with considerable benefit. Most exploration missions force operators to choose between some tradeoff between data sources or image parameters. For example, the LITA scientists had to choose between multiple daily panoramas or more FI sequences. Rarely is the operator team entirely happy with the tradeoff they decide on. Providing optimized data collection will provide the operator team with more options to utilize their bandwidth.

The model also provides insights for future developments in robotics. The future for robots suggests algorithms for autonomous data collection. During the 2005 LITA investigation, Zoë implemented a "science-on-the-fly system" that allowed her to collect a full FI sequence if the initial FI collection detected a positive chlorophyll signal. Estlin *et al.* [43] developed and demonstrated the onboard autonomous science investigation system on the Rocky 8 using data from the MER mission Field Integrated Design and Operations field tests. These autonomous collection protocols may benefit from incorporation of the model. It can help control the bandwidth contributed to autonomous collections while ensuring that the rover collects sufficient data for the user to interpret its final meaning. Depending on the image processing capabilities of the autonomous collection systems, it also provides the opportunity for simple algorithms to quantify the information contained in an image and determine whether it deserves a portion of the daily bandwidth.

APPENDIX
VALUE AND PROBABILITY TABLES

TABLE V
PROBABILITY ESTIMATES FOR EACH TARGET CLASS IN IMAGE REGIONS 0–2

			At the Robot's Feet											
Target Class (A_k)		Value (0-1)	$P(r_0)$				$P(r_1)$				$P(r_2)$			
			0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1
0	Sun	1	0	0	0	0	0	0	0	0	0	0	0	0
1	Clouds	0.8	0	0	0	0	0	0	0	0	0	0	0	0
2	Local Topographic Highs, Used for Localization (i.e. Mountains/Hills)	0.9	0	0	0	0	0	0	0	0	0	0	0	0
3	Slopes/Dropoffs (Non- or difficult to traverse regions	0.5	0	0	0	0	0	0	0	0	0	0	0	0
4	Drainages / Channels	0.75	0.2	0.4	0.45	0.5	0.2	0.4	0.45	0.5	0.2	0.4	0.45	0.5
5	Rocks >1m	0.3	0	0	0	0	0	0	0	0	0.5	0.55	0.6	0.65
6	Rocks <1m	0.3	1	1	1	1	1	1	1	1	1	1	1	1
7	Sediment	0.3	1	1	1	1	1	1	1	1	1	1	1	1

TABLE VI
PROBABILITY ESTIMATES FOR EACH TARGET CLASS IN IMAGE REGIONS 3–5

			Foreground											
Target Class (A_k)		Value (0-1)	$P(r_3)$				$P(r_4)$				$P(r_5)$			
			0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1
0	Sun	1	0	0	0	0	0	0	0	0	0	0	0	0
1	Clouds	0.8	0	0	0	0	0	0	0	0	0	0	0	0
2	Local Topographic Highs, Used for Localization (i.e. Mountains/Hills)	0.9	0	0	0	0	0	0	0	0	0	0	0	0
3	Slopes/Dropoffs (Non- or difficult to traverse regions	0.5	0	0	0	0	0	0	0.2	0.3	0	0	0.5	0.6
4	Drainages / Channels	0.75	0.4	0.6	0.65	0.7	0.4	0.6	0.65	0.7	0.4	0.6	0.65	0.7
5	Rocks >1m	0.3	1	1	1	1	1	1	1	1	1	1	1	1
6	Rocks <1m	0.3	1	1	1	1	1	1	1	1	0.8	0.9	1	1
7	Sediment	0.3	1	1	1	1	0.5	0.8	1	1	0	0	0	0

TABLE VII
PROBABILITY ESTIMATES FOR EACH TARGET CLASS IN IMAGE REGIONS 6–8

			Horizon				Sky							
Target Class (A_k)		Value (0-1)	$P(r_6)$				$P(r_7)$				$P(r_8)$			
			0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1
0	Sun	1	0	0	0	0	0	0	0	0	1	1	1	1
1	Clouds	0.8	0	0	0	0	0	0	0	0	1	1	1	1
2	Local Topographic Highs, Used for Localization (i.e. Mountains/Hills)	0.9	0.2	0.4	0.45	0.5	0.5	0.8	0.85	0.9	0.1	0.3	0.35	0.4
3	Slopes/Dropoffs (Non- or difficult to traverse regions	0.5	0	0	0.6	0.7	0	0	0.2	0.3	0	0	0	0
4	Drainages / Channels	0.75	0.4	0.6	0.65	0.7	0	0	0	0	0	0	0	0
5	Rocks >1m	0.3	0.5	0.6	0.8	0.85	0	0	0	0	0	0	0	0
6	Rocks <1m	0.3	0	0	0	0	0	0	0	0	0	0	0	0
7	Sediment	0.3	0	0	0	0	0	0	0	0	0	0	0	0

REFERENCES

[1] T. Blackmon, L. Ngyuen, C. Neveu, D. Rasmussen, E. Zbinden, M. Maimone, L. Matthies, S. Thayer, V. Broz, J. Teza, J. Osborn, M. Hebert, G. Thomas, and J. Steele, "Photorealistic virtual reality mapping system for Chernobyl accident site assessment," in *Proc. SPIE—Human Vision Electronic Imaging IV*, 1999, p. 3644.

[2] R. Murphy, "Rescue robotics for homeland security," *Commun. ACM*, vol. 27, no. 3, pp. 66–69, Mar. 2004.

[3] S. W. Squyres, R. E. Arvidson, E. T. Baumgartner, J. F. Bell, P. R. Christensen, S. Gorevan, K. E. Herkenhoff, G. Klingelhofer, M. B. Madsen, R. V. Morris, R. Rieder, and R. A. Romero, "Athena Mars rover science investigation," *J. Geophys. Res.—Planets*, vol. 108, no. E12, p. 8062, 2003.

[4] J. Casper and R. Murphy, "Human–robot interactions during the robot-assisted urban search and rescue response at the World Trade Center," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 33, no. 3, pp. 367–385, Jun. 2003.
[5] J. Glasgow, E. Pudenz, G. Thomas, P. Coppin, N. Cabrol, and D. Wettergreen, "Observations of a science team during an advanced planetary rover prototype field test," in *Proc. 14th IEEE Int. Workshop Robot Human Interaction Commun.*, Nashville, TN, 2005, pp. 137–142.
[6] E. Pudenz, G. Thomas, J. Glasgow, P. Coppin, D. Wettergreen, and N. Cabrol, "Searching for a quantitative proxy for rover science effectiveness," in *Proc. 1st Conf. Human–Robot Interaction*, 2006, pp. 18–25.
[7] J. M. Glasgow, G. Thomas, E. Pudenz, N. Cabrol, D. Wettergreen, and P. Coppin, "Panoramic image information utility for mobile robot exploration," in *Proc. IEEE Syst. Man Cybern. Conf.*, 2006, pp. 3216–3221.
[8] J. B. F. van Erp and P. Padmos, "Image parameters for driving with indirect viewing systems," *Ergonomics*, vol. 46, no. 15, pp. 1471–1499, Dec. 2003.
[9] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: MIT Press, 1992, pp. 157–161.
[10] M. Kontitsis, K. P. Valavanis, and R. Garcia, "A simple low cost vision system for small unmanned VTOL vehicles," in *Proc. IROS*, 2005, pp. 3480–3486.
[11] H. Schempf, E. Mutschler, C. Piepgras, J. Warwick, B. Chemel, S. Boehmke, W. Crowley, R. Fuchs, and J. Guyot, "Pandora: Autonomous urban robotic reconnaissance system," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1999, vol. 3, pp. 2315–2321.
[12] R. Volpe, J. Balaram, T. Ohm, and R. Ivlev, "The Rocky 7 Mars rover prototype," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Osaka, Japan, 1996, pp. 1558–1564.
[13] B. Yamauchi and P. Rudakevych, "Griffon: A man-portable hybrid UGV/UAV," *Ind. Rob.*, vol. 31, no. 5, pp. 443–450, 2004.
[14] K. E. Herkenhoff, S. W. Squyres, J. F. Bell, J. N. Maki, H. M Arneson, P. Bertelsen, D. I. Brown, S. A. Collins, A. Dingizian, S. T. Elliott, W. Goetz, E. C. Hagerott, A. G. Hayes, M. J. Johnson, R. L. Kirk, S. McLennan, R. V. Morris, L. M. Scherr, M. A. Schwochert, L. R. Shiraishi, G. H. Smith, L. A. Soderblom, J. N. Sohl-Dickstein, and M. V. Wadsworth, "Athena microscopic imager investigation," *J. Geophys. Res.—Planets*, vol. 108, no. E12, p. 8065, 2003.
[15] C. D. Wickens, S. E. Gordon, and Y. Liu, *An Introduction to Human Factors Engineering*. New York: Addison-Wesley, 1998.
[16] M. M. Chun, "Contextual cueing of visual attention," *Trends Cogn. Sci.*, vol. 4, no. 5, pp. 170–178, May 2000.
[17] R. Parasuraman, "Vigilance, monitoring and search," in *Handbook of Perception and Human Performance*, K. R. Boff, L. Kaufman, and J. P. Thomas, Eds. New York: Wiley, 1986.
[18] C. D. Wickens and J. G. Hollands, Eds., *Engineering Psychology and Human Performance*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 2000.
[19] A. L. Yarbus, *Eye Movements and Vision*. New York: Plenum, 1967.
[20] T. B. Morawski, C. G. Drury, and M. H. Karwan, "The optimum speed of visual inspection using a random search strategy," *IIE Trans.*, vol. 24, no. 5, pp. 122–133, 1992.
[21] M. H. Karwan, T. B. Morawski, and C. G. Drury, "Optimum speed of visual inspection using a systematic search strategy," *IIE Trans.*, vol. 27, no. 3, pp. 291–299, 1995.
[22] M. J. Wang, S. Lin, and C. G. Drury, "Training for strategy in visual search," *Int. J. Ind. Ergon.*, vol. 20, no. 2, pp. 101–108, Aug. 1997.
[23] S. K. Hong, "Human stopping strategies in multiple-target search," *Int. J. Ind. Ergon.*, vol. 35, no. 1, pp. 1–12, Jan. 2005.
[24] R. R. Murphy and J. J. Sprouse, "Strategies for searching an area with semi-autonomous mobile robots," in *Proc. ASCE Specialty Conf.*, 1996, pp. 22–28.
[25] T. Kawanishi, H. Murase, S. Takagi, and M. Werner, "Dynamic active search for quick object detection with pan-tilt-zoom camera," in *Proc. Int. Conf. Image Process.*, 2001, vol. 3, pp. 716–719.
[26] D. Bapna, E. Rollins, J. Murphy, M. Maimone, and W. Whittaker, "The Atacama desert track: Outcomes," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1998, vol. 1, pp. 597–604.
[27] G. Cheng and A. Zelinsky, "Supervised autonomy: A paradigm for teleoperating mobile robots," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 1997, vol. 2, pp. 1169–1176.
[28] H. A. Everett, G. A. Gilbreath, and D. A. Ciccimaro, "An advanced telereflexive tactical response robot," *Auton. Robots*, vol. 11, no. 1, pp. 39–47, Jul. 2001.
[29] T. Fong, C. Thorpe, and C. Baur, "Advanced interfaces for vehicle teleoperation: Collaborative control, sensor fusion displays, and remote driving tools," *Auton. Robots*, vol. 11, no. 1, pp. 77–85, Jul. 2001.
[30] K. Kawabata, T. Ishikawa, T. Fujii, H. Asama, and I. Endo, "Collaborative task execution by a human and an autonomous mobile robot in a teleoperated system," *Integr. Comput.-Aided Eng.*, vol. 6, no. 4, pp. 319–330, Dec. 1999.
[31] D. Woods, J. Tittle, M. Feil, and M. Roesler, "Envisioning human–robot coordination in future operations," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 34, no. 2, pp. 138–153, May 2004.
[32] M. R. Endsley, "The application of human factors to the development of expert systems for advanced cockpits," in *Proc. Human Factors Soc. 31st Annu. Meeting: Human Factors Ergonomics Soc.*, Santa Monica, CA, 1987, pp. 1388–1392.
[33] M. R. Endsley, "Design and evaluation for situation awareness enhancement," in *Proc. Human Factors Soc. 32nd Annu. Meeting: Human Factors Ergonomics Soc.*, Santa Monica, CA, 1988, pp. 97–101.
[34] M. R. Endsley, "Designing for situation awareness in complex systems," in *Proc. 2nd Int. Workshop Symbiosis Humans, Artifacts Environ.*, 2001, pp. 176–190.
[35] M. R. Endsley, D. B. Kaber, and E. Onal, "The impact of intermediate LOAs on situation awareness and performance in dynamic control systems," in *Global Perspectives of Human Factors in Power Generation*, D. I. Gertman, D. L. Schurman, and H. S. Blackman, Eds. New York: IEEE, 1997, pp. 7/7–7/12.
[36] D. B. Kaber, E. Onal, and M. R. Endsley, "Design of automation for telerobots and the effect on performance, operator situation awareness, and subjective workload," *Hum. Factors Ergon. Manuf.*, vol. 10, no. 4, pp. 409–429, 2000.
[37] J. C. Scholtz, B. Antonishek, and J. D. Young, "Implementation of a situation awareness assessment tool for evaluation of human–robot interfaces," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 35, no. 4, pp. 450–459, Jul. 2005.
[38] K. J. Vicente and J. Rasmussen, "Ecological interface design: Theoretical foundations," *IEEE Trans. Syst., Man, Cybern.*, vol. 22, no. 4, pp. 589–606, Jul./Aug. 1992.
[39] S. Weinstein, D. Pane, K. Warren-Rhodes, C. Cockell, L. A. Ernst, E. Minkley, G. Fisher, S. Emani, D. S. Wettergreen, M. Wagner, N. Cabrol, and A. S. Waggoner, "Use of a novel rover-mounted fluorescence imager and fluorescent probes to detect biological material in the Atacama desert in daylight," in *Proc. Lunar Planetary Sci. XXXVI*, 2005.
[40] C. G. Drury and C. F. Chi, "A test of economic models of stopping policy in visual search," *Inst. Ind. Eng. Trans.*, vol. 27, no. 3, pp. 382–393, 1995.
[41] J. Wagner, G. Thomas, J. Glasgow, R. C. Anderson, N. Cabrol, and E. Grin, "Error-associated behaviors and error rates for robotic geology," in *Proc. Human Factors Ergonom. Soc. 48th Annu. Meeting*, New Orleans, LA, 2004, pp. 444–447.
[42] F. M. Marchese and D. G. Sorrenti, "Omni-directional vision with a multi-part mirror," in *Proc. RoboCup: Robot Soccer World Cup IV*, 2000, vol. 864, p. 197.
[43] T. Estlin, R. Castano, R. C. Anderson, D. Gaines, F. Fisher, and M. Judd, "Learning and planning for Mars Rover science," in *Proc. IJCAI Workshop Notes Issues Designing Phys. Agents Dynamic Real-Time Environments: World Modeling, Planning, Learning, Commun.*, 2003.

Justin M. Glasgow received the B.S. degree in biomedical engineering and the M.S. degree in industrial engineering from the University of Iowa, Iowa City, in 2005 and 2006, respectively, where he is currently working toward a combined M.D./Ph.D. degree in the Department of Mechanical and Industrial Engineering, focusing his research on improving vaccine design.

From 2003 to 2006, he was a Research Assistant with the Graphical Representation of Knowledge Laboratory, University of Iowa, which was run by Geb Thomas. From 2004 to 2006, he was working on the Life in the Atacama Project, from which will be his third publication. His research interests include data acquisition and system optimization.

Mr. Glasgow was inducted into the IA-B chapter of Tau Beta Pi National Engineering Honor Society in the spring of 2002, has served as Chapter President from 2004 to 2005, and has been serving as a Chapter Adviser since 2006.

Geb Thomas (M'98) received the Ph.D. degree in industrial engineering from The Pennsylvania State University, University Park, in 1996.

He was a National Research Council (NRC) Postdoctoral Research Associate with NASA Ames Research Center, Mountain View, CA, in 1997. In 1998, he was with the Faculty of The University of Iowa, Iowa City, where he is currently an Associate Professor with the Department of Mechanical and Industrial Engineering. His research interests are in the areas of virtual reality and human–robot interaction.

Dr. Thomas is currently a Cochair of the IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS technical committee on human–computer interaction.

Erin Pudenz received the B.S. degree in psychology and the B.S.E. and M.S. degrees in industrial engineering with a focus in human factors from the University of Iowa, Iowa City, in 2004 and 2006, respectively.

She participated and served as President for both the Human Factors and Ergonomics Society Student Chapter and the Institute of Industrial Engineers Student Chapter, University of Iowa. While acting as President of the Institute of Industrial Engineers Student Chapter, the chapter received the Gold Award. After graduation, she was with the consulting company Userthink, Peachtree City, GA. From 2004 to 2006, she was a Research Assistant with the Graphical Representation of Knowledge Laboratory, University of Iowa, which was run by Geb Thomas. During this time, she was working on the Life in the Atacama Project and has several related publications. She is currently a Human Factors Specialist with the Department of Mechanical and Industrial Engineering, The University of Iowa. Her research interests include human factors and human-interaction modeling.

Ms. Pudenz was a member of the Alpha Pi Mu, Industrial Engineering Honors Society, and the Women in Science and Engineering.

Nathalie Cabrol received the M.S. and Ph.D. degrees from Sorbonne University, Paris, France, in 1986 and 1991, respectively.

Since 1998, she has been a SETI Institute Principal Investigator with the NASA Ames Research Center, Moffett Field, CA. Her research interests relate to aqueous environments favorable to life on Mars, their exploration (robotic and human), and the study of terrestrial analogs.

David Wettergreen received the Ph.D. degree in robotics from Carnegie Mellon University, Pittsburgh, PA, in 1995.

From 1996 to 1997, he was an NRC Postdoctoral Research Associate with the NASA Ames Research Center, Mountain View, CA. From 1998 to 2000, he was a Research Fellow with the Australian National University, Canberra, Australia. Since 2000, he has been with the Faculty of Carnegie Mellon University, where he is currently an Associate Research Professor with the Robotics Institute. He creates robotic technologies for exploration underwater, on the surface, and in air and space. His research develops perception, planning, and learning for robot autonomy with recent focus on merging multiple scales of sensing for navigation and in autonomously interpreting and acting upon scientific observations. For nearly two decades, he has created robotic explorers for challenging environments in the Arctic and Antarctic, in Alaska, Australia, Chile, and, recently, Mexico.

Peter Coppin received the M.F.A. degree (the terminal degree in art) in electronic and time-based media from Carnegie Mellon University, Pittsburgh, PA, in 1998.

From 1999 to 2006, he was a Research Fellow with the STUDIO for Creative Inquiry, Carnegie Mellon University, where he developed the NASA-funded EventScope Project to enable the public to remotely experience faraway places through robots. He is currently a Project Scientist with the Robotics Institute, Carnegie Mellon University. For over a decade, he has created information systems to enable people to remotely experience real places and situations from a distance, from venues such as the home, schools, science museums, through the web, and more recently, rover operations centers. His research interest is in the areas of virtual presence, science visualization, and interfaces to support remote exploration.