

MARS EXPLORATION ROVER PERFORMANCE AS A BASELINE FOR FLIGHT ROVER AUTONOMY TECHNOLOGY ASSESSMENT

Edward Tunstel, Ayanna Howard, Mark Maimone, Ashitey Trebi-Ollennu

NASA Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109 USA,

Email: tunstel@robotics.jpl.nasa.gov

ABSTRACT

Technology assessments rely on performance metrics to establish a basis for rating technologies. Metrics are also used to measure relative merit of similar technologies to state-of-the-art technology. Functional performance metrics are presented for mobility and robotic arm autonomy exercised on the Mars Exploration Rovers (MER) surface mission thus far. The metrics are used to apply an existing technology assessment method to establish a baseline for assessing future flight rover technologies. The methodology decomposes robotic activities into operational functions and addresses how technologies, based on performance metrics, can be systematically related to increases in science return. Considering the basic mission objective to maximize scientific yield, we can assess how relative performance of future technologies might impact science return. We provide a useful set of metrics and present an example application of the method to assess merit of hypothetical future Mars rover performance relative to the MER baseline.

1. INTRODUCTION

The utility of autonomous rovers is a function of their ability to move about and explore without frequent contact with Earth-based mission operators. More increasingly, robotic vehicle autonomy is required to achieve aspects of overall success for planetary surface missions such as the 2003 Mars Exploration Rovers (MER) and the planned 2009 Mars Science Laboratory (MSL) and ExoMars missions. With MER, NASA landed two twin rovers, named *Spirit* and *Opportunity*, on Mars in January of 2004 (Fig. 1). These rovers were explicitly required to use robotic mobility and manipulator arm positioning functionality to achieve exploration mission objectives by serving as surrogate robotic field geologists for a science team on Earth. MSL and ExoMars may be required to do the same, and perhaps more, with greater demand on autonomy and lifetime.

Investments in associated autonomy research or technology products are based in part on potential to maximize scientific yield of the missions. Mission

planners or systems engineers tend to justify the inclusion of new technology by determining its effect on the utility of the mission, often computed by combining the utility of outcome with the probability of achieving the outcome [1]. Utility of autonomy technology should therefore be linked to achievement of mission goals by evaluating impact on science return. However, few systematic methodologies exist that quantify the concept of science return due to autonomy technology components.



Fig. 1. *Spirit* and lander (computer models combined with 3-D surface data acquired by *Spirit*'s cameras).

Recent work has proposed and developed a framework that can systematically relate technologies to science return in a structured fashion [2, 3]. Mission objectives are quantified in terms of science return and achievement of objectives is represented as a set of mission operational functions. Technologies are then linked to the mission through several levels of abstraction that associate performance metrics with science return. A fundamental requirement for applying this method is availability of a set of computed performance metrics that are associated with a selected baseline technology. For purposes of technology assessment, the baseline is viewed as the state-of-the-art (SOA) against which a new or alternative technology can be evaluated with respect to relative impact on science return.

MER represents the longest deployment of planetary surface robots and a new benchmark in planetary robotic autonomy. As such, it is important to capture and document the rovers' performance in ways that facilitate evaluation of similar technologies relative to

the SOA established by the MER benchmark. This paper lays groundwork necessary to establish mobility and related robotic arm autonomy used during MER surface operations as a SOA baseline useful for applying the technology assessment methodology in [2]. Performance metrics for MER surface navigation and robotic arm instrument placement are computed based on actual performance data from *Spirit* and *Opportunity* on Mars.

2. AUTONOMY TECHNOLOGY ASSESSMENT

The first step toward autonomy technology assessment for current and future space flight missions is to decompose mission scenarios into mission operational functions [2]. Operational functions can further be subdivided into functional sequences and functional sequences into functional steps. This decomposition results in a hierarchy that links technologies to mission science objectives by associating performance metrics more directly with science return.

2.1 Mission Operational Functions

For MER, surface robotic autonomy scenarios employ operational functions that can be separated into functional sequences including *Navigation*, *Approach* to science targets, and *Instrument Placement* (IP) onto science targets. Navigation functional steps include: specify surface goal, plan traverse toward goal, and execute traverse to goal. Approach and IP steps include: specify nearby science target, perform short approach to target, deploy robotic arm, and place instrument on target. MER software functionality enables robotic execution of these functional steps and sequences. Robotic tasks include wheel motion control and vision-guided autonomous navigation functions of varying complexity for traversing the Martian surface, as well as robotic arm motion control functions for accurate placement of science instruments onto rocks and soil. Combined and repeated use of mobility and robotic arm capabilities enables acquisition of desired high priority science measurements.

To enable science return, technological capabilities must directly address at least one of three attributes of scientific measurements: *quality*, *quantity*, or *diversity*. Based on this notion, the value of a technology component can be assessed by determining how it impacts each of these three science return attributes [2]. In order to pursue this for MER surface robotic autonomy, performance metrics associated with the functional steps mentioned above must be formulated and computed. A technology impact score, calculated at the functional level, can then be propagated through several levels of abstraction from technology to science to determine quantitative overall impact on science

return (Sec. 3). One example of such levels is illustrated in Fig. 2 where the highest level is science return and the three science return components are located at the science return subclass level. Below that subclass level is operations, which represents the functional sequences associated with an operational function of the mission. At the next level are technologies that execute steps associated with a functional sequence. Technologies that enable achievement of functional steps reside at the last level.

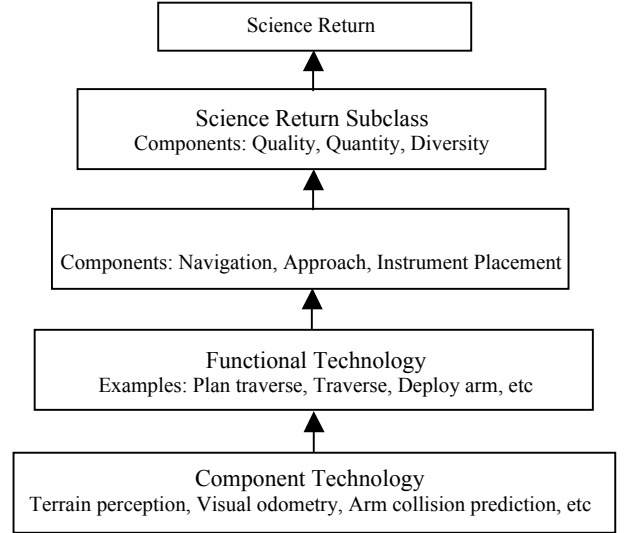


Fig. 2. Technology levels linking to science return.

3. IMPACT ON SCIENCE RETURN

Performance metrics capture important attributes of each technology, as relate to their associated technology levels in Fig 2. To enable reasonable comparison of different technologies, a template that characterizes technologies in terms of four types of metrics is used [2]. The four types of performance metrics are task dependencies (inputs), task results (outputs), environmental constraints (e.g., terrain-related), and resource constraints. In addition, performance values are normalized to facilitate multi-attribute assessment and understanding of the relative impact of technologies on the mission. A technology impact score, S_T , for technology component T , is thus calculated as an averaged normalized difference between performance values for the technology of note and for the selected baseline SOA technology, such that:

$$S_T = \frac{1}{w} \sum_{p=1}^w \frac{\square_p \square_p \square_p}{\square_p} \quad (1)$$

where w is the number of performance metrics used in the assessment of the technology component, and \square_p and \square_p are SOA and technology component capability

measures with respect to performance metric, p . This difference value is summed over all employed performance metrics that are common to the technology. Note that the sign of the technology impact score of Eq. 1 depends on the type of metric and its numerical interpretation. Since output metrics are dependent on inputs metrics, $-S_T$ is used for impact scores of input technology metrics. In addition, if the performance metric value for a given technology is interpreted such that the smaller the number, the better the improvement, then $-S_T$ is used rather than $+S_T$. Once S_T is calculated, the technology *level* impact score, S_L , for all technology components at a given level, L , of the hierarchy is then determined as

$$S_L = \frac{1}{w_L} \sum S_T \quad (2)$$

where w_L is the number of technology components resident at that level and individual technology impact scores are summed over all components resident at level L . Technology impact scores are propagated up the hierarchy to establish overall impact on science return as follows [2]. Scores at level L , S_T and S_L , are used recursively to compute the same for the next level above. This chain of calculations is repeated until the top level (science return) is reached. The following general formulation is used for each technology component, i , resident at upper level, $\square = 2, 3, \dots$

$$S_T(i, \square) = \frac{1}{w_i} \sum S_T(i, \square-1) \quad (3)$$

$$S_\square = \frac{1}{w_\square} \sum S_T(i, \square) \quad (4)$$

In Eqs. 3 and 4, w_i is the impact weight of technology component i computed as the number of related components from the previous level divided by the number of components the technology effects. As such, the technology impact score for component i is summed over all components at the $(\square-1)^{\text{th}}$ level that are linked to component i at level \square . The technology *level* impact score is then computed by summing over all technology impact scores for components resident at level \square and averaging via division by, w_\square , the total number of technology components resident at that level. Since the highest level of the hierarchy includes the science return components, technology impact scores calculated at the top level represents a measure of the overall technology impact on science return. When assessing the combined impact of multiple technologies we consider the science return for the mission to be represented by the average technology impact score for all technologies involved.

This technology assessment methodology is one of a number of possible approaches that could be employed. Its formulation is one of the few proposed that provide a systematic approach to technology assessment. It is not without drawbacks in its current form, however, so its key limitation and underlying assumption are noted. The key limitation of the assessment methodology is that technologies are assumed to be independent. For example, consider a vision algorithm that enables hazard detection in a navigation algorithm. If we wish to assess these two technologies independently, the performance associated with the navigation algorithm would implicitly benefit from the performance derived from the vision algorithm. In this regard, the navigation algorithm would lead to a higher science return, even though it requires the vision algorithm to fully achieve its potential. To avoid this, we would need to extract navigation performance output values that would not be based on vision; this then becomes a less realistic scenario. A key underlying assumption of the methodology is that all metrics are assumed equally weighted. This is not true in all cases. For example, if safety is of more value to the mission than distance traversed, we would like to apply a higher weighting factor to safety than to the distance output. Required changes to the formulation would include a means to incorporate different weighting factors that are associated with the different performance metrics.

4. PERFORMANCE METRICS

The above procedure is initiated by applying selected performance metrics that are used to establish the technology impact on science return for a selected baseline or SOA system. The same metrics are applied to a new or alternative system to aid in assessment of its science return potential relative to the SOA system.

We present a set of performance metrics to establish a numerical baseline of technology impact scores that could be used to assess similar rover technologies. The set presented here is by no means complete but only a small subset of many that could be applied to rover systems. It should be noted that any reasonable set of performance metrics could be used for this purpose, and it should be emphasized that any relative comparison with other systems are only proper if they employ the same metrics. This is a strict necessity for valid relative assessments using the method of Sec. 3.

Throughout the many months of the MER surface mission, combined and repeated use of mobility and robotic arm capabilities has enabled acquisition of high priority science measurements. The mobility and robotic arm software runs onboard the rovers' computers to consistently perform integral parts of various exploration tasks. Each MER computer is a 20

MHz RAD-6000 processor (radiation-hardened version of a PowerPC chip) running the VxWorks real-time operating system, with 128 MB of DRAM and 256 MB flash memory and EEPROM, embedded in a VME chassis. Robotic tasks are specified in command loads uplinked to the rovers by engineers who plan their daily robotic activities on Earth given desired science activities prescribed by scientists. Given a set of command sequences that would implement exploration activities for a given day, *Spirit* and *Opportunity* set out to autonomously perform the necessary operational functions. Recalling the MER functional sequences including Navigation, Approach, and IP, we define a set of simple performance metrics for these surface mobility and robotic arm activities. In addition, we make use of a wealth of actual telemetry returned by *Spirit* and *Opportunity* over hundreds of sols (Martian days) to compute baseline values for the metrics.

4.1 Autonomous Navigation

The MER vehicles execute a navigation algorithm called GESTALT (Grid-based Estimation of Surface Traversability Applied to Local Terrain), which is documented in [4]. GESTALT performs stereo vision-based perception, local terrain hazard mapping, traversability assessment, and incremental goal-directed path selection through its local traversability map. In addition to these forms of rover autonomy for navigation, the rovers' onboard software performs visual odometry when commanded and low-level reactive fault protection based on proprioceptive sensing to achieve self-localized and safeguarded mobility. For this work, we present performance metrics related to: autonomous traverse rates, percentage of traverses performed autonomously, time required to make autonomous navigation decisions, and the mean number of self-localizations per traverse.

The average rate at which a rover can traverse autonomously depends on the traversability of the terrain over which such a measurement applies (among other things). A given rover may traverse flat and hazard-free terrain at a faster average rate than it would a sloped and rocky terrain due to the increased deliberation required in the latter case. Therefore, a metric for traverse rate might ideally account for terrain type in some meaningful way. We do not attempt to achieve such an ideal formulation here. Instead, we choose a formulation that somewhat reduces dependence on terrain type by considering performance within a single regional terrain type such as that at a given landing site (assuming local variability is not dramatic from place to place). The following metric is a relative measure of average to maximum traverse rate achieved in a given regional

terrain type which we will refer to as the *autonomous traverse speed ratio* (ATSR).

$$ATSR = \frac{AverageAutonavSpeed}{MaximumAutonavSpeed} \quad (5)$$

Spirit's average and maximum autonomous traverse rates at its Mars landing site (Gusev Crater) thus far are 17.5 m/hr and 34.35 m/hr, respectively. *Opportunity's* average and maximum autonomous traverse rates at its landing site (Meridiani Planum) thus far are 25.5 m/hr and 36 m/hr, respectively. Herein, we make use of an average of the ATSR for *Spirit* and *Opportunity* as the MER ATSR.

The MER traverses are commanded using a variety of mobility modes ranging from sequenced segments of primitive driving motions without hazard avoidance enabled (i.e., blind) to full autonomous navigation, sometimes including self-localization using visual odometry. We present a measure of the percentage of overall traverses performed autonomously (versus using manually sequenced drive primitives with autonomy disabled). Vision-based hazard detection and avoidance along with visual odometry are the main autonomy modes of local surface navigation for MER. While the rovers are capable of executing them simultaneously, MER mobility planners rarely commanded both capabilities to execute at the same time. As such, we compute the *percent autonomous traverse* (PAT), or percent of mission traverse performed using onboard vision-based autonomy, as follows, where d_{auto} , d_{visod} , and d_{blind} are traverse distances performed using autonomous navigation, visual odometry, and blind mode, respectively.

$$PAT = \left[\frac{d_{auto} + d_{visod}}{d_{blind} + d_{auto} + d_{visod}} \right] * 100 \quad (6)$$

At the time of this writing, *Spirit* has traversed 2616 km blind, 1326.47 km autonomously, and 457.77 km using visual odometry for a PAT of 41%. *Opportunity* has traversed 3299.97 km blind, 1257.77 km autonomously, and 497.60 km using visual odometry for a PAT of 35%. Again, we make use of an average of the PAT for *Spirit* and *Opportunity* to designate a MER PAT of 38%.

The next metric is a simple time duration required for GESTALT to perceive local terrain, detect hazards, select a hazard-free path/direction, and execute a 35 cm step (nominally) along the selected path. This is the time required per autonomous navigation step. The *navigation decision time* for *Spirit* was 97 seconds on its 20 MHz RAD-6000 processor during its 90-sol prime mission (the duration increased with each patch

of improved flight software due to inclusion of additional onboard safety checks each time). For *Opportunity* the duration is approximately 1.5 times longer due in large part to use of its mast-mounted stereo Navigation Cameras, which have a higher resolution view of the terrain than the nominally-used, body-mounted stereo Hazard Cameras, which *Spirit* used. *Opportunity* used its Navigation Cameras because the higher resolution was required for good stereo correlation of near-textureless images of the smooth, uniform soil at Meridiani Planum. As such, there were more raw image data to process for local mapping and traversability assessment.

Oftentimes, high priority *in situ* science targets are located on portions of terrain that are difficult to traverse such as steep hills, soft soils, and excessively rocky/rough areas. In traverse attempts in such areas, the mobility system encounters reduced traction or high slip regimes during which onboard position estimates are severely compromised. Visual odometry is employed on occasion in such situations to best maintain position estimates by vision-based self-localization. As a generally more robust alternative to wheel odometry in rough natural terrain, it would be ideal to use visual odometry at all times. Its use on the MER vehicles was only commanded a small amount of time on average. Both *Spirit* and *Opportunity* traversed their share of rocky and soil-covered slopes on Mars as well as relatively flat and benign terrain. Thus far, the mean number of position self-localizations performed per traverse sol, when onboard visual odometry was enabled, is 12.1 for *Spirit* and 13.1 for *Opportunity*. To reduce the dependence of a self-localization usage metric on terrain types, we take an average of the commanded usage of visual odometry for both rovers.

Calculation of each autonomous navigation metric presented was done using telemetry and data products from *Spirit* and *Opportunity*. They are tabulated in Table 1 and represent baseline performance metric evaluations for MER autonomous navigation per the specific metrics as formulated above. Next, we present metrics for the operational functions of science target Approach and Instrument Placement.

Table 1. Calculated navigation metrics for MER

Performance Metric	Value
Autonomous traverse speed ratio	0.609
Percent autonomous traverse	38
Navigation decision time (secs)	97
Average self-localizations per sol	12.6

4.2 Approach and Instrument Placement

Primary among surface exploration mission goals is moving from place to place, and performing

measurements and investigations of a wide range of rocks and soils with *in situ* instruments. Therefore, an instrument positioning system (IPS) with the ability to perform precision placement of *in situ* instruments from mobile platforms is essential. The MER IPS is comprised of a robotic arm also known as the Instrument Deployment Device (IDD). This five degree-of-freedom mechanism includes a rotary turret as an end-effector to which science instruments are mounted. The purpose of the IDD is to place those instruments onto science targets within its kinematic work volume. Given a set of 3-D coordinates to reach, specified by commands in one or more sequences, it achieves this autonomously (including switching of instrument use) by realizing a combination of kinematic configurations that are pre-taught and/or newly-commanded. Cartesian- and joint-space motions are determined via onboard calculation of inverse kinematics and position error compensation (due to mechanical compliance of an as-built flexible link assembly and the effect of Mars' gravity given rover attitude). Automatic predictive collision-checking is also performed to validate contemplated arm motions; however, this software was run by IDD sequence developers on Earth for most of the operations on Mars. IDD (and mobility) operation and performance of both MER vehicles is documented in [5, 6].

An approach traverse refers to a one on the order of ten meters or less that is intended to terminate with a specific science target within the IDD work volume or in close proximity to the rover. The science target is selected and designated by mission operators in imagery acquired prior to the approach. A successful target approach is typically followed by placement of instruments onto the target using the IDD. MER operational guidelines required human confirmation of the rover position prior to each IDD use, making each approach and instrument placement take at least two sols, but future missions or software upgrades may make same-sol deployment possible. Depending on approach distance from a target, complexity of terrain between rover and target, and other considerations, target approach executions do not always succeed on first attempts. On occasion, more than one sol is needed to reach certain targets. One possible figure of merit for approach traverse performance considers distance to targets and the number of sols that were necessary to reach the targets. We employ such a measure as an average approach distance achieved, $d_{approach}$, for N targets per unit sol needed, n_{sols} , to successfully approach the N targets as follows.

$$Approachability = \frac{1}{N} \sum_{i=1}^N \frac{d_{approach}}{n_{sols}} \quad (7)$$

The approachability for *Spirit* and *Opportunity* during their 90-sol prime mission was determined for a small set of ten approach traverses and found to be 3.56 m/sol for *Spirit* and 2.29 m/sol for *Opportunity*. Based on that limited set of data an overall average approachability for MER is 2.93 m/sol.

Salient requirements for *in situ* IPS performance include dexterity, absolute positioning accuracy, and repeatability associated with placement of instruments on targets of interest including rock/soil targets and rover-mounted targets. Dexterity of a manipulator is characteristic of the mobility of any of its tools in all linear and angular directions. It refers to the motion a manipulator can achieve at various positions in its workspace and the ease with which the tool tip can move along a position and orientation in 3-D space. Positioning accuracy in this context refers to the arm's ability to position and orient its tools at a specified absolute location in its workspace. Repeatability can be defined in two ways, first as a characterization of error in positioning and orienting a tool when making incremental (small, relative) motions, and second, as the difference between the initial and final positions and orientations of a tool when moved back and forth between an initial position and a designated position. Several performance metrics for these attributes of IPS performance are presented. They are a small subset of many that could be applied to the functional steps involved in the approach and IP functional sequences.

For an *in situ* IPS, the dexterous workspace includes physical science targets on a planetary surface as well as rover-mounted targets that can be reached by all *in situ* instruments attached at the end-effector. Thus, the bigger the dexterous workspace the better, since the amount of fine rover positioning required to approach a science target would be minimized. In general, the larger the rover (footprint and height of the manipulator base above the ground plane) the more fine-positioning maneuvers it may require to best position the dexterous workspace (discounting mobility system steerability advantages such as provided by crab-steering that may minimize necessary maneuvers). As such, we define a *mobile-manipulability* metric as a ratio of dexterous workspace volume, v_{manip} , to mobility platform volume, v_{mobile} :

$$MobileManipulability = \frac{v_{manip}}{v_{mobile}} \quad (8)$$

For the MER vehicles, the IPS's nominal science target workspace is defined by a cylindrical volume, which is 0.5 m in diameter and 0.7 m in height and oriented vertically. The IDD is mounted towards the front of the rover and is capable of reaching out approximately 0.75 meters in front of the rover at full extent [7]. For

the purposes of this study, we consider the cylindrical volume of 0.14 m^3 as the MER dexterous workspace. With a wheelbase of 1.22 m, wheel-track of 1.06 m, and the IDD mounted at 0.43 m above the ground-plane we compute the mobility platform volume as 0.556 m^3 . Thus the mobile-manipulability metric for MER used in this study per Eq. 8 is 0.25.

The absolute positioning requirement for the MER IPS was split equally between the ability of the front Hazard Camera stereo pair to resolve the 3-D position and surface normal of a science target and the ability of the IDD to achieve a certain instrument position and orientation. As a result, the IDD was required to be capable of achieving a position accuracy of 5 mm and an angular accuracy of 5° in free space within its dexterous workspace. The IDD repeatability was required to be 4 mm in position and 3° in orientation. Performance results from *Spirit* and *Opportunity* during surface operations revealed an absolute positioning accuracy of 0.8 mm, and a repeatability of approximately 1 mm in position and 1° in orientation [7]. This performance is derived based on telemetry and stereo image range data evaluations of 422 placements of all instruments by *Spirit* and 439 placements of all instruments by *Opportunity* on rock, soil, and rover-mounted targets from sols 1 to 365 for both rovers [5]. Additional requirements were imposed that are not discussed here; see [7] for details.

For the purposes of this study, we define IPS performance metrics for absolute positioning accuracy and repeatability as ratios of required capability to actual demonstrated capability. For simplicity, we only consider the position component of each. The IPS accuracy ratio, IPS repeatability ratio, approachability, and mobile-manipulability metrics as computed for MER are tabulated in Table 2. They represent baseline performance metric evaluations for MER approach and instrument placement per these specific metrics.

Table 2. Computed approach and IP metrics for MER

Performance Metric	Value
Average approachability (m/sol)	2.93
Mobile-manipulability	0.25
IPS positioning accuracy ratio	6.25
IPS repeatability ratio	4.00

5. FUTURE ROVER PERFORMANCE

We will now apply the performance metrics defined in the previous section to a hypothetical future Mars rover. The performance values of this future rover will be used to demonstrate the technology assessment approach discussed earlier, and its utility for rover technology assessment relative to MER surface

robotics autonomy as a baseline. Illustrative examples will be presented in the next section. We consider a hypothetical future rover (HFR) with capabilities similar to tentative and projected performance capabilities for proposed future rovers. That is, for added realism, some relevant numbers and specifications are drawn from rovers planned for the next decade such as the MSL and ExoMars rovers [8, 9]. Other required performance data for the HFR are based on educated assumptions.

The ATSR metric for the HFR is based on average and maximum autonomous traverse speeds of 89 m/hour and 100 m/hour, respectively. Navigation decision time for MER is largely a function of processor speed and computing resources. Future rover processors will likely run at least an order of magnitude faster [4]. We make the linear assumption here that navigation decision time for the HFR would scale down linearly with processor speed. An HFR with a 133 MHz processor would then potentially support a navigation decision time of about 15 seconds (assuming an algorithm of similar complexity as GESTALT).

To prescribe a PAT metric for the HFR we assume that rovers designed for future missions will traverse more challenging terrain and tend to use self-localization more frequently (perhaps relying on it for effective autonomous navigation in complex terrain). As such, we assume that a greater percentage of traverses would be performed using visual odometry or a similar self-localization approach. We further assume that a greater percentage of traverses would be performed autonomously in general, and a lesser percentage blind. Relative to *Spirit*, a hypothetical 20% increase in autonomous traversal, 10% increase in use of self-localization, and 20% decrease in blind traverses would roughly equate to a PAT of 50 for the HFR. Based on the assumed 10% increase in use of self-localization, we project that the HFR might perform 10% more self-localizations on average than the MER rovers have thus far, that is, 13.9 per sol.

The HFR is assumed to be capable of approaching a science target and using its IPS to place an instrument onto the target within 3 sols from a distance of 20 m away, and thus, capable of an average approachability of 6.67 m/sol. With a platform footprint and manipulator dexterous workspace assumed to be 20% larger than those for MER, the HFR would have a mobile-manipulability of 0.3. Since the MER IPS positioning accuracy and repeatability are quite good and do not leave substantial room for improvement, we assume a 10% improvement in these capabilities for the HFR. Table 3 summarizes the performance of the HFR per the metrics defined in Sec. 4.

Table 3. Metrics for a hypothetical future rover

Performance Metric	Value
Autonomous traverse speed ratio	0.89
Percent autonomous traverse	50
Navigation decision time (secs)	15
Average self-localizations per sol	13.90
Average approachability (m/sol)	6.67
Mobile-manipulability	0.30
IPS positioning accuracy ratio	6.88
IPS repeatability ratio	4.40

6. RELATIVE TECHNOLOGY ASSESSMENT

The performance metrics for MER and the HFR will be used to demonstrate the technology assessment approach. As an example, we use the method to determine the impact of each system's autonomy on surface mission science return. The performance metrics enable propagation of technology impact scores, calculated at the functional level, upward to determine impact on science return (Fig. 2). At the Science Return Subclass Level (Level 4), the autonomy technologies of both MER and the HFR improve the quality, quantity, and diversity of science by enabling access to many and various high-priority science targets of opportunity. At the Operations Level (Level 3), technology components affect *Navigation*, *Approach* and *Instrument Placement* operations. Each of these autonomy technologies affects all of the functional steps associated with their respective operational functions (Level 2). Level 1 represents the technology impact score computed using Eq. 1. Given this construct, Tables 4 and 5 summarize the quantified assessment for these rover autonomy technologies for the HFR relative to MER as the SOA baseline. These tables are screen-shots of a spreadsheet tool designed to perform the calculations outlined in Sec. 3.

Table 4. Navigation technology assessment

PERFORMANCE PARAMETERS	Units	SOA value	Technology value	Comparison (S _T)
Inputs				
Processor Speed	MHz	20.00	133.00	-0.85
Outputs				
Autonomous traverse speed ratio	unitless	0.61	0.89	0.32
Percent autonomous traverse	unitless	38.00	50.00	0.24
Average self-localizations per sol	unitless	12.60	13.90	0.09
Environment				
No improvements	unitless			
Resources				
Navigation decision time	secs	97.00	15.00	5.47
Technology Impact Score, S ₁				1.32
Technology Impact Score, S ₂ (Functional)				1.32
Technology Level Impact Score, S ₃ (Operations)				0.44
Technology Level Impact Score, S ₄ (Science Return Subclass)				0.44
Technology Impact on Science Return				0.44

Table 5. Approach/IP technology assessment

PERFORMANCE PARAMETERS	Units	SOA value	Technology value	Comparison (S _T)
Inputs				
No dependencies				
Outputs				
Average approachability	m/sol	2.93	6.67	0.56
Mobile-manipulability	unitless	0.25	0.30	0.17
IPS positioning accuracy ratio	unitless	6.25	6.88	0.09
IPS repeatability ratio	unitless	4.00	4.40	0.09
Environment				
No improvements	unitless			
Resources				
No resource dependencies	unitless			
Technology Impact Score, S ₁				0.23
Technology Impact Score, S ₂ (Functional)				0.23
Technology Level Impact Score, S ₃ (Operations)				0.15
Technology Level Impact Score, S ₄ (Science Return Subclass)				0.15
Technology Impact on Science Return				0.15

An input for processor speed is incorporated in Table 4 since the navigation decision time performance metric (listed under Resources) is dependent on processing speed. From the bottom rows of Tables 4 and 5 we observe that the HFR autonomy technology provides a higher relative impact on science return than the MER autonomy technology. The relative increase in impact on science return is more substantial for the Navigation technology than for the Approach and Instrument Placement technology considered for both rovers.

The methodology outlined above also allows one to suggest what level of technology capabilities would be needed beyond the SOA by a future system to achieve a desired science return potential. For example, suppose we were interested in a measure of how much advancement in Approach and Instrument Placement technology might be necessary beyond MER to achieve a 50% relative increase in science return potential as compared to the HFR technology of Table 5. For Approach and Instrument Placement, the metrics that we would desire possible improvements in are the average approachability and mobile-manipulability. We can explore “what if” scenarios by increasing these values in various combinations within the spreadsheet to yield the desired relative increase. For this example, an HFR with an average approachability of 7.0 and mobile-manipulability of 0.6 as documented in Table 6 should achieve the desired increase in science return potential.

Table 6. Alternate approach/IP technology assessment

PERFORMANCE PARAMETERS	Units	SOA value	Technology value	Comparison (S _T)
Inputs				
No dependencies				
Outputs				
Average approachability	m/sol	2.93	7.00	0.58
Mobile-manipulability	unitless	0.25	0.60	0.58
IPS positioning accuracy ratio	unitless	6.25	6.88	0.09
IPS repeatability ratio	unitless	4.00	4.40	0.09
Environment				
No improvements	unitless			
Resources				
No resource dependencies	unitless			
Technology Impact Score, S ₁				0.34
Technology Impact Score, S ₂ (Functional)				0.34
Technology Level Impact Score, S ₃ (Operations)				0.22
Technology Level Impact Score, S ₄ (Science Return Subclass)				0.22
Technology Impact on Science Return				0.225

These examples have not made use of metrics related to environmental effects on exploration performance. Such effects can be factored into the assessment by incorporating appropriate metrics in the Environment category of each table. An example of a relevant environment metric is one that relates to accessibility of science targets on the terrain. This would apply in a scenario where the most interesting science targets are concentrated on steep slopes only, in which case the rover’s capabilities to climb and accurately navigate on slopes would significantly impact science return.

Such “what if” scenarios could be useful to technology program managers and technologists for supporting claims of expected impact of new technologies on future missions. They may also be useful to flight systems engineers for making more educated trades between existing (heritage) technology and more recent technology developments relevant to their missions. These examples hint at the utility of the technology assessment methodology given MER as a baseline.

7. SUMMARY AND CONCLUSIONS

An overview of a systematic rover technology assessment methodology was presented with associated performance metrics. Autonomy performance values were used to apply the method and assess related technologies relative to MER as a baseline. Illustrative examples demonstrated the application and utility of the technology assessment methodology for flight rover autonomy. Navigation, approach, and instrument placement autonomy are considered here; however, additional autonomy technologies may be factored into the assessment given suitable metrics.

Development of a formulation that is more generally applicable to many rover systems would warrant deeper and more thorough study than reflected here. Among other forms of added rigor, a more informed technology assessment application would include a more complete set of performance metrics, success probabilities from risk models, and would factor in science instrument types and measurement capabilities.

This work will assist rover autonomy technologists in quantifying scientific benefits that their technology brings to a mission, thus providing a means to validate, to mission designers and systems engineers, the need to infuse autonomy technology into mission capabilities. It also provides a foundation for examining related “what if” scenarios that may aid in decision making for technology sponsors and flight systems engineers.

8. REFERENCES

1. Lincoln, W.P., Elfes, A., Huntsberger, T., Rodriguez, G. and Weisbin, C.R., Relative Benefits of Potential Autonomy Technology Investments, *Intl. Conf. on Space Mission Challenges for Information Technology*, Pasadena, CA, July 2003.
2. Howard A. and Rodriguez G., Validating Mission Relevance of Autonomy Technologies through Increased Science Return, *Intl. Conf. on Machine Learning, Workshop on Machine Learning in Space Systems*, Washington, D.C., 2003.
3. Weisbin C.R., Rodriguez G., Elfes, A. and Smith, J.A., Toward a Systematic Approach for Selection of NASA Technology Portfolios, *Systems Engineering Journal*, Vol. 7, No. 4, 285-301, 2004.
4. Goldberg, S.B., Maimone, M.W. and Matthies, L., Stereo Vision and Rover Navigation Software for Planetary Exploration, *IEEE Aerospace Conf*, Big Sky, MT, 2002.
5. Trebi-Ollennu, A. Baumgartner, E.T., Leger, C. and Bonitz, R.G., Robotic Arm In-Situ Operations for the Mars Exploration Rovers Surface Mission, *IEEE Intl. Conf. on Systems, Man, and Cybernetics*, October 2005, Waikoloa HI, (in press).
6. Tunstel, E. et al. Mars Exploration Rover Mobility and Robotic Arm Operational Performance, *IEEE Intl. Conf. on Systems, Man, and Cybernetics*, October 2005, Waikoloa HI, (in press).
7. Baumgartner, E.T. et al. The Mars Exploration Rover Instrument Positioning System, *IEEE Aerospace Conf*, Big Sky, MT, March 2005.
8. Heninger, R., Mars Science Laboratory Mission 2009 Landed Science Payload, Proposal Information Package, JPL Document D-27202, MSL-225-0502, NASA/JPL/Caltech, Apr.14, 2004.
9. ExoMars09, CDF Study Report: CDF-14(A), ESTEC Concurrent Design Facility (CDF), ESA, August 2002.