

Smart Robots on Mars: Deciding Where to Go and What to See

Kiri L. Wagstaff

September 16, 2008

INTRODUCTION

The planet Mars has much in common with the Earth, including a similar rocky composition and familiar terrain features such as volcanoes, cliffs, valleys, buttes, and so on. It has an active (if thin) atmosphere, with clouds, ice, dust devils, and storm systems, and its frozen polar caps advance and retreat with the seasons. Yet Mars also differs from the Earth in many ways. It is 50% further from the Sun, with a correspondingly longer year (687 Earth days). The Martian day is slightly longer than Earth's, at 24 hours and 40 minutes. Mars itself is only half the size of Earth, with 1/10th of the mass, resulting in much lower gravity on the surface (about 1/3 Earth's gravity). The atmosphere is primarily composed of CO₂ rather than nitrogen and oxygen, and it is much thinner, with only 1/100th the pressure of Earth's atmosphere. As a result, the daily temperature swings are much more dramatic, from 23 F down to -125 F at night.

Because of these tantalizing similarities and curious differences, Mars is the subject of intense study and exploration. We seek to understand how it was formed, whether it might once have harbored life, and what caused it to reach its currently inhospitable state. It is not known conclusively whether life may exist on Mars today, but if it does, it would likely be hidden away in a subsurface haven heated by hydrothermal activity, or otherwise ensconced in a protected niche. To help us learn about and explore Mars, we have dispatched several robots—orbiters, landers, and rovers—to collect observations from orbit and on the ground.

Mars is so distant (100-200 million miles, depending on the positions of the planets in their orbits) that directly controlling these spacecraft from Earth is impossible. Much of the extensive exploration that has already been conducted has been achieved by careful pre-scripting of spacecraft activities, usually weeks in advance for orbiters and a day in advance for rovers. However, a fundamental property of exploration

is that we commonly encounter the unexpected. The Mars Exploration Rovers have the ability to drive farther in one day than can be observed, and planned for, the day before. It has become increasingly evident that relying on pre-scripted activity imposes significant limitations on what can be achieved. Therefore, several innovative methods have been developed to enable autonomous exploration and decision making for our robots on Mars.

EXPLORING MARS ON THE GROUND

The twin Mars Exploration Rovers landed on Mars in January 2004. With an original surface mission timeline of 90 days, some people are surprised to learn that the rovers are still going, five and a half years later. To be sure, they have experienced some degradation in capability: Spirit now drives backwards while dragging its right front wheel, which is no longer functioning. At the time this article was written, Spirit wasn't doing any driving at all, since it was stuck in a deposit of soft sand. Opportunity has had multiple problems with its robotic arm, including a stuck heater switch and stalls in the shoulder joint. However, Opportunity passed the 10-mile mark in terms of total distance covered in May 2009, and new information gained from both rovers has greatly impacted our understanding of the history of water on Mars.




Figure 1. Mars Exploration Rovers, artist's conception (left). Sand grains imaged by Spirit (middle). Opportunity's tracks leaving Victoria Crater, with the Cape Verde promontory visible (right). Images courtesy of NASA/JPL-Caltech/Cornell University.

It is on the ground that autonomy is most needed by remote exploration vehicles. Once a spacecraft attains its target orbit, the environment is for the most part stable and does not present many surprises or changes. Challenges do occur, such as cosmic radiation (which can corrupt memory and computation) and even erroneous commands sent by human operators on Earth, which led to the demise of the Mars Global Surveyor orbiter in late 2006. However, orbiters do not face the continual need to navigate their environment and respond to new information and obstacles.


The key components of any exploration, whether a family trip to Disneyland or a robotic investigation of Mars, are first, deciding where to go (and how to get there safely), and second, determining what pictures to take and what to send home. Humans are natural explorers, driven by curiosity; rovers must be provided with exploration goals and guidance as to how to achieve them. Like humans on a family vacation, the rovers have a limited time to explore a given area, and they additionally must operate under tight limits in terms of how much data they can store and how much can be transmitted back to Earth. Our goal is to enable rovers to make decisions about mobility and data collection efficiently and accurately, without waiting for the round-trip communication time to Earth for every decision that must be made.

DECIDING WHERE TO GO

The Mars Exploration Rover program employs several people who work as “rover drivers,” but this title is something of a misnomer. With a round-trip communication time that varies from ten to 40 minutes, driving the rovers in real-time is impossible. Instead, rover drivers assemble scripted driving sequences that are uploaded to the rover for the next day’s traverse. They create these sequences using images and other data collected by the rover in its current location as well as contextual views of the region collected by orbiters. By carefully tracking and avoiding visible obstacles, rover drivers can lay out a safe plan for the rover’s initial work the following day.


However, both rovers are capable of driving further in a single day than can be safely envisioned and planned for using the previous day’s images. They therefore use a combination of two driving modes. The first, blind driving, involves the rover simply executing a pre-planned sequence of drives and turns. An example would be “drive forward 4.5 meters, turn right 37 degrees, drive forward 2.7 meters, and stop.” The rover does not look where it is going but instead trusts that the sequence has been correctly assembled and tested. The second mode is autonomous navigation, in which the rover takes an image of the terrain in front of it, determines how to safely navigate it, moves forward about one meter, and then repeats the process. Blind driving is, of course, much more efficient than the stop-and-go process of autonomous navigation, although neither one is fast by human standards. Blind driving has a top speed of 5 cm/sec, and typically operates closer to 3.75 cm/sec to avoid consuming too much power¹. Autonomous navigation, in addition to frequent stops, is also slowed by the time it takes to process each image and compute a safe path; it generally achieves speeds of less than 1 cm/sec¹. However, since autonomous navigation can dramatically extend the range of the rover for a given day, rover drivers generally do use it as a fallback after the pre-planned blind driving sequence runs out.

Autonomous navigation depends on three key technologies: real-time stereo analysis of navigation images, traversability map computation, and obstacle avoidance. Stereo analysis of image pairs allows the rover to determine the 3D locations of rocks, ridges, and other features of its environment. Below is the stereo information that the Spirit rover computed for a rock called Humphrey that is 50 cm tall. It takes the rovers about 30 seconds to process each stereo pair¹.



**Figure 2. Stereo results from the Spirit navigation cameras, looking at Humphrey rock in Gusev Crater.
From Maimone et al., 2004¹.**

The rover next computes a traversability map that identifies which areas are safe to drive through and which are not. The area around the rover is divided into a grid. Given the stereo information, each grid cell is classified as safe, questionable, or dangerous (shown as green, yellow, or red below, respectively). The GESTALT system¹ onboard the rover considers 48 candidate trajectories that the rover may follow, including 23 that move forward, 23 that move backward, and two in-place turns (clockwise and counter-clockwise). The forward trajectories are shown below in white. GESTALT selects the trajectory that is safest, remaining in green grid cells as much as possible, while also moving the rover toward its next waypoint (as specified by the rover drivers). It takes the rover about 70 seconds to analyze all 48 trajectories and select the best one¹. The rover follows the selected trajectory about one meter, then stops and repeats the analysis and trajectory selection procedure. In 2007, the rovers received an updated trajectory analysis procedure called Field D-Star, which permits the automated planning of drives as much as 50 meters long, depending on visibility.



**Figure 3. GESTALT candidate trajectories and traversability map, from Maimone et al., 2004¹.
Opportunity looking back at an autonomous 15.8-meter drive conducted with Field D-Star onboard navigation; image courtesy of NASA/JPL-Caltech/Cornell University (right).**

More sophisticated methods for assessing traversability, which use machine learning to develop trained models, are currently under development². Examples to train the model can be generated offline by a human operator who labels regions within a navigation image as traversable or not. It can also be done online through teleoperation, in which anything successfully driven over becomes a traversable example, and regions where the operator prematurely stopped the rover from driving become examples of obstacles.

DECIDING WHAT TO SEND BACK HOME

The purpose of the rovers' exploration is to collect and send back observations to enable scientists to better understand the Martian environment. The rovers have a variety of instruments that allow them to collect images (panoramic, navigation, and microscopic), and spectrometer readings (thermal emission, Mössbauer, and alpha particle x-rays). However, the rovers have only a limited amount of communication bandwidth available for transmitting data back to the Earth. All missions beyond Earth orbit communicate with the Earth through the Deep Space Network, which consists of telescopes positioned at three main locations around the globe: Canberra, Australia; Goldstone, California; and Madrid, Spain. Telescope transmitter and receiver time must currently be split between the competing needs of missions such as New Horizons (on its way to Pluto), Cassini (at Saturn), Mars Reconnaissance Orbiter, Mars Express, the Mars rovers, MESSENGER (at Mercury), and more. Like most missions, the rovers have the capability to collect far more data than they can transmit. This is especially true for images. Therefore, it is often necessary to carefully select which data to collect, and which data to transmit.

One way to solve the problem of deciding what to send is to carefully pre-plan each data collection activity (images or other instrument readings). Since the bandwidth allocation is known in advance, mission operators can create a schedule that will fill up, but not exceed, the available capacity. However, as with pre-planned driving, this approach does not permit any exploration of the unknown, or a reaction to an unexpected observation. If we knew ahead of time exactly when and where to point our cameras, then little exploration would be needed. In the new environments we find on Mars, however, there are many surprises. Therefore, missions are now looking at ways to permit the rover itself to take far more images than it can transmit, then analyze the data and make decisions about how to rank the images for download.

Dust devils have provided a strong motivation for this kind of operation. To passively capture images of the whirling vortices, initially the Mars rovers reserved dedicated time to sit stationary and do nothing but

take a sequence of images of the horizon, hoping that one or more might contain a dust devil. The rover then transmitted all of the images to Earth, and a mission scientist examined each one to identify the lucky finds. Of course, all of the “empty” images took up both downlink bandwidth and mission time, since the rover could not conduct any other activities while collecting the images.




Figure 4. Dust devil imaged by Spirit, August 2005 (left). Dust devil imaged by Phoenix, September 2008 (right). Images courtesy of NASA/JPL-Caltech/Cornell University.

Therefore, a technique has been developed to detect motion inside images, inferred to indicate dust devils or other interesting phenomena, and flag those images as high-priority for download³. The algorithm was uploaded to the Mars rovers in 2007. Over 26 sols (Martian days), it detected dust devils in 30% of the images that were collected. Restricting download to only those images could save up to 70% of the current bandwidth consumption. The onboard analysis continues to be used by both rovers when dust devils are in season.

Dust devils can be detected in other ways, too. Cameras are good for detecting distant dust devils. However, sometimes the dust devils cross over the rover itself (often with the beneficial side effect of clearing dust from the solar panels). In those cases, the rover can detect the event simply by monitoring the current readings of an atmospheric pressure sensor. The dust devil causes a local, temporary drop in pressure. The Mars Pathfinder lander had such a sensor, and in the 83 sols of the mission, 79 crossing dust devils were detected in later analysis of the data. If continuous data collection had been employed, it is estimated that 210 to 349 dust devils would have been observed during that same time period. We adapted a post-analysis technique for use in an onboard, online setting, and applied it to the Mars Pathfinder data⁴. The Mars Exploration Rovers did not have atmospheric pressure sensors, so they could not employ this technique. However, the Mars Science Laboratory rover, scheduled for launch in 2011,

will have such a sensor and could employ onboard detection of crossing dust devils to trigger a follow-up image to be taken by a camera, possibly capturing the dust devil as it departs.

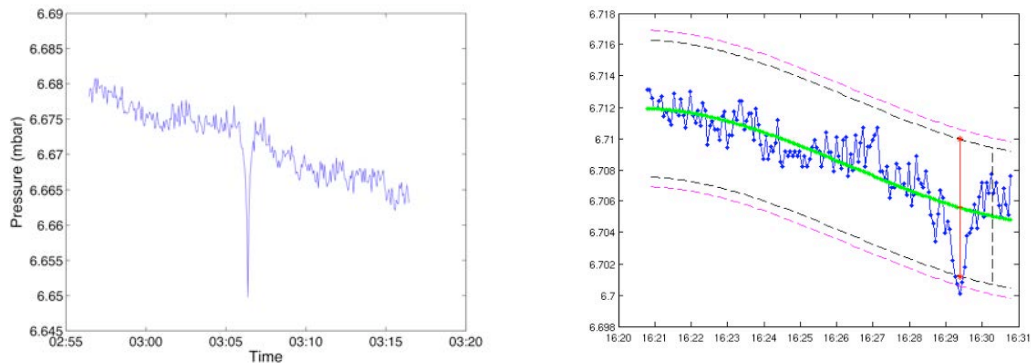


Figure 5. Sample dust devil (drop in pressure) observed in Mars Pathfinder atmospheric sensor data (left). Realtime analysis of the atmospheric sensor data, showing the moving trend (green) and a successful dust devil detection (red), from Wagstaff et al.⁴ (right).

In the future, we anticipate that rovers would be able to do more than just decide which data to send home; they might also make decisions about which images to collect. The Onboard Autonomous Science Investigation System (OASIS) could both prioritize data and issue “science alerts” onboard the rover to recommend a deviation from the pre-planned trajectory. This is useful if a particular type of rock, previously designated as high-priority (such as a carbonate), is detected while the rover is en route to another location. The rover could respond to the OASIS alert by stopping to take an additional image of the unusual rock and then proceeding on its pre-planned course⁵. The rover would take into consideration resource constraints before approving the deviation. Further, the same technology could be used to aid instruments that require precise pointing by selecting individual targets.

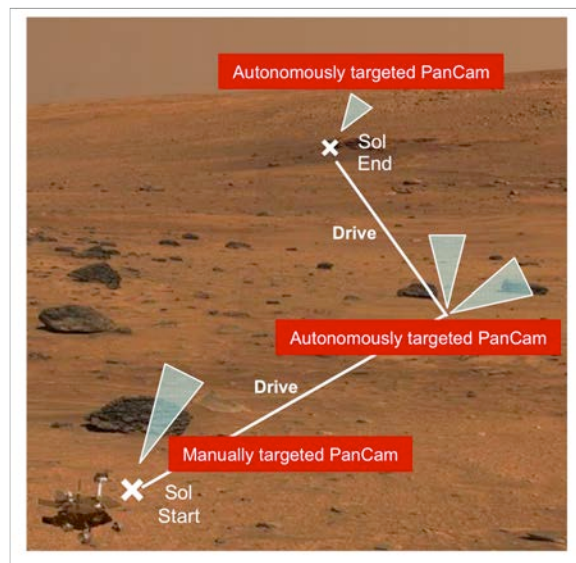


Figure 6. AEGIS concept for autonomous target selection in support of the PanCam imager on a Mars rover, from Estlin⁶.

The AEGIS (Autonomous Exploration for Gathering Increased Science) software currently under development would allow rovers to detect interesting or unusual rocks and then deploy spectrometers or thermal imagers to acquire additional information about those rocks. This software will be uploaded to the Mars rovers this year in a testing phase, and ultimately it is planned for inclusion on the Mars Science Laboratory rover in 2011⁶.

SUMMARY

The “robots” we have on Mars (our landers, rovers, and orbiters) have provided one insight after another about this planet, which is simultaneously similar and yet very different from the Earth. In exploring new environments, they face two main challenges: deciding where to go (and how to do it safely) and what data to collect and transmit back to Earth. Operating under time, power, memory, CPU, and bandwidth constraints, these robots rely more and more on autonomous operation, both in terms of navigation and in onboard data analysis. Advances in onboard autonomy are what could enable future missions to make the most of their limited resources and time, and help answer the remaining big questions about Mars.

ACKNOWLEDGMENTS

I wish to thank all of the researchers and mission operators for their hard work and innovative developments that have enabled the astonishing discoveries we’ve already found on Mars, as well as the future ones to come. This paper was written at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2009, California Institute of Technology. Government sponsorship acknowledged.

NOTES

1. Mark Maimone, Andrew Johnson, Yang Cheng, Reg Wilson, and Larry Matthies, “Autonomous navigation results from the Mars Exploration Rover (MER) mission,” *Proceedings of the 9th International Symposium on Experimental Robotics*, (2004): 3-13.
2. Max Bajracharya, Andrew Howard, Larry H. Matthies, Benyang Tang, and Michael Turmon, “Autonomous off-road navigation with end-to-end learning for the LAGR program,” *Journal of Field Robotics*, 26 (2008): 3-25.
3. Andres Castaño, Alex Fukunaga, Jeffrey Biesiadecki, Lynn Neakrase, Patrick Whelley, Ronald Greeley, Mark Lemmon, Rebecca Castaño, and Steve Chien, “Automatic detection of dust devils and clouds on Mars,” *Machine Vision and Applications*, 19 (2008): 467-482.
4. Kiri L. Wagstaff, Rebecca Castaño, and Steve Chien, “Real-time in-situ detection of dust devils using pressure sensors,” *Eos Transactions of the American Geophysical Union*, 87 (2006): Abstract #IN52A-03.
5. Rebecca Castaño, Tara Estlin, Dan Gaines, Ben Bornstein, Robert C. Anderson, Brian Bue, and Michele Judd, “Experiments in onboard rover traverse science,” *Proceedings of the IEEE Aerospace Conference* (2008).

6. Tara Estlin, “Automated targeting for the MER rovers,” *Proceedings of the 3rd IEEE International Conference on Space Mission Challenges for Information Technology* (2009).

Jet Propulsion Laboratory, California Institute of Technology, kiri.wagstaff@jpl.nasa.gov, (Depts. of IT and Communications)