


 NASA Anniversary

Autonomy for Mars Rovers: Past, Present, and Future

Max Bajracharya, Mark W. Maimone, and Daniel Helmick
 Jet Propulsion Laboratory, California Institute of Technology

The vehicles used to explore the Martian surface require a high degree of autonomy to navigate challenging and unknown terrain, investigate targets, and detect scientific events. Increased autonomy will be critical to the success of future missions.

In July 1997, as part of NASA's Mars Pathfinder mission, the *Sojourner* rover became the first spacecraft to autonomously drive on another planet. The twin Mars Exploration Rovers (MER) vehicles landed in January 2004, and after four years *Spirit* had driven more than four miles and *Opportunity* more than seven miles—lasting well past their projected three-month lifetime and expected distances traveled. The newest member of the Mars rover family will have the ability to autonomously approach and inspect a target and automatically detect interesting scientific events. In fall 2009, NASA plans to launch the Mars Science Laboratory (MSL) rover, with a primary mission of two years of surface exploration and the ability to acquire and process rock samples. Figure 1 shows mockups of all three rovers.

In the near future, the Mars Sample Return (MSR) mission, a cooperative project of NASA and the European Space Agency, will likely use a lightweight rover to drive out and collect samples and bring them back to an Earth return vehicle. This rover will use an unprecedented level of autonomy because of the limited lifetime of a return rocket on the Martian surface and the desire to obtain samples from distant crater walls.

AUTONOMY

All spacecraft must utilize a high degree of autonomy due to their low-bandwidth, high-latency communication channels to Earth. Traditional autonomous capabilities include responding to faults, pointing antennas, orientation control using star tracking, and onboard data storage and retransmission.

Spacecraft used to explore space or orbit planetary bodies travel through a relatively simple environment and need only perform navigation corrections tens to hundreds of times during their multiyear missions. However, the rovers used to explore the surface of Mars must constantly interact with challenging and unknown terrain. For example, in their first three years, the MER vehicles performed some 60,000 coordinated motions—powering either steering or drive motors continuously.

Autonomy software allows a rover to make decisions and command actuators based on its observations of the environment or sensor feedback. Without it, human operators must verify the state of the vehicle before and after it performs any action.

High-bandwidth, low-latency communications make direct teleoperation of a vehicle possible. Unfortunately, the round-trip communication delay between Earth and Mars ranges from 8 to 42 minutes for MER, and communication using the Deep Space Network antennas is only available several times during a Martian solar day (sol). Consequently, the operations team must plan and uplink an entire sol of vehicle operations and then wait for the vehicle to execute it without any human monitoring or confirmation. Increased autonomy will be critical to the success of future missions, in which rovers will be expected to travel over long distances in a short period of time and handle dynamic processes such as taking a core sample from a rock while slipping on a slope.

In addition to the communications limitations, exploring the surface of another planet presents many mobility and sensing challenges. A rover must sense the terrain in widely varying lighting conditions, interact with terrain

that might not be fully characterized, and account for uncertainties in sensing or system faults.

Further, planetary vehicles have extremely limited computational resources due to the high radiation levels and large temperature changes of space. *Sojourner* had a 0.1-MHz Intel 80C85 CPU with 512 Kbytes of RAM and 176 Kbytes of flash memory, the MER vehicles have a 20-MHz RAD6000 CPU with 128 Mbytes of RAM and 256 Mbytes of flash memory, and the MSL vehicle will have a 200-MHz RAD750 PowerPC with 256 Mbytes of RAM and 512 Mbytes of flash memory. The MER and MSL vehicles use the VxWorks operating system and run many parallel tasks continuously, leaving autonomy software with less than 75 percent of the available CPU time.

SOJOURNER

Sojourner's autonomous capabilities consisted of terrain navigation, contingency response, and resource management.¹ As Figure 2a shows, the rover could autonomously navigate through flat but rocky Martian terrain to position locations specified by Earth-based operators. It used stereo cameras and five infrared laser stripers to detect hazards. Each of the cameras recorded laser spots using four scanlines, providing 20 terrain height measurements. If the vehicle detected a hazard, it would turn in place until the hazard was no longer visible, drive forward past it, and then resume driving to its waypoint.

The vehicle did not build or maintain a permanent terrain map and was consequently a purely reactive system. It did, however, have specific behaviors such as a thread-the-needle action that could be triggered in certain situations. It also had contingency capabilities such as reacting to contact bump sensors or not reaching a destination. Other commands, such as "find rock," let the vehicle servo to a rock target, as Figure 2b shows.

MARS EXPLORATION ROVERS

Compared to *Sojourner*, the MER vehicles have a significantly more sophisticated suite of autonomous capabilities that they used even during the first 90 sols comprising the Primary Mission. Updates made to the vehicles' software during the Extended Mission incrementally enhanced

these capabilities, as well as added new ones. The initial capabilities included safe terrain navigation while avoiding geometric hazards, visual pose estimation, and absolute orientation sensing. The upgrades included a global path planner for improved navigation and new abilities to autonomously approach and place an instrument on a target and automatically detect science events.²

Terrain navigation

The MER vehicles use stereo camera pairs mounted on the body and on a pointable mast as the primary sensors for most autonomous surface capabilities. Unlike *Sojourner*, the MER vehicles perform full stereo processing of the camera pairs, providing a dense 3D reconstruction of the terrain geometry. While *Sojourner* only sensed 20 points in 3D at each navigation step, the MER vehicles sense 15,000 to 40,000 3D points from each pair of stereo images.

To navigate to operator-defined waypoints, the vehicles detect geometric hazards from the stereo data and then choose the best path to execute.³ They detect hazards by fitting rover-sized patches of 3D stereo data to a plane and then looking for deviations from the plane and



Figure 1. Mockups of the current family of Mars rovers: (center) 1997 Mars Pathfinder vehicle, (left) 2004 Mars Exploration Rover vehicle, and (right) 2009 Mars Science Laboratory vehicle.




Figure 2. The *Sojourner* rover (a) driving to the Mermaid dunes and (b) inspecting the Yogi rock.




Figure 3. Example terrain assessment and path selection by MER vehicle. The terrain model is built up from stereo imagery. The colored grid represents the traversability map, and the white lines show the considered paths.

its slope. The vehicles use the hazards to populate a surrounding map of terrain traversability; from this map, they select the best motion command to execute, as Figure 3 shows. The motions consist of arcs or turn-in-place maneuvers, and their cost is calculated by integrating the traversability along the path, the steering and drive motor rotations, and the end distance to the goal. The vehicles repeat the process of detecting hazards and then executing a motion until they reach the goal.

The initial version of the navigation software could only plan one or two motions ahead. This was sufficient to navigate around small obstacles and has helped keep both rovers safe through nearly five years of driving. It enabled *Spirit* to reach the Columbia Hills 50 percent faster because the rover could be commanded to continue driving in previously unseen areas, as long as enough power was available. However, the system would occasionally fail to make forward progress when the vehicle encountered a larger, extended obstacle. JPL scientists addressed this problem by adding a global path planner, which made it possible to maintain a much larger map and plan optimal paths through it.⁴

Visual pose estimation

To maintain a local navigation map, the MER vehicles must accurately estimate their relative pose between navigation steps. They do this by combining

- the wheel odometry, which provides a measure of how far each wheel has traveled;
- an inertial measurement unit (IMU), which provides linear accelerations and angular velocities; and
- optional vision-based pose estimates.

On flat, hard terrain, the wheel odometry and IMU are generally sufficient to estimate vehicle pose. However, because wheel odometry cannot account for wheel

slip, the rovers must use visual techniques in sloped, sandy, or rocky terrain.

Visual pose estimation, or visual odometry, autonomously detects and tracks scientific features in sequential stereo camera pair images, then uses the 3D motion between the features to update the pose estimate.⁵ It enables precise middrive imaging of a nearby feature, speeds up driving to a feature in difficult terrain by enabling intelligent autonomous replanning, and greatly enhances vehicle safety by performing slip checks and implementing human-specified “keep-out zones.” Due to its high computational cost, visual odometry is only used when necessary.

Absolute orientation sensing

To detect their absolute heading on the surface, the MER rovers use their pointable mast cameras to detect the sun’s location and then use information about the solar time and the direction of gravity measured by the IMU to compute their orientation.

Target investigation

One of MER’s primary goals was to use scientific instruments mounted to each rover’s robotic arm to investigate and analyze Martian terrain. This requires the rovers to drive up to a designated target, position themselves to reach the target, and deploy the arm onto the target to perform the investigation. During most of the mission, this was a very manual process, requiring confirmation from a human operator at each step to ensure the safety of the vehicle, arm, and instruments.

Later in the mission, JPL scientists added new autonomous capabilities to the system to reduce the number of sols required to reach and inspect a target. As Figure 4 shows, these included the rover’s ability to track a target while approaching it, automatically determine the best configuration to position itself, safely deploy the arm and inspect a target without human confirmation, and use visual feedback to correct for arm positioning errors.

Target tracking. The MER rovers can use their pointable mast-mounted engineering cameras to track a target while approaching it.⁶ Visual target tracking lets the vehicles servo to a designated target in a closed-loop fashion rather than simply driving to a specified waypoint, which might not be possible due to slip. Unlike visual odometry, which tracks dozens to hundreds of automatically detected features, visual target tracking focuses on a single designated target and can consequently execute much faster.

Rover positioning. When a rover reaches the target’s vicinity, it must place itself so that the arm can reach the target and place the instrument. This requires accounting for the terrain in front of the target as well as the surface geometry of the target itself. Using a terrain and target model built from the stereo cameras, the rover can




Figure 4. Autonomous capabilities of MER vehicles during target investigation: (a) features tracked during visual odometry; (b) initial and final image of a target tracked during approach; (c) vehicle and terrain model used for collision detection during arm deployment; (d) example of using visual feedback to correct for arm position error.

automatically determine the optimal position by searching over all positions to maximize the target's manipulability.⁷

Arm deployment. The MER rovers use the terrain model, along with a known model of the vehicle, to predict collisions with the terrain during arm deployment. They also can correct arm positioning errors using hand-eye coordination.⁸ The rovers can use their stereo camera to detect a marker on the end of the arm; visually compute the arm's position offline; compare this to the desired position, which might not have been achieved due to gravity sagging of the arm or positioning control error; and correct the final position.

Science event detection

Later in the MER mission, the rovers acquired the ability to opportunistically tag scientific observations from their mast cameras. This allowed them to automatically

detect interesting, rare, and dynamic scientific events, such as dust devils or cloud formations, and then send these observations back to Earth.⁹ Previously, JPL scientists had to predict when an event might happen, take many images, and send all of them back to Earth even though the event would only be in a few of them—an inefficient use of communication resources. The latest software enables the rover to detect the presence of dust devils by processing and comparing images taken over a short period of time, and to detect the presence of clouds by locating the horizon and finding high-intensity variance regions in the sky.

FUTURE ROVERS

The MER vehicles' autonomous capabilities have enabled them to explore more of the Martian surface than was ever imagined, and the upcoming Mars Science Laboratory rovers will rely on these capabilities to




Figure 5. Terrain prediction. (a) Predicted terrain classes from a research rover panorama of the JPL Mars Yard. The color overlay indicates the terrain class. (b) The rover is traversing the terrain after using the predicted terrain classes to compute the path with the least slip.

ensure the same level of performance. However, future missions will require more autonomy to achieve even minimal success.

The Mars Sample Return baseline scenario is to land an ascent vehicle along with a rover to obtain a diverse set of samples from the surface and place them into Mars's orbit so that an orbiter can collect them and return to Earth. Unfortunately, thermal cycling of its fuel limits the ascent vehicle's lifetime on the surface to about one year. Consequently, the MSR rover must be able to drive to many locations up to 5 km from the landing site in a very short period of time. Further, the most interesting and informative samples are believed to be on crater walls or steep slopes, and given the rover's need to be as light as possible to offset the ascent vehicle's size, the rover must be able to core into hard rock from a potentially unstable position.

Terrain prediction

Navigating in rough or high-slip terrain involves highly complex interactions between the vehicle and terrain. Because these interactions are both difficult to detect and model, deciding how to navigate is difficult. For instance, on a sandy slope, the way a vehicle drives and the type of sand affect how much and in what direction the vehicle slips, but because it is sinking and sliding in loose material, predicting the actual motion is difficult.

One way to handle this problem is to endow the vehicle with the ability to automatically learn properties

of the terrain while driving over it and then use this knowledge to predict and react to new terrain.¹⁰ As Figure 5a shows, the vehicle can learn terrain properties by sensing the amount and direction of slip it experiences and correlating this to the terrain's appearance and geometry. This "self-supervised" approach lets the vehicle predict the slip of new terrain it can see but has not yet traversed.

Combining knowledge of the terrain properties with an accurate vehicle model, a dynamics engine can predict how the vehicle will react on this terrain and choose the best path to navigate it, as Figure 5b shows. The vehicle's ability to learn about the environment also allows it to react to terrain or obstacles it previously could not predict—for instance, a crust of hard soil with soft soil below it that might cause the vehicle to sink.

Navigation speed

To increase overall navigation speed during an MSR mission, the rover can take advantage of the fact that it will always return to the lander to deliver samples to the ascent vehicle. This lets the rover use knowledge obtained during a traverse out to a sampling site on the return traverse. If the vehicle has safely navigated to a location, it can remember the route it took and then return by the same path. However, to do this, the vehicle must be able to detect the path or localize itself with respect to the path. It can achieve this by either directly detecting the path from imagery or remembering a map of features in the environment and then matching these features during the return traverse.¹¹

Autonomous science

Learning can also be used to increase autonomous science. As a vehicle explores an area and takes images or scientific measurements, it can learn what measurements to expect. Once it has learned the nominal expected measurements, it can automatically detect interesting or off-nominal measurements.¹² Unlike cloud and dust devil detection, which the vehicle is specifically programmed to detect, this capability would let the vehicle automatically detect novel science events from many different instruments.

Core sampling

One of MSR's goals is to take rock core samples from various locations on crater walls. This will require a light-

weight rover about the size of the MER vehicles to drill several centimeters into outcropped rocks in potentially sandy terrain. To prevent the coring tool from binding or breaking, which would be mission-ending, the MSR vehicle must be able to compensate for inevitable shifts and slips while the tool is embedded in a hard rock.

Autonomy is the only possible solution to this problem. The vehicle must sense the slip, using its cameras to detect visual motion or force sensors in the coring tool to sense binding or undesired torques. Using these sensory inputs, the vehicle can then use both arm and vehicle motions to adjust the tool to compensate for the motion,¹³ as Figure 6 shows.

All of these future capabilities require more computational resources than are currently available for planetary rovers. While some improvement in radiation-hardened CPU speeds and memory can be expected, some constraints will continue to be a challenge for increasing vehicle autonomy. Slow sensor acquisition and readout speed, poor branch prediction and cache management strategies, and other real-time operating system effects can slow down numerically complex autonomy software. However, JPL researchers are pursuing alternate methods of increasing computation.

The stereo vision and visual odometry algorithms, which are the basis of many other algorithms but generally dominate their execution time, are amenable to custom hardware implementation—in an FPGA, for example. Not only would this allow the algorithm to run faster, but it offloads the processing from the main CPU. Another method of increasing computation in the harsh environment of space is to use a faster, but less safe, coprocessor for certain computations. When coupled with a fully fault-tolerant main processor, a coprocessor prone to radiation-induced single-event upsets is still useful for noncritical processes, including most relating to autonomy. Alternatively, a multicore architecture might allow many processors to be robust to faults by computing the same algorithm in parallel.

Despite the computational constraints imposed by the harsh environment on Mars, the MER vehicles exploring its surface have made significant use of autonomy, and their success has raised the bar for autonomous operations for the current MSL mission. Further, the experience has proven that autonomy can be used as a mission's critical element. Consequently, increased autonomy will enable future missions, such as the Mars Sample Return mission, that otherwise would have been impossible. ■

Acknowledgments

The work described in this article was carried out at the Jet Propulsion Laboratory, California Institute of



Figure 6. A research rover coring on a slope in the JPL Mars Yard uses visual and force feedback to compensate for slippage and sinkage.

Technology, under a contract to the National Aeronautics and Space Administration, with funding provided by the NASA MER project, Mars Technology Program, Software Intelligent Systems and Modeling program, and New Millennium Program.

References

1. A.H. Mishkin et al., "Experiences with Operations and Autonomy of the Mars Pathfinder Microrover," *Proc. 1998 IEEE Aerospace Conf.*, vol. 2, IEEE Press, 1998, pp. 337-351.
2. M.W. Maimone, P.C. Leger, and J.J. Biesiadecki, "Overview of the Mars Exploration Rovers' Autonomous Mobility and Vision Capabilities," *Proc. 2007 IEEE Int'l Conf. Robotics and Automation Space Robotics Workshop*, IEEE Press, 2007; www-robotics.jpl.nasa.gov/publications/Mark_Maimone/mer_autonomy_icra_2007.pdf.
3. J.J. Biesiadecki and M.W. Maimone, "The Mars Exploration Rover Surface Mobility Flight Software: Driving Ambition," *Proc. 2006 IEEE Aerospace Conf.*, IEEE Press, 2006; www-robotics.jpl.nasa.gov/publications/Mark_Maimone/biesiadecki_maimone06.pdf.
4. J. Carsten et al., "Global Path Planning on Board the Mars Exploration Rovers," *Proc. 2007 IEEE Aerospace Conf.*, IEEE Press, 2007; http://marstech.jpl.nasa.gov/publications/z02_0102.pdf.
5. M. Maimone, Y. Cheng, and L. Matthies, "Two Years of Visual Odometry on the Mars Exploration Rovers," *J. Field Robotics*, Mar. 2007, pp. 169-186.

Call for Articles

IEEE **Software** seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable information to software developers and managers to help them stay on top of rapid technology change. Submissions must be original and no more than 5,400 words, including 200 words for each table and figure.

Author guidelines: www.computer.org/software/author.htm
Further details: software@computer.org
www.computer.org/software

IEEE Software

6. W.S. Kim et al., "Targeted Driving Using Visual Target Tracking on Mars: From Research to Flight," to appear in *J. Field Robotics*, 2008.
7. P. Backes et al., "Automated Rover Positioning and Instrument Placement," *Proc. 2005 IEEE Aerospace Conf.*, IEEE Press, 2005, pp. 60-71.
8. M. Bajracharya et al., "Visual End-Effector Position Error Compensation for Planetary Robotics," *J. Field Robotics*, May 2007, pp. 399-420.
9. A. Castano et al., "Automatic Detection of Dust Devils and Clouds on Mars," *Machine Vision and Applications J.*, Oct. 2008, pp. 467-482.
10. D. Helmick, A. Angelova, and L. Matthies, "Terrain Adaptive Navigation for Planetary Rovers," to appear in *J. Field Robotics*, 2008.
11. M. Bajracharya and P. Backes, "Enabling Reconnaissance and Return for Mars Rovers," to appear in *Proc. 9th Int'l Symp. Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS 08)*, 2008.
12. R. Castano et al., "Oasis: Onboard Autonomous Science Investigation System for Opportunistic Rover Science," *J. Field Robotics*, May 2007, pp. 379-397.
13. P. Backes et al., "Results of Coring from a Low Mass Rover," *Proc. 2008 IEEE Aerospace Conf.*, IEEE Press, Mar. 2008, pp. 1-7.

Max Bajracharya is a senior member of the Mobility and Manipulation Group of the Mobility and Robotics Systems Section at the Jet Propulsion Laboratory, California Institute of Technology. His research interests include vision-based tracking, manipulation, navigation, and pose estimation for Mars rovers. Bajracharya received an MS in computer science and electrical engineering from the Massachusetts Institute of Technology. Contact him at maxb@robotics.jpl.nasa.gov.

Mark W. Maimone is a senior member of the Advanced Robotic Controls Group of the Mobility and Robotics Systems Section at the Jet Propulsion Laboratory. His research interests include robotic navigational autonomy, stereo vision, and camera calibration. Maimone received a PhD in computer science from Carnegie Mellon University. Contact him at mwm@robotics.jpl.nasa.gov.

Daniel Helmick is a senior member of the Mobility and Manipulation Group of the Mobility and Robotics Systems Section at the Jet Propulsion Laboratory. His research interests include sensor-based control of robots, sensor fusion and state estimation, and rover navigation and mobility. Helmick received an MS in mechanical engineering from the Georgia Institute of Technology. Contact him at dhelmick@robotics.jpl.nasa.gov.

IEEE Annals of the History of Computing

IEEE Annals of the History of Computing is an active center for the collection and dissemination of information on historical projects and organizations, oral history activities, and international conferences.

www.computer.org/annals