

Autonomous Navigation Results from the Mars Exploration Rover (MER) Mission

Mark Maimone, Andrew Johnson, Yang Cheng, Reg Willson, and Larry Matthies

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA, 91109, USA
firstname.lastname@jpl.nasa.gov

Abstract. In January, 2004, the Mars Exploration Rover (MER) mission landed two rovers, Spirit and Opportunity, on the surface of Mars. Several autonomous navigation capabilities were employed in space for the first time in this mission. In the Entry, Descent, and Landing (EDL) phase, both landers used a vision system called the Descent Image Motion Estimation System (DIMES) to estimate horizontal velocity during the last 2000 meters (m) of descent, by tracking features on the ground with a downlooking camera, in order to control retro-rocket firing to reduce horizontal velocity before impact. During surface operations, the rovers navigate autonomously using stereo vision for local terrain mapping and a local, reactive planning algorithm called Grid-based Estimation of Surface Traversability Applied to Local Terrain (GESTALT) for obstacle avoidance. In areas of high slip, stereo vision-based visual odometry has been used to estimate rover motion. As of mid-June, Spirit had traversed 3405 m, of which 1253 m were done autonomously; Opportunity had traversed 1264 m, of which 224 m were autonomous. These results have contributed substantially to the success of the mission and paved the way for increased levels of autonomy in future missions.

1 Introduction

Searching for signs of fossil or extant life is a major goal of Mars rover exploration. Given the central role of water in life as we know it, as well as its likely roles in the geology, climate, and resource inventory of the planet, “follow the water” has emerged as a key theme of Mars exploration. The MER mission sent rovers to the Gusev Crater and Meridiani Planum landing sites to attempt to verify that water had played significant roles in the history of these regions. The mission was designed for the rovers to be able to traverse about one kilometer in the course of a 90 day primary mission, operating for about four hours a day (one Mars day is about 24 hours and 40 minutes). They have far exceeded these goals, both in distance traveled and in lifetime. Given the round-trip communication latency between Earth and Mars of 20 minutes or more, this required autonomous navigation capabilities in the rovers. The Sojourner rover in the 1997 Mars Pathfinder mission used structured light for 3-D perception and obstacle avoidance with an extremely limited onboard computer [1]. With greater computational power and the need for more extensive traverse capability, the MER rovers use

stereo vision and a more sophisticated local map and obstacle avoidance algorithm. Visual odometry for position estimation was not in the baseline software design, but was integrated as an “extra credit” item and has grown in importance as surface operations progressed. The DIMES system for horizontal velocity estimation was added late in the mission development phase to address a pressing need to increase the probability of safe landing. This paper gives an overview of the autonomous navigation capabilities of this mission, including the DIMES system used in EDL (section 2), the rover hardware design (section 3), the rover stereo vision and obstacle avoidance system (section 4), and the rover visual odometry system (section 5). Performance of these systems has generally been very good; we show examples drawn from surface operations.

2 Descent Image Motion Estimation System (DIMES)

As shown in Fig. 1, after dropping the heat shield, the MER descent system consisted of a parachute, the backshell from the interplanetary cruise stage, which also held retro-rockets for landing, and the lander itself, which was spooled out on a “bridle” to get it away from the rocket plumes. The retro-rockets (“RADs”, for rocket assisted descent) were designed to bring the vertical velocity to zero 10 to 20 m above the surface, after which the bridle was cut and the lander’s fall was cushioned with airbags. The backshell also incorporated a Transverse Impulse Rocket System (TIRS), which could fire laterally to ensure that the backshell was vertical before RAD firing, so that the RADs did not add to horizontal velocity.

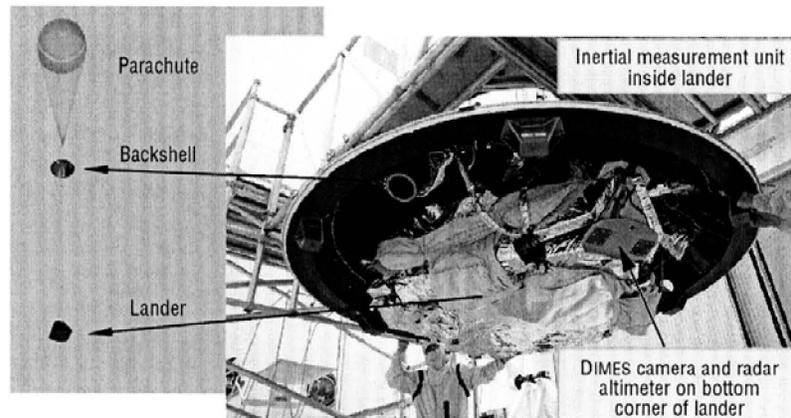


Fig. 1. MER descent system. Left: artist’s illustration of the parachute, backshell, and lander in the last few km of descent. Right: photo of the underside of the lander, before integrating the heat shield.

During mission development, the EDL design team learned that steady-state winds near the martian surface could induce a lander horizontal velocity that

might result in impact forces exceeding the strength of the airbags. The TIRS system could be adapted to tilt the backshell so that the RADs could offset the wind shear, if a horizontal velocity measurement was available. At the time this was learned, it was too late to incorporate a traditional doppler radar velocity sensor; however, it was possible to add a downlooking camera and software to use the camera to track terrain features to estimate the horizontal velocity. We briefly summarize the algorithm design of that system here; more detail appears in [2].

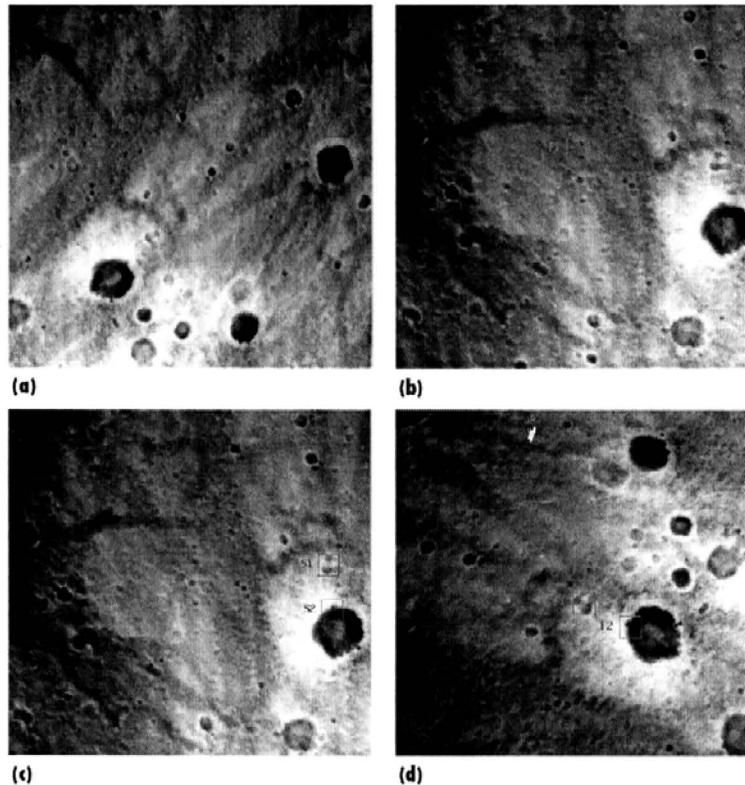


Fig. 2. DIMES results from the Spirit landing at Gusev Crater: (a) first descent image with selected features shown as the red squares; (b) second descent image with the matched features shown as the green squares; (c) second image again, showing the second set of selected features; (d) third image showing the matched features from the second image. All feature matches were correct.

The lander had an inertial measurement unit (IMU) and a radar altimeter that could measure angular velocity and vertical velocity. Thus, in principle, tracking one feature through two images would add enough information to estimate the entire velocity vector. Using just one feature is not very reliable, of course, but the onboard computer was too slow to do much more than this in real-time during de-

scents. The following scheme was adopted to maximize reliability within the available computing resources. An interest operator picked two features in the first image, acquired at about 2000 m above the surface. A multi-resolution correlation algorithm matched those features in a second image, acquired about 1700 m above the surface. A variety of consistency checks were used to validate the matches. This gave one velocity estimate. Two more features were chosen from this image and matched in a third image, acquired at about 1400 m altitude, to give a second velocity estimate. These two velocity estimates gave an estimate for acceleration for the intervening interval, which was checked against accelerations measured with the IMU as a final consistency check. If all consistency checks passed, the velocity estimate was used to determine whether and how to use the TIRS to offset horizontal velocity.

DIMES determined that TIRS firing was needed for Spirit, but not needed for Opportunity. After-the-fact reconstructions of the landing events showed that without DIMES, the Spirit impact velocity would have been right on the edge of the tested airbag performance limits. Fig. 2 shows the descent images from Spirit, with the selected and tracked features for each pair of images.

3 Rover Hardware Overview

Before discussing the rover navigation capabilities, we give a brief overview of the rover hardware. More details are available in press release material on the web [3]. Each vehicle weighs about 174 kg, has a wheelbase of 1.1 m, and is 1.5 m tall to the top of the camera mast. Locomotion is achieved with a rocker bogie system very similar to that used in the 1997 Mars Pathfinder mission, with six driven wheels that are all kept in contact with the ground by passive pivot joints in the rocker bogie suspension. The outer four wheels are steerable.

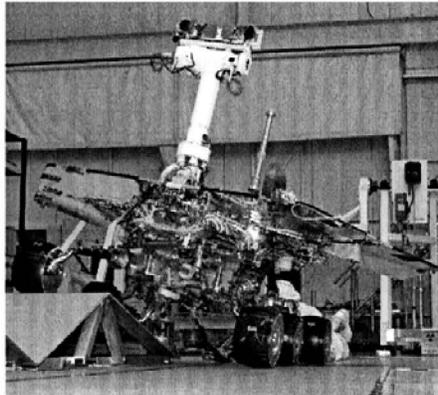


Fig. 3. One of the MER rovers in testing.

The rovers are solar powered, with a rechargeable lithium ion battery for nighttime science and communication operations. The onboard computer is a 20 MHz RAD6000, which has an early PowerPC instruction set, with 128 MB of RAM and 256 MB flash memory. Navigation is done with three sets of stereo camera pairs: one pair of "hazcams" (hazard cameras) looking forward under the solar panel in front, another pair of hazcams looking backward under the solar panel in the back, and a pair of "navcams" (navigation cameras) on the mast. All cameras have

1024x1024 pixel CCD arrays that create 12 bit greyscale images. The hazcams have a 126 degree field of view (FOV) and baseline of 10 cm; the navcams have a 45 degree FOV and baseline of 20 cm [4]. Each rover has a five degree of freedom arm in front which carries a science instrument payload with a microscopic imager, Mossbauer spectrometer, alpha/proton/x-ray backscatter spectrometer (APXS), and a rock abrasion tool (RAT). The camera mast has two additional science instruments: a stereo pair of "pancams" (panoramic cameras) and the "mini-TES" (thermal emission spectrometer). The pancams have filter wheels for multispectral visible and near-infrared imaging for mineral classification. They have the highest angular and range resolution of all cameras on the rover, with a 16 degree field of view and 30 cm baseline. The mini-TES acquires 167 bands between 5 and 29 μm in a single pixel. All cameras on the mast and the mini-TES are pointable by pan/tilt motors.

4 Obstacle Detection and Avoidance

The long round-trip communication latency between Earth and Mars and scheduling constraints on the Deep Space Network make it difficult to control long distance rover traverses from Earth, which necessitates some degree of onboard autonomous navigation to improve operational efficiency, reduce operations cost, and increase mission safety. However, the limited onboard computing power, the need for safety, and other factors constrain the level of sophistication in the algorithms that can be put onboard. Owing to budget and schedule constraints, the baseline autonomous navigation system includes only local obstacle avoidance with stereo vision; that is, there are no onboard global mapping, global path planning, or global localization functions. Stereo vision is used as the range sensor for obstacle avoidance because mature algorithms and reasonably compact, low-power, flight-qualified cameras are available for this, whereas flight-qualified versions of alternate sensors (e.g. lidar) with acceptable performance and form factor were not available.

4.1 Stereo Vision

Details of the stereo algorithm are described in [5]; here we give a brief overview of the algorithm, discuss some details that are specific to the MER implementation, and discuss its performance on Mars.

The stereo algorithm is a typical area-based algorithm using the sum of absolute differences (SAD) criterion for matching. Due to very slow readout from the flight cameras (about 5 seconds per frame for full resolution, 1024x1024 pixel imagery), images are generally binned vertically within the CCD cameras and read out at 256x1024 pixel resolution. This is reduced by averaging to 256x256. Stereo matching is then performed at 256x256 resolution. The images are rectified and bandpass filtered before the SAD operation, a variety of stereo matching consis-

tency checks are applied after the SAD operation, subpixel disparities are computed by quadratic interpolation of SAD scores, and XYZ ranges images are produced as the final result. On the MER flight processor, this can take 24 to 30 seconds per image pair to compute.

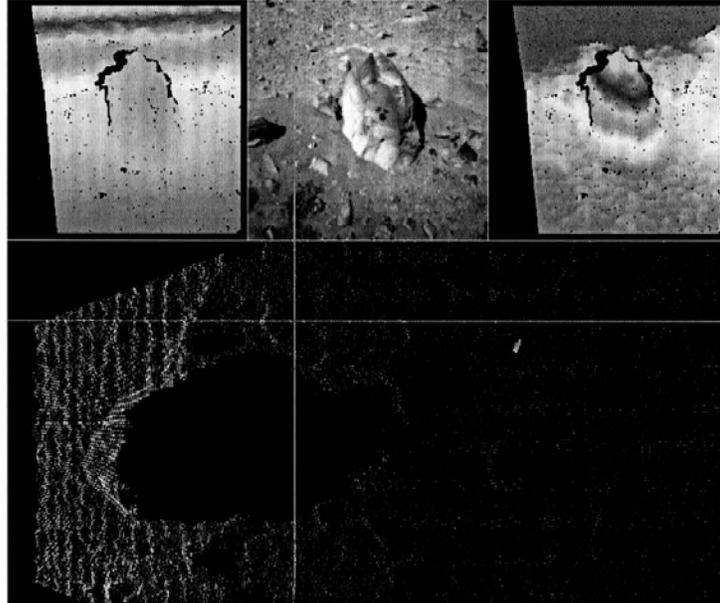


Fig. 4. Stereo results from the Spirit navcams, looking at Humphrey rock in Gusev Crater. The rock is about 0.5 m tall. Upper left: false color range image (red is closest, magenta is furthest). Upper right: false color height image (red is lowest, magenta is highest). Bottom: elevation plot, seen from above, where the cameras are at the left looking right (same color coding as the height image). Green cross-hairs are not significant in this image.

Either the hazcams or the navcams can be used for autonomous navigation. The wide field of view (FOV) of the hazcams was designed to see more than the full width of the rover a short distance ahead of the rover, which is important for obstacle avoidance and to verify the safety of turn-in-place operations. However, the useful lookahead distance with the hazcams is at most 3-4 meters, due to their wide FOV and narrow baseline. The navcams can see further with their narrower FOV and wider baseline, but the FOV is only wide enough to verify the traversability of one candidate path several meters ahead of the vehicle. The pan-cams are not used for autonomous navigation.

Spirit uses the hazcams for stereo. Opportunity was unable to get acceptable range data with the hazcams, because the finer texture in the rock-free soil at Meridiani produced inadequate texture in hazcam images for stereo matching. Adequate results could be obtained with the navcams. Since Meridiani is largely ob-

stacle-free, it has been sufficient to use navcam stereo to check the traversability of the nominal path forward and stop the vehicle if a hazard is detected.

Fig. 4 shows sample stereo results from the Spirit navcams looking at a rock that was studied by the science team.

4.2 GESTALT

The GESTALT obstacle avoidance algorithm is also described in detail in [5]; here we give a brief overview and discuss implementation and performance issues specific to MER. A higher level description of the overall rover driving software architecture appears in [6].

Range images from stereo are converted to “goodness” or “traversability” maps with 20 cm cells in a 10x10 m grid centered on the rover. For each range image, the complete set of range points is analyzed for traversability by fitting planar patches centered on each map cell in turn, where each patch is a circle with the diameter of the rover (nominally 2.6 m). The surface normal, RMS residual, and minimum and maximum elevation difference from the best fit plane determine a “goodness” factor for that map cell that characterizes its traversability. Goodness maps from each range image are registered and accumulated over time with the usual modulo map indexing arithmetic to avoid the need to scroll map data to keep the map bounded. Where new data overlaps old data, the new data overwrites the old data in the map. The merged goodness map is then used to evaluate traversability of a fixed set of candidate steering trajectories, which are circular arcs of varying radius. 23 forward arcs, 23 backward arcs, and two point turns are evaluated in each driving cycle. Evaluation amounts to adding up the goodness scores along each arc, with nearby cells given higher weight. The result is a set of traversability votes for all arcs. These votes are input to an arbiter, which also takes input from waypoints provided by human operators during mission planning. The rover drives a fixed distance along the winning arc before stopping to acquire new images for the next driving cycle. The distance per cycle is set by human operators at anywhere from 35 cm to 1 m or more depending a variety of operational factors, including terrain difficulty and overall distance goals for the day.

Typical computing time per cycle of GESTALT is around 70 seconds. While the rover is driving, its peak speed is 5 cm/sec, but it is typically operated at less than that (3.75 cm/sec) for power reasons. With computing time, the median net driving speed was about 0.6 cm/sec. Because this is so slow and the science team desired to cover large distances to Endurance Crater and the Columbia Hills, a hybrid daily driving scheme was designed in which human operators use navcam and pancam stereo imagery to plan each day’s traverse manually as far as they can see it to be safe. The rover drives this segment blindly, then switches to autonomous navigation to drive for whatever time remains in the day. Blind drives can be several 10’s of meters. The maximum daily traverse for Spirit up to June 14, 2004, was 124 m, of which 62 m were autonomous.

After some tuning of the algorithm to overcome excessively conservative behavior at slope changes, GESTALT has performed well. Fig. 5 shows some results

of obstacle avoidance from Spirit. The terrain on the plains at Meridiani was benign enough that the only “obstacles” were occasional hollows from small, in-filled craters. Within Eagle and Endurance Craters at Meridiani, the main navigation issues were slopes and slippage, not obstacles *per se*.

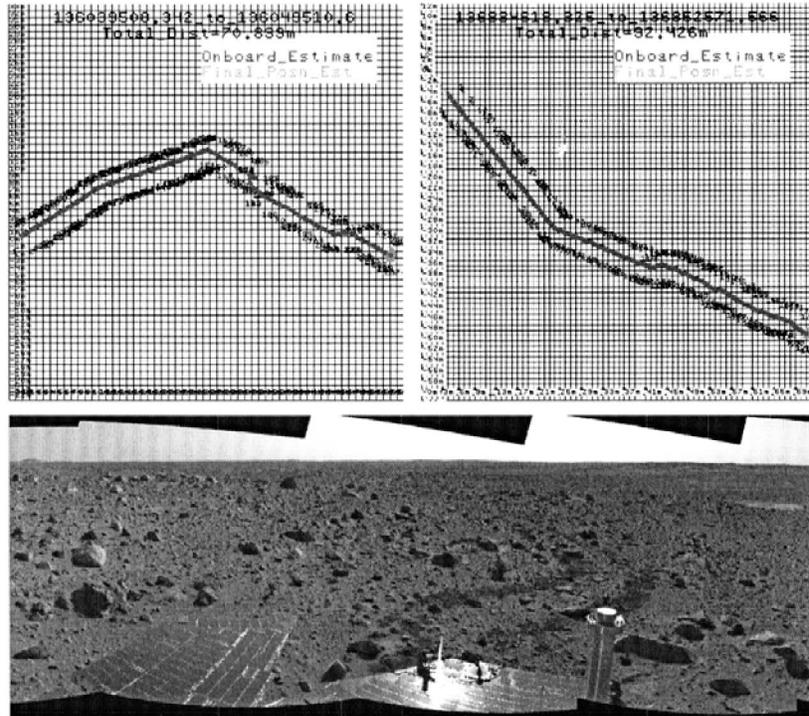


Fig. 5. GESTALT results from Spirit. Top: two examples of daily traverse plots that started with blind drives and ended with autonomous portions that included some rock avoidance maneuvers, covering total distances of 70.9 m (left) and 92.4 m (right). Bottom: a mosaic of navcam imagery looking back on a day’s traverse that ended with some rock avoidance, as can be seen from the rover tracks.

5 Visual Odometry

In routine operation, onboard position estimation is done by dead reckoning with the wheel encoders and IMU, with occasional heading updates by sun sensing with the pancams. Long distance localization is done on Earth using bundle adjustment from manually matched tie points in panoramic imagery [7]. On the plains at Gusev Crater, interpreting the bundle adjustment results as nominal ground truth, dead reckoning errors have been only a few percent of distance over

several kilometers of travel. However, both rovers have experienced large slippage on slopes in the Columbia Hills, Eagle Crater, and Endurance Crater; in fact, up to 125% in one case in the Columbia Hills (ie. the rover slipped backwards in an attempted forward drive). Slip is becoming a bigger issue for Spirit as it attempts to drive with one wheel locked, due to likely imminent failure of that wheel's drive motor. Stereo vision-based visual odometry is part of the flight software, but was not routinely used in most of the mission to maximize driving speed. It has been used to assess slippage of Opportunity in Eagle Crater and is now being used by Spirit to measure and counter the effects of slip on slopes and of dragging the locked wheel.

Our visual odometry algorithm is described in detail elsewhere [8]. In a nutshell, it selects point features, uses multi-resolution area correlation to match them in stereo, tracks them in subsequent stereo pairs, and uses the tracking results to estimate the six degree of freedom rover motion between consecutive stereo pairs. This runs in about 160 sec. We have evaluated its performance against accurate ground truth on Earth-based rover testbeds and found that it can achieve 2% or better of distance over 30 m of travel. Fig. 6 shows sample results from Opportunity in Eagle Crater, where visual odometry correctly detected a 50% slip.

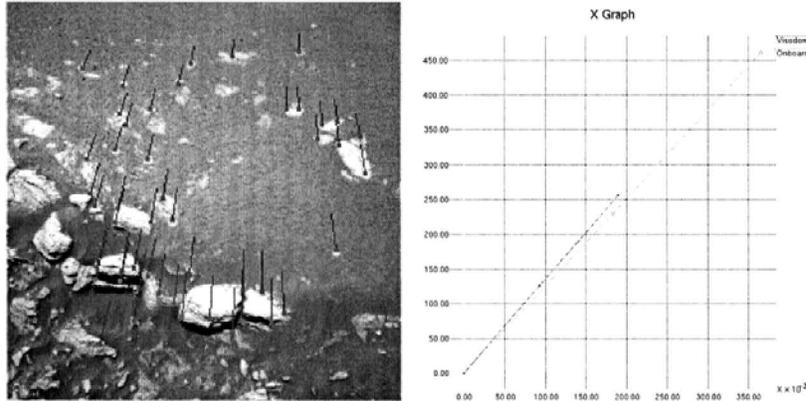


Fig. 6. Visual odometry results from Opportunity in Eagle Crater. Left: features selected (red dots) in an image near a rock outcrop; blue lines show optical flow to the feature position in the next image. Right: plots of onboard dead reckoned position versus visual odometry (in mm) for a case where visual odometry correctly inferred a 50% slip.

6 Discussion

The MER mission was the first use of stereo vision, local map-based obstacle avoidance, and visual odometry for autonomous rover navigation in a planetary exploration mission, as well as the first use of computer vision in an autonomous,

real-time function (horizontal velocity estimation) during landing in a planetary mission. The algorithms are competent and emphasize reliability within the constraints a very slow onboard computer. Algorithms for analogous functions on Earth-based research vehicles can have more sophistication because of the much greater computing resources often available, but are rarely, if ever, designed or tested to reach the same level of fault tolerance. These algorithms have performed well and contributed to the success of the mission; in particular, DIMES may have saved Spirit from a disastrous landing.

In this mission, the most valuable science results have been found in rock outcrops in sloping, slippery terrain, either inside craters or on hills; moreover, the rovers have had to drive much further than anticipated prior to landing to reach such outcrops. Thus, for future missions, key issues include increasing the speed of the vision and planning algorithms and integrating visual odometry inside the driving and steering loop to enable safe, efficient traversal on slippery terrain. Incorporating path planning algorithms for longer lookahead and rougher terrain will also be valuable. For landers, the focus of vision algorithm research will switch to enabling precision landing.

References

- [1] Wilcox, B., and Nguyen, T. Sojourner on Mars and Lessons Learned for Future Planetary Rovers. 28th Int'l Conf. on Environmental Systems, Danvers, MA, July 1998.
- [2] Cheng, Y., Goguen, J., Johnson, A., Leger, C., Matthies, L., San Martin, M., and Willson, R. The Mars Exploration Rovers Descent Image Motion Estimation System. IEEE Intelligent Systems Magazine, May/June 2004.
- [3] <http://marsrovers.jpl.nasa.gov/newsroom/presskits.html>.
- [4] Maki, J., et al. Mars Exploration Rover Engineering Cameras. Journal of Geophysical Research, vol. 108 E12, December 2003.
- [5] Goldberg, S., Maimone, M., and Matthies, L. Stereo Vision and Rover Navigation Software for Planetary Exploration. IEEE Aerospace Conference, Big Sky, Montana, March 2002.
- [6] Maimone, M. and Biesiadicki, J. The Mars Exploration Rover Surface Mobility Flight Software: Driving Ambition. IEEE Aerospace Conference, Big Sky, Montana, March 2005.
- [7] Li, R., Di, K., Matthies, L., Arvidson, R., Folkner, W., Archinal, B. Rover Localization and Landing Site Mapping Technology for the 2003 Mars Exploration Rover Mission. Journal of Photogrammetric Engineering and Remote Sensing, January 2004.
- [8] Olson, C., Matthies, L., and Schoppers, M. Rover Navigation using Stereo Ego-motion. Robotics and Autonomous Systems, 43(4), June 2003.