Visual Odometry on the Mars Exploration Rovers

A Tool to Ensure Accurate Driving and Science Imaging

BY YANG CHENG, MARK W. MAIMONE, AND LARRY MATTHIES

Keeping track of a vehicle's location is one of the most challenging aspects of planetary rover operations. The Mars Exploration Rover (MER) vehicles are typically commanded only once per Martian solar day (sol) using a prescheduled sequence of precise metrically specified commands (e.g., "drive forward 2.34 m, turn in place 0.3567 radians to the right, drive to location X,Y, take color pictures of the terrain at location X,Y,Z" [2]), so having an accurate position estimate on board during the execution of *all* terrain-based commands is of critical importance.
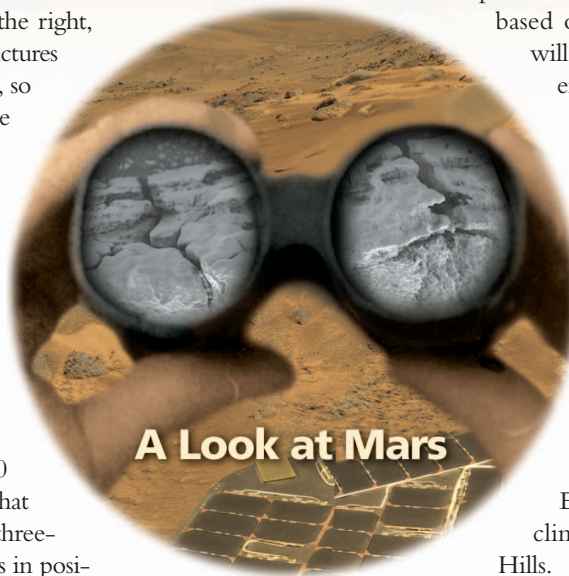
MER rovers onboard position and attitude estimates were updated at 8 Hz nearly every time the wheels or rover arm (instrument deployment device, or IDD) was actuated. Changes in attitude (roll, pitch, yaw) were measured using a Litton LN-200 inertial measurement unit (IMU) that has three-axis accelerometers and three-axis angular rate sensors, and changes in position were estimated based on encoder readings of how much the wheels turned (wheel odometry) [1].

After moving a small amount on a slippery surface, the rovers were often commanded to use camera-based visual odometry to correct any errors in the initial wheel odometry-based estimate that occur when the wheels lose traction on large rocks and steep slopes. Our visual odometry system computes an update to the six-degree-of-freedom (6-DOF) rover pose (x,y,z, roll, pitch, yaw) by tracking the motion of "interesting" terrain features between two pairs of stereo images in both two-dimensional (2-D) pixel coordinates and three-

dimensional (3-D) world coordinates. A maximum likelihood estimator applied to the computed 3-D offsets produces the final motion estimate. However, if any internal consistency check fails, too few feature points are tracked, or the estimation fails to converge, then no motion estimate update will be produced, and the initial estimate (nominally based on wheel odometry and the IMU) will be maintained. The rover will then either continue driving or halt if some preset number of maximum allowed nonconvergences occur [4].

NASA's twin MER rovers *Spirit* and *Opportunity* landed on the surface of Mars in January 2004. As shown in the blue lines of the traverse plots in Figures 1 and 2, human rover drivers have commanded extensive use of the visual odometry software during high-tilt operations: driving *Opportunity* inside Eagle and Endurance craters and climbing *Spirit* through the Columbia Hills.

In the first year since landing, the rovers have driven over terrain with as much as 31° of tilt and over textures composed of slippery sandy material, hard-packed rocky material, and mixtures of both. Engineering models of vehicle slip in sandy terrain developed during Earth-based testing correlated remarkably well with certain sandy Meridiani terrains. However, slip was extremely difficult to predict when the rover was driven over nonhomogeneous terrains (e.g., climbing over rock for one part of a drive and loose soil for another). Early on, the uncertainty in the amount of slip resulting from drives on high slopes or loose soils forced the

A Look at Mars

operations team to spend several days driving toward some targets, even those just a few meters away. But through the rest of the first year of operations, visual odometry software has enabled precision drives (e.g., ending with the science target being directly reachable by the IDD) over distances as long as 8 m, on slopes greater than 20°.

Algorithm

Our approach to position estimation is to find features in a stereo image pair and track them from one frame to the next. This approach, known as visual odometry or ego-motion estimation, was originally developed by Matthies [11]. Following his work, some minor variations and modifications helped improve its robustness and accuracy [14]. Related work using stereo images for localization can be found in [8], [13], and using a single omnidirectional camera in [5]. The key idea of the present method is to determine the change in position and attitude for two or more pairs of stereo images using maximum likelihood estimation. The basic steps of this method are described as follows.
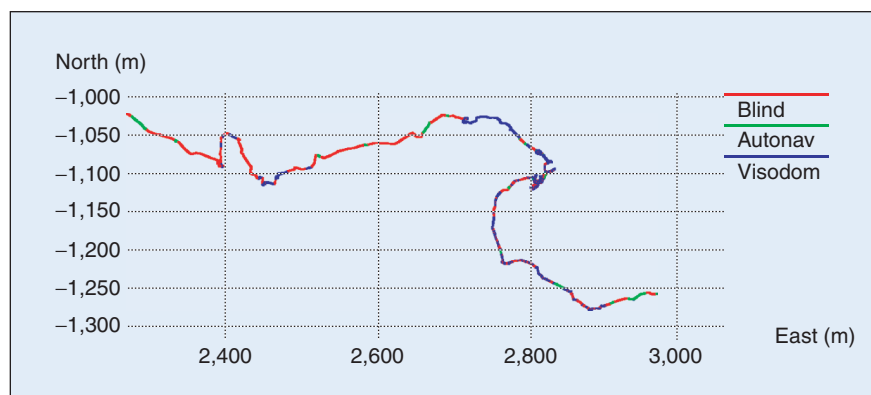
Feature Detection

First, features that can be easily matched between stereo pairs and tracked across a single motion step are selected. An interest operator tuned for corner detection (e.g., Forstner or Harris) is applied to an image pair, and pixels with the highest interest values are selected. To reduce the computational cost, a grid with cells smaller than the minimum dis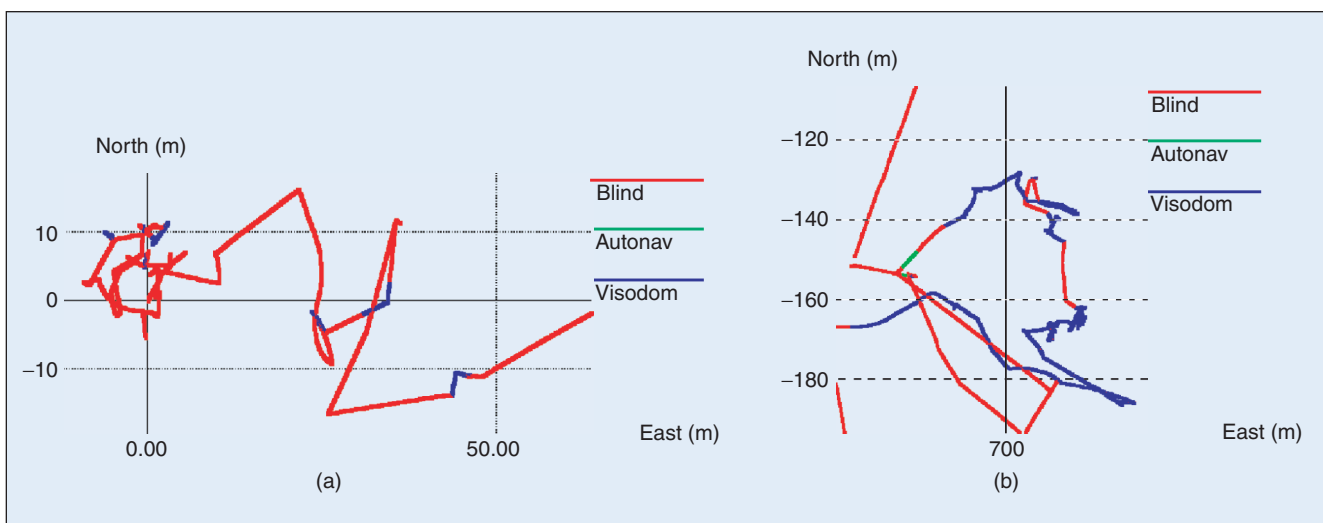tance between features is superimposed on the left image. The feature with strongest corner response in each grid cell is selected as a viable candidate. A fixed number of features having the highest interest operator responses is selected, subject to a minimum distance constraint to ensure that features span the image.

Feature-Based Stereo Matching

Each selected feature's 3-D position is computed by stereo matching. Because the stereo cameras are well calibrated, the stereo matching is done strictly along the epipolar line with only a few pixels of offset buffer above and below it. We use pseudo-normalized correlation to determine the best match. In order to obtain subpixel accuracy, a biquadratic polynomial is fit to a 3×3 neighborhood of correlation scores, and the peak of this polynomial is chosen as the correlation peak.

Figure 1. Plot of Spirit's traverse history using visual odometry in the Columbia Hills from Sols 178–585. Units are in meters from the landing site origin, as measured on board the rovers. Red lines indicate directly commanded "blind" drives, green lines indicate autonomous hazard detection, and blue lines indicate visual odometry. Spirit only used visual odometry within the Columbia Hills, not during its 3-km trek to reach them.

Figure 2. Plot of Opportunity's traverse history using visual odometry. (a) The drive in and around 20-m diameter Eagle Crater from Sols 1–70. (b) The drive in and around Endurance Crater from Sols 133–312. Units are in meters from the landing site origin, as measured on board the rovers. Red lines indicate directly commanded "blind" drives, green lines indicate autonomous hazard detection, and blue lines indicate visual odometry.

The 3-D positions of these selected features are determined by intersecting rays projected through the camera models. Under perfect conditions, the rays of the same feature in the left and right images should intersect at a point in space. However, due to image noise, camera model uncertainty, and matching error, they do not always intersect. The shortest distance "gap" between the two rays indicates the goodness of the stereo match: features with large gaps are thrown out.

The quality of a 3-D feature is a function of its relative location, the gap between the two stereo rays, and the sharpness of the correlation peak. The covariance computation fully reflects these three factors, and we compute the covariance associated with each feature. Assume the stereo cameras are located at $C_1(X_1, Y_1, Z_1)$ and $C_2(X_2, Y_2, Z_2)$. Let r_1 and r_2 be two unit vectors connecting C_1 and C_2 to the same feature in both images. Because of noise, r_1 and r_2 do not always intersect precisely at a point. But as long as the rays converge in front of the lens and are not exactly parallel, there will always be a shortest line segment connecting these two rays. Letting P_1 and P_2 be the endpoints of that line segment, we have the following constraints on the distance from each camera to that feature (where $m_1 = |P_1 C_1|$, $m_2 = |P_2 C_2|$):

$$P_1 = C_1 + r_1 m_1 \tag{1}$$

$$P_2 = C_2 + r_2 m_2. \tag{2}$$

Therefore we have

$$(P_2 - P_1) r_1 = (C_2 - C_1 + r_2 m_2 - r_1 m_1) \cdot r_1$$
$$= 0 \tag{3}$$

$$(P_2 - P_1) r_2 = (C_2 - C_1 + r_2 m_2 - r_1 m_1) \cdot r_2$$
$$= 0. \tag{4}$$

Then we have

$$m_1 = \frac{B r_1 - (B \cdot r_2)(r_1 \cdot r_2)}{1 - (r_1 \cdot r_2)^2} \tag{5}$$

$$m_2 = (r_1 \cdot r_2) m_1 - B r_2 \tag{6}$$

$$P = (P_1 + P_2)/2, \tag{7}$$

where $B = C_2 - C_1$ is the stereo baseline, and m_1 and m_2 are functions of feature locations in both images whose partial derivatives are:

$$m_1' = \{[B \cdot r_1' - (B \cdot r_2')(r_1 \cdot r_2)$$
$$- (B \cdot r_2)(r_1' \cdot r_2 + r_1 \cdot r_2')][1 - (r_1 \cdot r_2)^2]$$
$$+ 2[B \cdot r_1 - (B \cdot r_2)(r_1 \cdot r_2)]$$
$$\times [(r_1 \cdot r_2)(r_1' \cdot r_2 + r_1 \cdot r_2')]\}/[1 - (r_1 \cdot r_2)^2]^2 \tag{8}$$

$$m_2' = (r_1 \cdot r_2) m_1' + (r_1' \cdot r_2 + r_1 \cdot r_2') m_1 - B r_2' \tag{9}$$

$$P' = (r_1' m_1 + r_1 m_1' + r_2' m_2 + r_2 m_2')/2. \tag{10}$$

Further note that the covariance of P is

$$\sum_P = P' \begin{bmatrix} \sum_l & 0 \\ 0 & \sum_r \end{bmatrix} P'^T, \tag{11}$$

where P' is the Jacobian matrix or the first partial derivative of P with respect to the 2-D feature locations in the left and right images, and $\sum_l$ and $\sum_r$ are 2×2 matrices whose elements are the curvatures of the biquadratic polynomial along the vertical, horizontal, and diagonal directions, which can be obtained directly from subpixel interpolation.

Feature Tracking

After the rover moves a short distance, a second pair of stereo images is acquired. The features selected from the previous image are projected into the second pair using the approximate motion provided by onboard wheel odometry (see Figure 3). Then a correlation-based search reestablishes the 2-D positions precisely in the second image pair. Stereo matching of these tracked features determines their new 3-D positions. Because the 3-D positions of those tracked features are already known from the previous step, the stereo matching search range can be greatly reduced. Features whose initial and final 3-D positions differ by too large an amount are filtered out.

Robust Motion Estimation

If the initial motion is accurate, the difference between two estimated 3-D positions should be within the error ellipse. However, when the initial motion is off, the difference between the two estimated positions reflects the error of the initial motion, and it can be used to determine the change of rover position.

The motion estimation is done in two steps. First, a less accurate motion is estimated by least squares estimation. The error residual is

$$e_j = P_{cj} - R P_{pj} - T, \tag{12}$$

and the cost expression is

$$M(R, T) = \sum w_j e_j^T e_j \tag{13}$$

$$w_j = \left(det\left(\sum_{pj}\right) + det\left(\sum_{cj}\right) \right)^{-1}. \tag{14}$$

There is a closed form solution for this least squares estimation [12]. The advantage of this least squares method is that it is simple, fast, and robust. Its disadvantage is that it is less accurate because it only takes the quality (the volume of the error ellipsoid) of the observations as a weight factor.

Because it is an inexpensive operation, we embed it within a RANSAC (random sample consensus) process to do outlier removal:

1) A small set of features (e.g., 6) is randomly selected and the motion is then estimated using the least squares estimation method.
2) All features from previous steps are projected to the current image frame by the newly estimated motion. If the gap between a reprojected feature and its correspondent is less than a threshold (e.g., 0.5), the score of this iteration will be incremented once for each viable feature.
3) Steps 1 and 2 repeat for a fixed number of iterations, and the motion with the highest score is selected. All features that pass this iteration will be used in the following more accurate estimation—the maximum likelihood motion estimation.

The maximum likelihood motion estimation considers the 3-D position difference and associated error models when estimating position. Let P_{pj} and P_{cj} be the observed feature positions before and after the current robot motion. Then

$$P_{cj} = RP_{pj} + T + e_j, \qquad (15)$$

where R and T are the rotation and translation of the robot, and e_j is the combined error in the observed positions of jth features. In this estimation, three axis rotations θ_R and translation T are directly determined by minimizing the summation in the exponents

$$\sum r_j^T W_j r_j \qquad (16)$$
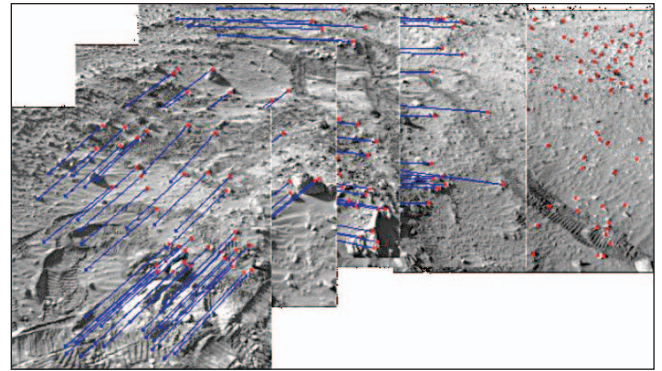
$$r_j = P_{cj} - RP_{pj} - T, \qquad (17)$$

where W_j is the inverse covariance matrix of e_j. The minimization of this nonlinear problem is done by linearization and an iterative process [12]. Two nice properties of the maximum-likelihood estimation make the algorithm powerful. First, it estimates the three axis rotations θ_R directly so that it eliminates the error caused by rotation matrix estimation done by the least squares estimation. Secondly, it fully incorporates error models (the shape of the ellipsoid) into the estimation, which greatly improves the accuracy of the final motion estimate.

As of the February 2005 version of MER flight software, optional constraints can also be placed on the final motion estimate to provide additional sanity checking. The magnitude of the 3-D update vector, its X and Y site frame components, the magnitude of the roll, pitch and yaw updates, and the angular deviation from a purely downslope vector can all be restricted. Any update violating the active set of constraints is treated as a failure to converge; the number of acceptable failures is another optional constraint.
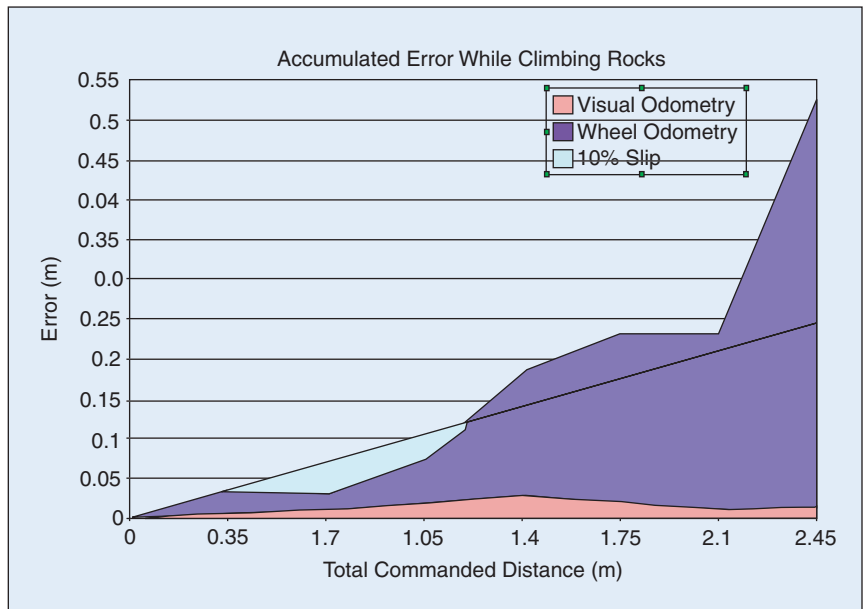
Ground-Based Validation

This visual odometry software has been tested on numerous rover platforms. The latest tests were conducted on the Jet Propulsion Laboratory's (JPL) Rocky 8 rover at the JPL Marsyard and in Johnson Valley, California [7]. Rocky 8 has two pairs of hazard avoidance stereo cameras mounted on the front and rear of the rover body about 50 cm above the ground. The image resolution is 640 × 480, horizontal and vertical fields of view are 80° horizontal by 64° vertical, and the baseline is about 8.4 cm. The Johnson Valley site had slopes of loose granular sand where the rover experienced substantial slip, tilt, and roll during the test.

In order to evaluate visual odometry performance, high-precision ground-truth data (position and attitude) was also collected using a total station (like a surveyor's theodolite with

Figure 3. Feature tracking occurs between every pair of images. In this view, several images from Spirit's Sol 178 drive and their tracked features have been superimposed.

Figure 4. Visual odometry error measured during a 2.45-m drive using HAZCAMs on the MER Surface System Testbed Lite rover. The rover was driven over several large nonobstacle rocks, each less than 20-cm tall, in 35-cm steps. The vehicle was held in place during the final step, so the wheel odometry error for that step is artificially large, yet the visual odometry error remains small.

The improved drive accuracy in new or mixed-soil terrains also yielded a greater number of science observations, by reducing the number of sols needed to make targets reachable by the instrument arm (IDD).

a laser range sensor). By tracking four prisms on top of the rover, the rover's position and attitude were measured with high precision (<2 mm in position and <0.2° in attitude). The absolute position errors were less than 2.5% over the 24-m Marsyard course and less than 1.5% over the 29-m Johnson Valley course. The rotation error was less than 5° in each case.
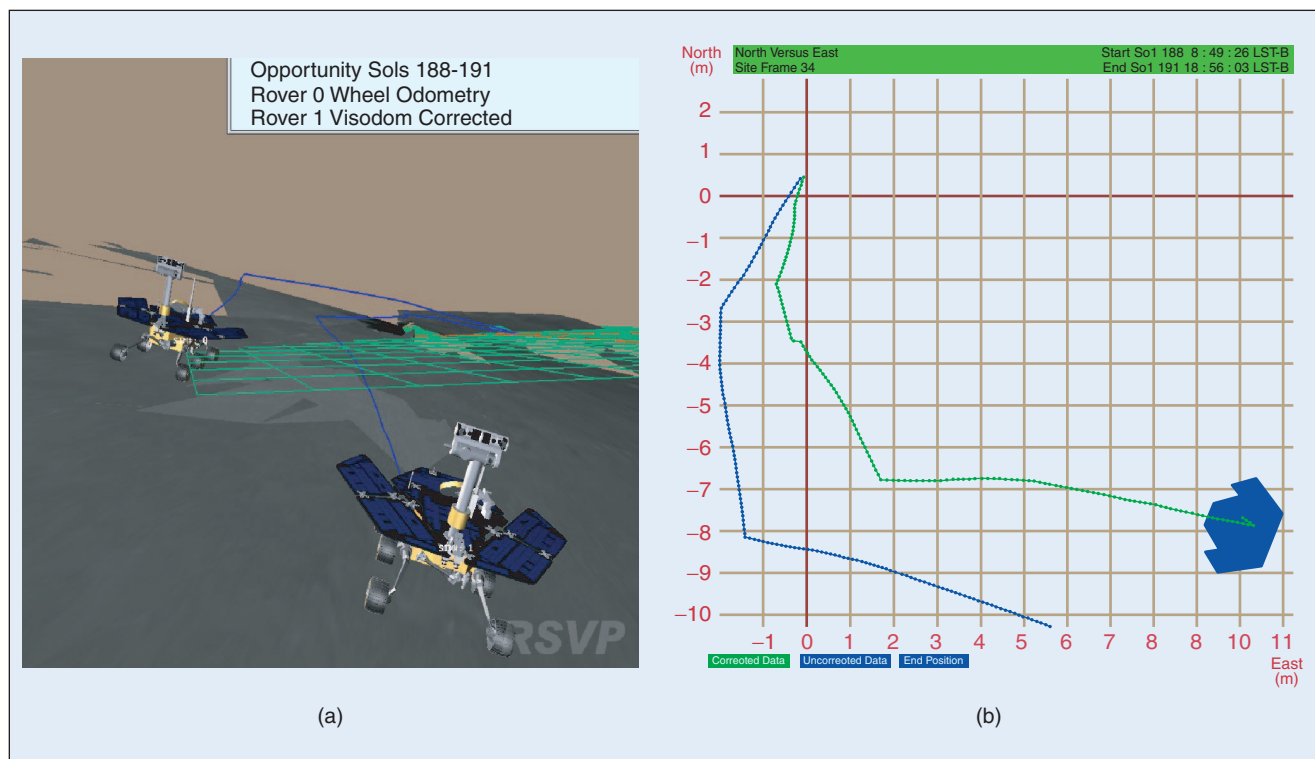
Tests were also run on the MER Surface System Testbed Lite rover in an indoor sandbox test area. Ground truth was acquired using a total station to measure the vehicle's 6-DOF motion estimate by tracking three points at each step. During these tests visual odometry processing took place using images from the 120° field-of-view (FOV) HAZCAM sensors (but on Mars only the 45° FOV NAVCAMs are commanded). Several tests were run in which visual odometry was found to be as good as wheel odometry on simple terrain (within the design specifications) and much better in complex terrain.

Figure 4 shows the position estimation error that resulted from the most slip-inducing test run: a 2.45-m rock-laden course driven in 35-cm steps. The straight line and light-blue background represent the design goal of at most 10% error in the position estimate. The dark curve represents the error that accrued when the position was estimated using only the IMU and wheel odometry; after 1.4 m of driving, the accumulated error had already gone beyond the desired 10% curve. Finally, the light curve at the bottom represents the error remaining after visual odometry processing has completed. Even after 2.45 m of driving over rough obstacles with as much as 85% slip, the visual odometry error remained small, less than 1% of the total traverse distance.

Using Visual Odometry on Mars

Visual odometry processing was performed on both MER rovers using mast-mounted NAVCAM imagery. NAVCAMs have a 45° FOV and sit 1.5 m above the ground plane [10], so all visual odometry drives were split into small steps to ensure at least 60% overlap between adjacent images. During each step, the rover was typically commanded to drive no more than 75 cm in a straight line or curved arc and when turning in place was commanded to change heading by no more than 18° per step. Motions outside these bounds forced the process to reinitialize.

Figure 5. *Views of* Opportunity's *19-m drive from Sol 188 through Sol 191. The inside path—green in (b) the overhead image— shows the correct visual odometry updated location. The outside path—blue in (b) the overhead image—shows how its path* would *have been estimated from the IMU and wheel encoders alone. Each cell represents one square meter.*

Although visual odometry processing could have been beneficial during all rover motion, each step required 2–3 min of computation time on the MER's 20-MHz RAD6000 CPU, and thus it was only commanded during relatively short drives (typically less than 15 m) that occurred either on steep slopes (typically more than $10°$) or in situations where a wheel was being dragged (digging a trench or conserving drive motor lifetime on *Spirit*'s right front wheel). The onboard IMU exhibited a very small drift rate (usually less than $3°$ per hour of operation) and, therefore, maintained attitude knowledge very well. So, during the first year of operations from January 2004 through January 2005, visual odometry was typically used to update rover position only. The tradeoffs that went into the daily decision of whether to use visual odometry are discussed in [3].

There were some instances in which visual odometry did not converge to a solution. These are mostly attributable to either too large a motion (e.g., a $40°$ turn in place, resulting in too little image overlap) or to a lack of features in the imaged terrain; but see the section "False Positives" too. It has successfully measured slips as high as 125% on Sol 206 when *Spirit* tried to climb a more than $25°$ slope.

Several benefits were realized from visual odometry. Vehicle safety was maintained by having the rover terminate a planned drive early if it realized via visual odometry that it was making insufficient progress toward its goal or was nearing the prespecified location of an obstacle. The improved drive accuracy in new or mixed-soil terrains also yielded a greater number of science observations, by reducing the number of sols needed to make targets reachable by the instrument arm (IDD). And PANCAM (panoramic camera) and MiniTES science observations requiring precision pointing of the mast were often scheduled in the middle of a drive, using visual odometry to eliminate the need for human confirmation of the pointing angle.
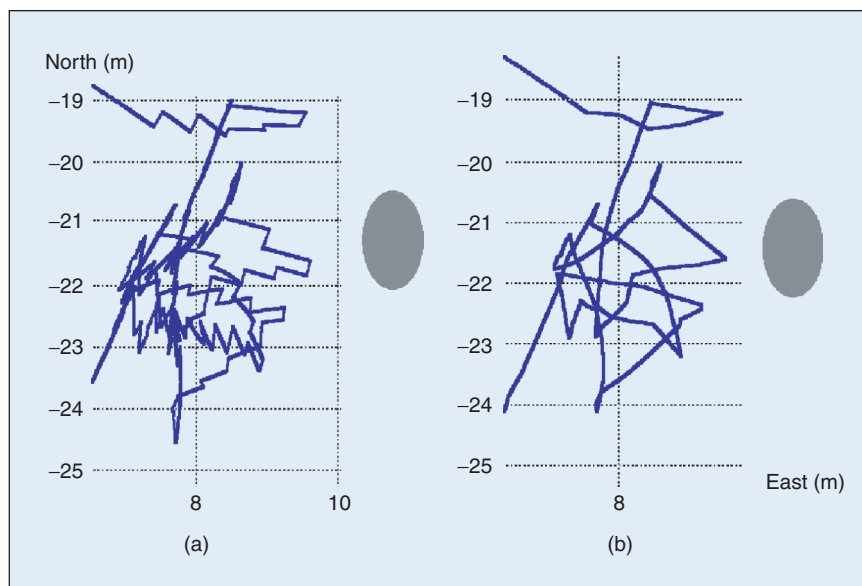
Meridiani Planum: Opportunity Rover

The terrain at Meridiani Planum is a challenging one for visual odometry. It is often difficult or impossible to find a patch of nearby terrain that has enough texture for visual odometry processing to successfully find and track features because much terrain is covered by a thick layer of extremely fine particles. Fortunately, areas that have this smooth, featureless appearance tend to be very flat, and in those areas, the IMU and encoder-based position estimation has performed well enough that visual odometry was not needed. Terrain that exhibits higher slope (and consequently more position uncertainty) almost always

has a distinctive appearance (e.g., bedrock outcrop) or is near enough to interesting features that visual odometry can be employed successfully.

The path predicted by wheel odometry alone can be quite different from the actual path. Figure 5 shows two views of the trajectory taken by *Opportunity* during Sols 188–191. The rover was driven uphill and across slope over a real distance of 19 m, but wheel odometry underestimated it by 1.6 m and failed to measure the slip-induced elevation change. The outside path indicates the course as estimated solely by wheel odometry, and the inside path shows the visual odometry-corrected course plot that was actually generated on board. The final positions differ by nearly 5 m.
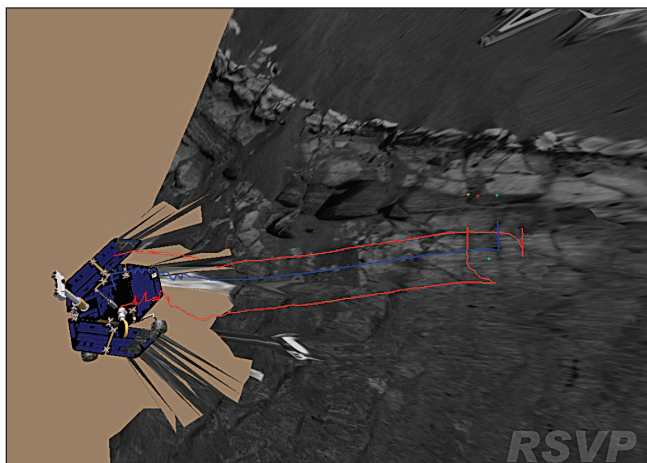

Figure 6. *Wopmay, an obstacle inside Endurance Crater, is 60 cm tall, 90 cm wide, and 150 cm long.*
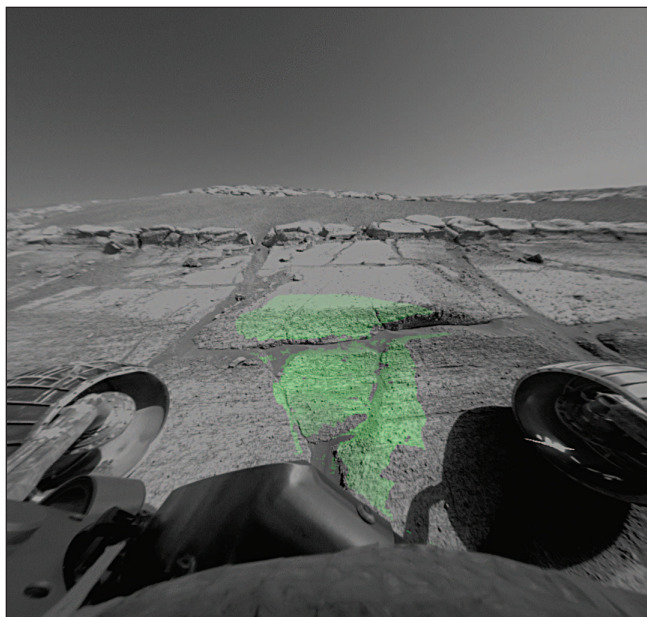

Figure 7. Opportunity*'s 15-sol trajectory near Wopmay (see Figure 6, approximately indicated by the gray ellipse), first driving toward and then trying to get around or away from it. The downslope is up and to the right. In (a), the "jumps" that point up to the right are the result of visual odometry adjusting the vehicle's position down slope. Visual odometry only corrects the rover's position at the* end *of each step of less than 1 m. In (b) is shown the same course with the visual odometry jumps approximately corrected.*

The earliest benefit from visual odometry came inside the 20-m diameter Eagle Crater, *Opportunity*'s landing site. Most driving inside Eagle Crater was meticulously planned by human drivers, predicting slip using tables generated by the mechanical team from Earth-based tests of a rover driving in sand. But while those tables worked well for predicting purely upslope and cross-slope slips on pure sand, no model was available for driving on pure bedrock outcrop, mixtures of bedrock and loose sand, or at angles other than 0°, 45°, and 90° from the gradient. In those circumstances visual odometry was sometimes used to drive to the proper target or ensure that high-resolution PANCAM images of science targets taken after a drive would be pointed right on target [see Figure 2(a)].

Figure 8. Opportunity*'s planned 8.7-m drive along a 20–24° slope on Sol 304.*

Figure 9. *After* Opportunity*'s 8.7-m slope drive on Sol 304, the goal area is perfectly reachable inside the IDD work volume, indicated in green.*

But the most extensive use of visual odometry was made by *Opportunity* inside 130-m diameter Endurance Crater from Sol 133 to Sol 312 (see the right-hand side of Figure 2). Except for a 12-m approach and return at the lowest point (with lowest rover tilt) on Sols 201 and 203 and a 17-m drive on Sol 249, visual odometry was used virtually continuously throughout. Had it not been available on board, many more sols would have been needed to approach targets, and fewer targets would have been achieved. But visual odometry not only improved target approach efficiency, it also proved crucial to maintaining vehicle safety.

From Sols 249–265, *Opportunity* kept finding itself near a 1.5-m long rock called *Wopmay* (see Figures 6 and 7). Although Wopmay was originally considered a science target, it also proved to be a most difficult obstacle to avoid. It was located downhill from a 17–20° downslope area composed of loose sand and buried rocks. Several attempts to drive around it were thwarted not only by very high slip but also by the unseen rocks buried just beneath the surface. Fortunately, the human-commanded sequences took into account the possibility that the rover might slip, so *Opportunity* halted its planned drives prematurely (and *correctly*) when it realized that it was moving too close to Wopmay.

Visual odometry also enabled more precise approaches to difficult targets. On Sol 304, a drive of over 8 m was planned on an outcrop whose slope varied from 20–24°. Because the drive plan took a wide range of potential slips into account, *Opportunity* was able to drive just far enough across slope then turn and drive just far enough upslope to perfectly position the desired target within the IDD work volume in a single sol. Figure 8 illustrates the planned drive, and Figure 9 shows the final image from the body-mounted front Hazard cameras (HAZCAMs), showing the target area perfectly located between the front wheels.

Visual odometry results are summarized in Table 1. As of March 2005, *Opportunity* has thrived for 394 sols. Visual odometry was used more here than on *Spirit* because *Opportunity* spent more of its first year on slippery surfaces. It has converged to a solution 95% (828/875) of the time.

Gusev Crater: Spirit *Rover*

The terrain at Gusev Crater is well suited for visual odometry processing. The rock abundances there matched predicted distributions [6], resulting in a generally feature-rich landscape with detailed textures composed of rocks of different sizes and brightnesses. When planning for drives using visual odometry, rover drivers typically only had to bear in mind the restriction that adjacent frames should have at least 60% image overlap, although they sometimes also had to avoid pointing the cameras at (the relatively infrequently occurring) sand dunes. As a result, *Spirit*'s visual odometry software has performed admirably. An overview of the types of driving modes employed on *Spirit* can be found in [9].

One unique driving mode that benefited a great deal from visual odometry on *Spirit* was wheel dragging. The right front wheel was found to draw more current while driving than

any of the other wheels starting on Sol 125. This concern led to the development of a driving strategy to conserve motor lifetime, during which that wheel would be dragged while all the others were driven. Although this was found to enable reasonable progress on relatively flat terrain, error in the position estimate grew substantially in this mode. The visual odometry capability meant that not only could progress be made, but also the error added to the onboard position estimate could be bounded as well.

Relatively little slip was seen during the first six months of *Spirit*'s mission. But once the base of the Columbia Hills was reached, drives in the hills were found to exhibit much more unpredictable slip. Thus visual odometry has been used during most of the drives in the Columbia Hills, especially to ensure that *Spirit* keeps far enough away from nearby rock obstacles. The average tilt of the rover during those times that visual odometry was commanded was 14.4°, counting 625 samples spanning an absolute range from 2–30°.

Visual odometry results are summarized in Table 1. As of March 2005, *Spirit* has thrived for 414 sols. Visual odometry was only used on *Spirit* after it had reached the Columbia Hills, nearly six months into its mission. But since then, it has converged to a solution 97% (590/609) of the time.

False Positives

Although we had never seen visual odometry converge to an inaccurate solution during testing, on *Opportunity* Sols 137 and 141, several unreasonable position updates were computed on board. These are attributable to an improper parameter setting; at that time, the minimum separation between features was too small. As a result, the set of detected features was allowed to cluster tightly around a small planar but feature-

Science observations requiring precision pointing of the mast often used visual odometry to eliminate the need for human confirmation.

rich area. Increasing that parameter was all that was needed to allow the software to find additional out-of-plane features and converge to a reasonable solution.

As of March 2005, the only other instance of a false positive solution was on Sol 235. During that sol, the NAVCAMs were pointed at two small and widely separated rocks. Although features were found on those rocks, many more features were found on the largest shape in the image, the rover's shadow. So even though forward drive progress was made, the onboard estimator assumed that the shadow (having more distinct features spread throughout the image) better reflected actual motion and, therefore, produced an incorrect estimate. This problem would not have arisen had there been more interesting texture around and under the shadow, and since then, human drivers have had to take the rover shadow into account whenever planning visual odometry drives.

Conclusions

Visual odometry has been a highly effective tool for maintaining vehicle safety while driving near obstacles on slopes, achieving difficult drive approaches in fewer sols, and ensuring accurate science imaging. Although it requires active pointing by human drivers in feature-poor terrain, the improved position knowledge enables more autonomous capability and better science return during planetary operations.

Table 1. Results of running visual odometry on Mars. Expressions *m +/− s* indicate mean *m* and standard deviation *s*.				
	Spirit		**Opportunity**	
Lifetime as of 4 March 2005	414	sols	394	sols
Total drive distance	4161	m	3158	m
Days spent driving	184	sols	172	sols
Days using visual odometry	52	sols	75	sols
Nominal evaluation steps	609	pairs	875	pairs
Nominal initialization steps	57	pairs	75	pairs
Forced initialization by large turn	10	pairs	11	pairs
Forced initialization by planned repointing	5	pairs	9	pairs
Forced initialization by driving too far	1	pairs	2	pairs
Total visodom image pairs processed	682	pairs	972	pairs
Successful (noninitial) convergences	590	pairs	828	pairs
Iterations needed (assume convergence)	6.4 +/− 1.7	iterations	8.4 +/− 5.2	iterations
Features tracked at each step	73.4 +/− 29.3	features	87.4 +/− 34.1	features
Nonconvergences	19	pairs	47	pairs
Mean updates per drive sol	12.0 +/− 8.5	pairs	12.9 +/− 10.2	pairs
Max updates per drive sol	33	pairs	59	pairs
Mean rover tilt during visodom	14.6 +/− 4.4	°	18.0 +/− 4.6	°
Absolute tilt range during visodom	2–30	°	0.8–31	°

Acknowledgments

Keywords

MER, Mars Exploration Rover, visual odometry, motion estimation, egomotion.

References

[1] K.S. Ali, C.A. Vanelli, J.J. Biesiadecki, M.W. Maimone, Y. Cheng, M.S. Martin, and J.W. Alexander, "Attitude and position esimation on the Mars exploration rovers," in *Proc. IEEE Conf. Syst., Man, Cybern.*, The Big Island, HI, Oct. 2005, pp. 20–27.

[2] J.J. Biesiadecki, E.T. Baumgartner, R.G. Bonitz, B.K. Cooper, F.R. Hartman, P.C. Leger, M.W. Maimone, S.A. Maxwell, A. Trebi-Ollenu, E.W. Tunstel, and J.R. Wright, "Mars exploration rover surface operations: Driving opportunity at meridiani planum," in *Proc. IEEE Conf. Syst., Man, Cybern.*, The Big Island, HI, Oct. 2005, pp. 1823–1830.

[3] J.J. Biesiadecki, C. Leger, and M.W. Maimone, "Tradeoffs between directed and autonomous driving on the mars exploration rovers," in *Proc. Int. Symp. Robot. Res.*, San Francisco, CA, Oct. 2005.

[4] J.J. Biesiadecki and M.W. Maimone, "The Mars exploration rover surface mobility flight software: Driving ambition," in *Proc. IEEE Aerospace Conf.*, vol. 5, Big Sky, MT, Mar. 2005.

[5] P.I. Corke, D. Strelow, and S. Singh, "Omnidirectional visual odometry for a planetary rover," in *Proc. IROS 2004*, [Online]. Available: http://www. ri.cmu.edu/pubs/pub_4913.html

[6] M. Golombek and D. Rapp, "Size-frequency distributions of rocks on Mars and Earth analog sites: Implications for future landed missions," *J. Geophys. Res.—Planets*, vol. 102, no. E2, pp. 4117–4129, Feb. 1997.

[7] D.M. Helmick, Y. Cheng, S.I. Roumeliotis, D. Clouse, and L. Matthies, "Path following using visual odometry for a Mars rover in high-slip environments," in *Proc. IEEE Aerospace Conf.*, Big Sky, MT, Mar. 2004.

[8] S. Lacroix, A. Mallet, R. Chatila, and L. Gallo, "Rover self localization in planetary-like environments," in *Proc. Int. Symp. Artif. Intell., Robot., Automat. Space (i-SAIRAS)*, Noordwijk, Netherlands, June 1999, pp. 433–440 [Online]. Available: ftp://ftp.laas.fr/pub/ria/simon/isairas99.ps.gz

[9] C. Leger, A. Trebi-Ollenu, J. Wright, S. Maxwell, B. Bonitz, J. Biesiadecki, F. Hartman, B. Cooper, E. Baumgartner, and M. Maimone, "Mars exploration rover surface operations: Driving spirit at gusev crater," in *Proc. IEEE Conf. Syst., Man, Cybern.*, The Big Island, Hawaii, USA, Oct. 2005, pp. 1815–1822.

[10] J.N. Maki, J.F. Bell III, K.E. Herkenhoff, S.W. Squyres, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, R. Willson, A. Johnson, M. Maimone, E. Baumgartner, A. Collins, M. Wadsworth, S.T. Elliot, A. Dingizian, D. Brown, E.C. Hagerott, L. Scherr, R. Deen, D. Alexander, and J. Lorre, "Mars exploration rover engineering cameras," *J. Geophys. Res.*, vol. 108, no. E12, pp. 12–1–12–24, Dec. 2003 [Online].Available: http://www.agu.org/pubs/crossref/2003/2003JE002077.shtml

[11] L. Matthies and S. Shafer, "Error modelling in stereo navigation," *IEEE J. Robot. Automat.*, vol. RA-3, no. 3, June 1987.

[12] L. Matthies, *Dynamic Stereo Vision*. Ph.D. dissertation, Comput. Sci. Dept., Carnegie Mellon Univ., CMU-CS-89-195, Pittsburgh, PA, Oct. 1989.

[13] D. Nister, O. Naroditsky, and J. Bergen, "Visual odometry," in *Computer Vision and Pattern Recognition*. Los Alamitos: IEEE Comput. Soc. Press, 2004, pp. 652–659.

[14] C.F. Olson, L.H. Matthies, M. Schoppers, and M.W. Maimone, "Rover navigation using stereo ego-motion," *Robot. Auton. Syst.*, vol. 43, no. 4, pp. 215–229, June 2003.

Yang Cheng is a senior member of technical staff at NASA-Jet Propulsion Laboratory (JPL). He earned the Ph.D. in geographic information processing at the University of South Carolina, where he studied remote sensing and parallel computation. In 1993, he moved to the Oak Ridge National Laboratory, Tennessee, where he studed remote sensing and large-scale geographic information processing. In 1999, he came to JPL. His expertise covers areas of robotic vision, cartography, map projection, photogrammetry, remote sensing, and parallel computing. He was a main developer of the Descent Image Motion Estimation System (DIMES) and visual odometry of the NASA Mars Exploration Rover (MER) mission. He currently works on vision algorithm development for small body navigation, spacecraft safe and pinpoint landing, and others.

Mark W. Maimone is a navigation and machine vision researcher at the Jet Propulsion Laboratory (JPL). He earned his Ph.D. in computer science from the Computer Science Department of Carnegie Mellon University, Pennsylvania, in 1996, and was then a postdoctoral research associate at Carnegie Mellon's Robotics Institute. Since starting at JPL in 1997, he has worked on the several Mars rover research projects and a vision system for inspection of the Chernobyl reactor. As a member of the 2003 Mars Exploration Rover (MER) flight software team, he developed the vision and navigation subsystems for the MER vehicles. He is now part of the MER ground operations team and is developing the autonomous navigation software for the Mars Science Laboratory rover, NASA's next Mars rover mission.

Larry Matthies is a senior research scientist at the Jet Propulsion Laboratory (JPL) and is the supervisor of the computer vision group (3474) in the mobility and robotic systems section. He earned his Ph.D. in computer science at Carnegie Mellon University, Pennsylvania, in 1989. He is also an adjunct professor in computer science at the University of Southern California and is a member of the editorial boards for the *Autonomous Robots* journal and the *Journal of Field Robotics*. His research interests are in passive and active sensing systems for three-dimensional perception, motion estimation, terrain classification, and related aspects of perception for autonomous vehicles.

Address for Correspondence: Mark Maimone, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109 USA. E-mail: mark.maimone@jpl.nasa.gov.