

A General Approach to Terrain Relative Navigation for Planetary Landing

Andrew E. Johnson^{*}, Adnan Ansar[†] and Larry H. Matthies[‡]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

Nikolas Trawny[§], Anastasios I. Mourikis^{**} and Stergios I. Roumeliotis^{††}
University of Minnesota, Minneapolis, MN 55455

We describe an algorithm for navigation state estimation during planetary descent to enable precision landing. The algorithm automatically produces 2D-to-3D correspondences between descent images and a surface map and 2D-to-2D correspondences through a sequence of descent images. These correspondences are combined with inertial measurements in an extended Kalman filter that estimates lander position, velocity and attitude as well as the time varying biases of the inertial measurements. The filter tightly couples inertial and camera measurements in a resource-adaptive and hence real-time capable fashion. Results from a sounding rocket test, covering the dynamic profile of typical planetary landing scenarios, show estimation errors of magnitude 0.16 m/s in velocity and 6.4 m in position at touchdown. These results vastly improve current state of the art and meet the requirements of future planetary exploration missions.

I. Introduction

NASA's roadmap for solar system exploration¹ contains landing and sample return missions to the Moon, Mars, Europa, Titan, comets and asteroids, and states that autonomous precision landing systems are a key requirement for these missions. Pin-point landing would provide safe, affordable access to landing sites with the highest science value. Exploration of the Moon also requires precision landing to enable targeted landing of payloads near to each other and in-situ resources. The 10-meter position error possible with the algorithm presented in this paper also facilitates landing very near fresh surfaces features on Europa and Enceladus.

We present three new technologies for surface-relative state estimation during planetary descent. The first technology is a computer vision algorithm that automatically matches surface landmarks in known locations with respect to a surface map to landmarks extracted from a visual image taken during descent. These matches provide a measurement of horizontal position relative to a known landing site specified in the surface map thereby enabling precise position estimation. The second technology is a computer vision algorithm that automatically matches image features through a sequence of descent images. This persistent feature tracker can handle large changes in scale and orientation that are typical during planetary descent. The final technology is an extended Kalman filter that combines mapped landmarks, persistent feature tracks and inertial measurements to enable robust position and attitude determination at high rates during the entire descent.

As shown in Figure 2, the Vision Aided Inertial Navigation (VISINAV) filter starts by propagating inertial measurements (accelerations and angular rates) from an inertial measurement unit (IMU). When image collection begins, landmarks are extracted from a descent image and then matched to an ortho-rectified image of the scene that is coregistered with a digital elevation map (c.f. Figure 1). The output of this process is a set of mapped landmark (ML) measurements; MLs are required to make the position of the lander observable by the filter. The filter processes features tracks obtained from the sequence of images in two distinct ways. For a few persistent features (PF) the filter adds the global position of each feature to the filter state so that it can be estimated. This has the

^{*} Principal Engineer, Mobility and Robotics Systems Section, Member AIAA

[†] Senior Engineer, Mobility and Robotics Systems Section

[‡] Senior Research Scientist, Mobility and Robotics Systems Section.

[§] Graduate Research Assistant, Computer Science and Engineering Department, Member AIAA

^{**} Graduate Research Assistant, Computer Science and Engineering Department

^{††} Assistant Professor, Computer Science and Engineering Department

effect of integrating out noise in the feature track while adding stability to the filter. The remaining opportunistic features (OF) are processed in a separate more computationally efficient module that computes a five degree-of-freedom motion estimate (direction of translation and rotation) between consecutive images.¹⁶ This motion estimate, instead of the OFs, is the measurement that is input into the filter. By choosing the number of PFs versus OFs to process a balance can be achieved between processing time and filter performance.

In the remainder of this section we describe related work to the proposed new technologies. In Section 2 we briefly describe a sounding rocket test that was used to collect data for testing these technologies. In Sections 3, 4 and 5 we describe mapped landmark matching, persistent feature tracking and the VISINAV Filter approaches in more detail while presenting results from the sounding rocket test. Finally we conclude with a description of lessons learned and future development directions.

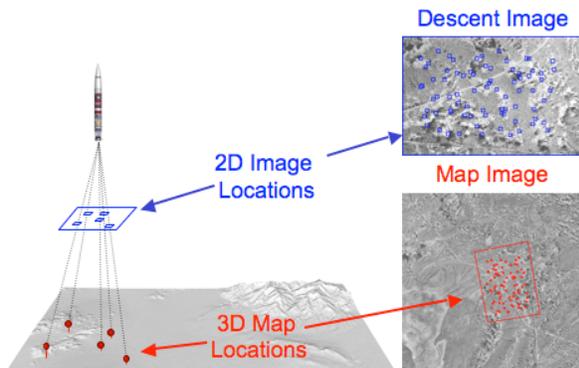


Figure 1. Mapped landmark matching concept.

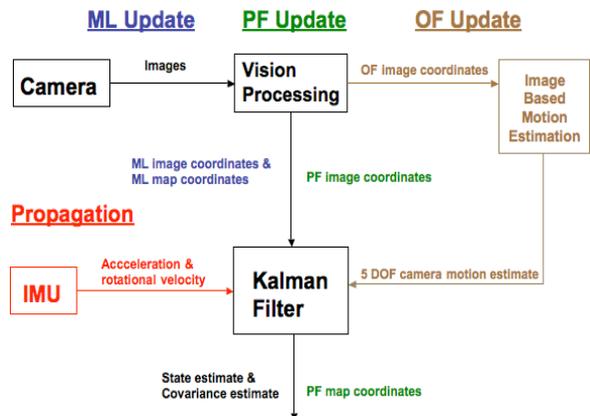


Figure 2. VISINAV block diagram.

A. Related Work

Craters are excellent landmarks for precision landing because crater shape follows a well-known geometric model, and this model can be used to reliably detect craters in imagery under a wide range of illumination conditions.⁶ This illumination invariance is important because there is no guarantee that the illumination conditions during descent imaging will be the same as the conditions when the orbital map imagery was acquired. Another advantage of crater landmarks is that they can be detected without knowledge of scene scale or relative orientation between the camera and surface.

Although craters are common on solar system bodies, targeted landing sites are not guaranteed to have enough (if any) craters to enable position estimation. Therefore a more general landmark representation is required. One possible approach that does not rely on a geometric landmark model is the Scale Invariant Feature Transform (SIFT)¹² which converts image data into a feature vector or keypoint that can be matched between a map and an image. All that is required to generate descriptive SIFT keypoints is unique image texture; unlike with craters, a model of terrain shape is not required or used. SIFT keypoints share the scale invariance properties of crater landmarks, but they do not have an intrinsic capability to handle illumination variations between the descent image and map. As with craters, SIFT keypoints can be matched without knowledge of the spacecraft altitude or orientation relative to the surface map.

Although the scale and rotational invariance properties of SIFT keypoints are useful properties in many recognition tasks, they are not required during planetary landing. Landers always have an altimeter and gyroscopes that provide on-board measurements of altitude and surface relative attitude (given an initial attitude estimate from celestial measurements). Using these measurements, descent images can be warped to a local level coordinate system so that normalized correlation¹¹ can be applied to find matches between the descent image and map. Image correlation is fundamentally more efficient than SIFT because it does not require a search over image scale. Because it has been implemented in numerous military and commercial products⁴, image correlation is more mature and widely accepted than SIFT. Since image correlation does not rely on a specific geometric model of landmarks to enable detection, as in the case for crater landmarks, it can be applied to any planetary terrain of sufficient texture making it a very general approach. Due to its maturity, efficiency, and general applicability, we have chosen 2D image correlation as the starting point for landmark detection and tracking.

The persistent feature tracker we have developed to augment mapped landmarks combines area correlation with homography image warping to account for all image motions that are possible during flight over a distant and flat scene. Unlike the widely used KLT tracker¹³ our tracker uses area correlation instead of gradient descent to find feature matches and therefore can handle larger motion between images. This tracker is also a significant improvement over the trackers used on recent flight missions. The tracker used in the Descent Image Motion Estimation Subsystem (DIMES) on the 2003 Mars Exploration Rover (MER)⁵ tracked four features across three descent images. Our new tracker generates tens of features tracks and maintains them through hundreds of images without relying on the use of on board measurements of attitude and altitude as was done in the DIMES algorithm. Deep Impact² located the nucleus of a comet in an onboard image but it did not perform tracking because it did not directly compare consecutive images frames to locate the nucleus. Instead it located the nucleus against the dark space background using centroiding of image brightness independently for each image.

Only a few recursive estimation approaches that utilize measurements of a priori known features have been proposed in the literature. Crassidis⁷ uses a statistical (zero-acceleration) model to propagate the pose estimate between ML observations. However, the use of a statistical model (rather than inertial measurements), limits the applicability of such an approach to maneuvers with slow dynamics that occur, for example, during spacecraft rendezvous and docking. Wu²¹ and Ivey⁸ fuse inertial measurements with observations of artificial rectangular targets and heading measurements from a magnetometer. In their work, the authors use measurements both of the coordinates of a target's projection, and of the area of this projection. Area measurements, though, may be difficult or unreliable when dealing with real planetary imagery, where visual features are less structured. Finally, Bayard³ fuses inertial measurements with bearing measurements to MLs, but the spacecraft's attitude is assumed to be perfectly known, which is not a valid assumption for EDL.

In addition to processing bearing angles to MLs, in this work we also utilize measurements to features that can be tracked in the images, but whose position in the global frame is not known in advance (OFs). The standard method of treating such features is to include their positions in the state vector, and to estimate them along with the vehicle's trajectory. This is the well-known Simultaneous Localization and Mapping (SLAM) problem. SLAM with visual measurements and inertial sensors has recently attracted significant interest in the robotics community^{7, 10, 9}. However, the need to maintain the landmark estimates in SLAM results in increased computational complexity (quadratic in the number of features for EKF-SLAM). Moreover, the main benefit of performing SLAM is the ability to achieve "loop closing" when revisiting an area. Due to the nature of EDL trajectories, loop closing is not an important consideration, and thus the quadratic computational complexity of SLAM does not appear to be justified in the context of precision landing. However as discussed above adding the depth to features to the state for a few features does improve filter stability and that is the approach we follow.

II. Sounding Rocket Test Description

In April 2006, a sounding rocket flight was conducted at White Sands Missile Range (WSMR)^{17,18}. Data was collected to test terrain-relative navigation algorithms by adding a descent camera to a rocket payload that already contained an IMU and GPS. The flight reached a maximum altitude of 123 km, deployed its parachute at 5.2 km AGL at launch plus 429 s, and landed 78 km downrange after a flight time of 794 s. The descent rate on the parachute was approximately 10 m/s. A commercially available analog machine vision camera (Pulnix TM-9701) provided descent imagery from parachute deployment to landing at 30 frames/s with a resolution of 768×484 pixels, 8 bits/pixel, and a field of view of 38×24°. This was downlinked in real-time during flight over an S-band telemetry channel and recorded on the ground. A GPS time tag was assigned to each image using a commercial timecode generator. 50 Hz IMU data and 10Hz GPS data was also downlinked. The common GPS time assigned to measurements was used to sync the data. After the flight, images were post processed to extract mapped landmarks and feature tracks. These image measurements were then combined with inertial measurements in the VISINAV filter. Results from the test are given in the following sections.

III. Mapped Landmark Matching

The mapped landmark matching algorithm automatically matches surface landmarks in known locations with respect to a surface map to landmarks extracted from a visible image taken during descent. These matches provide a measurement of position relative to a known landing site specified in the surface map, thereby enabling precision landing. The algorithm is general because it handles images taken at different scale and resolution relative to the map and can produce mapped landmark matches for any planetary terrain of sufficient texture. The algorithm has a built-in robustness due to the large number of mapped landmarks generated per image, which allows for automatic

detection and elimination of bad matches. The algorithm is accurate because it uses normalized correlation of grayscale images that produces accurate and precise sub-pixel matches.

Planetary descent and landing has some unique challenges that make it difficult to apply image-based correlation directly to the precision landing problem. However, these challenges (low terrain texture, large changes in scale and attitude) were also encountered during the development of DIMES. Our solution combines the space proven approaches from DIMES⁵: (i) feature selection to maximize the use of image texture and (ii) image rectification via the homography transform to deal with scale and orientation between image and map with (iii) 2D correlation. To improve computational efficiency, the mapped landmark algorithm is broken into two sub-algorithms: FFT Map Matching, which matches a single large template by correlation in the frequency domain, followed by Mapped Landmark Refinement, which matches many small templates by correlation in the spatial domain.

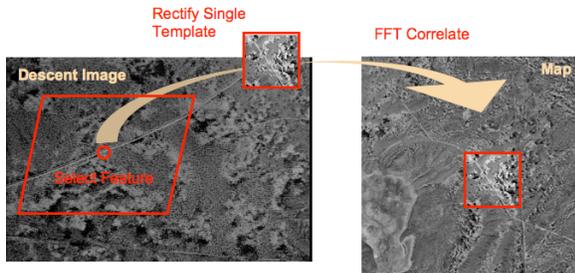


Figure 3. FFT Map Matching algorithm.

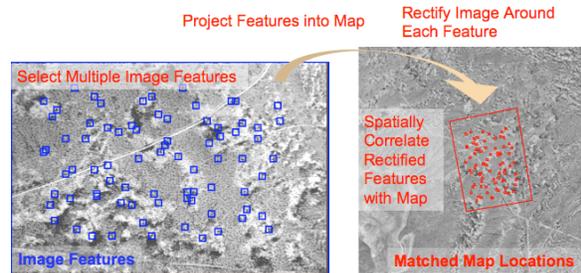


Figure 4. Mapped Landmark Refinement algorithm.

- Simple algorithms
- Direct extension of MER-DIMES
- Works for all terrain appearances
- Significant amount of testing and validation

Figure 5. Data sets used for testing the mapped landmark matching algorithm.

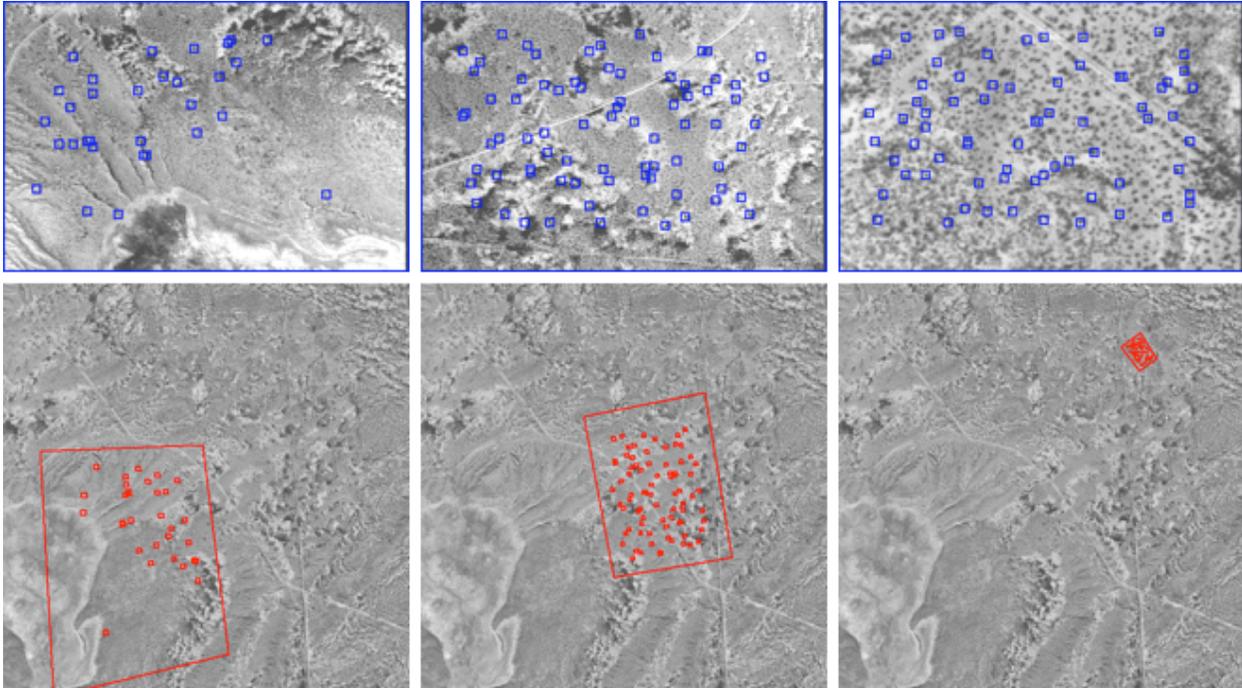


Figure 6. Mapped landmark matches between the map (red) and the descent image (blue) for three altitudes (left is 1600 m AGL, middle is 900 m AGL, right is 230 m AGL).

The first step in FFT Map Matching (c.f. Figure 3) is to select a single high-contrast template in the descent image using a standard interest operator. This template is then warped using a homography computed using onboard estimates of altitude and attitude (handles changes in scale, orientation and perspective effects). The template and map image are then normalized using a mean and variance filter that essentially acts as a high pass filter. The normalized template and map are then matched with correlation using the Fast Fourier Transform technique (FFT). Metrics on the correlation peak are then used to determine if the match is valid. If the match is valid the horizontal position of the lander relative to the map is computed using the match location, altitude and attitude measurements.

In Mapped Landmark Refinement (c.f., Figure 4), this position estimate is used to initialize the search for multiple matches between the descent image and map. In this stage multiple small templates that cover the image are selected from the descent image using a standard interest operator. These templates are warped to the map image using a homography computed for each template using the camera parameters and the position and attitude of the camera from the FFT stage. These warped templates are then correlated with the map using spatial correlation. If the correlation peak for a template is valid, a 2D to 3D mapped landmark correspondence is established between the 2D template location in the descent image and the 3D position obtained from the match location (horizontal dimensions) and the elevation value from the elevation map (vertical dimension) that is co-registered with the map image.

The Mapped Landmark algorithm has been tested with numerous data sets (cf. Figure 5) including image streams collected on a parachute drop platform (imaging from 40 km), and a sounding rocket (imaging from 4 km on down). The algorithm has also been tested with images of planetary surfaces and has shown robustness to changes in illumination and scene appearance. Furthermore a high fidelity camera simulator that models the DIMES descent camera was developed during MER²⁰ and this simulator has been used in a Monte Carlo analysis using Mars orbital images as input.

During the sounding rocket test flight in April 2006, two sets of mapped landmarks were extracted from imagery. For each set, the map image was a Digital Orthoimagery Quarter Quadrangle image produced by the USGS in 2001. Additional elevation data from the 1 arc second finished Shuttle Radar Topography Mission dataset provided 3D position coordinates for mapped features. In the first set, the entire 7x8 km map image was resampled to ~7m/pixel while for the second set the map was cropped to 2 km x 2 km and used at its base resolution of ~1 m/pixel. Figure 6 shows three examples of the mapped landmarks detected and extracted from the second set at altitudes of 1600 m above ground level (AGL), 900 m AGL and 230 m AGL. In these examples a single map is used for all the matches indicating the ability to match under a 5x change in scale as well as any roll about the optical

axis. There is clearly a change in appearance between the descent image and map, which is not unexpected given the five years between map and descent image acquisition. Nevertheless, numerous mapped landmarks are generated for each image, which demonstrates the robustness of the approach.

IV. Persistent Feature Tracking

In image feature tracking, a scene feature (e.g., the corner of an object) is first identified in an initial image. Tracking determines the location of this feature in consecutive images. The signal of the feature track is the direction and length of the feature track. The noise is the uncertainty in the location of the feature in the consecutive images, which is essentially constant as long as the feature in the initial image is being tracked. Consequently, the longer the feature track, the higher the signal to noise of the feature track. Since the signal-to-noise of the feature tracks drives the accuracy of the derived estimates for velocity and feature depth in the VISINAV filter, maintaining long-term feature tracks is very important.

Traditional feature trackers based on 2D correlation quickly lose track of the feature if there are rotations or changes in scale in the imagery. This is a fundamental limitation of 2D correlation. During planetary landing there are significant changes in image scale and orientation over short time periods due to the dynamic flight environment. Attitude changes can be large due to parachute oscillations and maneuvers during powered terminal descent. In addition, changes in vehicle altitude during landing or controlled flight will induce image scale changes.

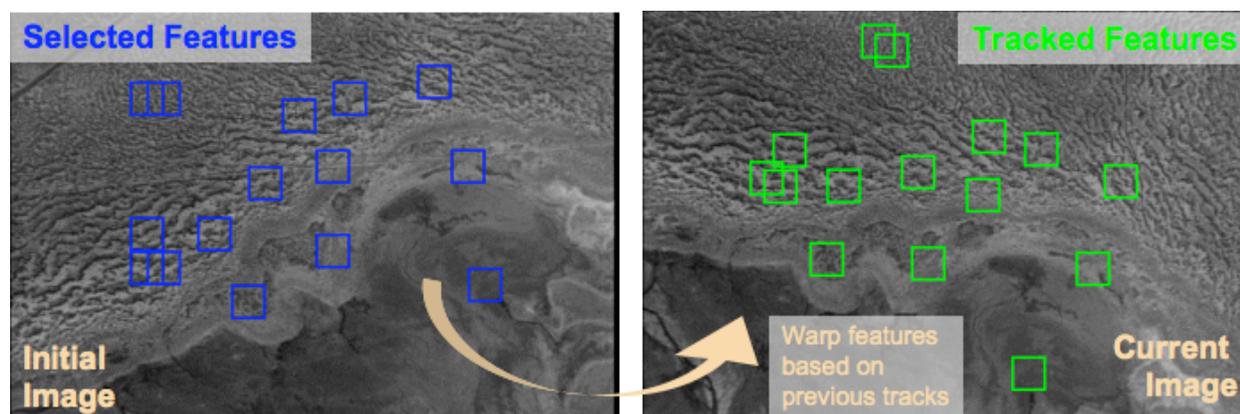


Figure 7. Feature tracking concept that applies image warping and local spatial correlation to maintain persistent feature tracks.

We have developed a feature tracker algorithm that is an improvement over 2D correlation because it can persistently track a feature through significant changes in image scale caused by motion toward the ground and significant changes in image orientation induced by changes in attitude of the vehicle during flight. Our solution starts with a 2D correlation approach to tracking and adds an image homography warping stage, to maintain persistent tracks. We assume that the scene is relatively flat in our approach, so a homography can be applied to warp the initial image to the consecutive image. If there are a set of features in one image and a set of matching features in a subsequent image, then a homography can be computed that aligns the features. The feature tracker first selects a set of features in the initial image using an interest operator. These features are then tracked between the first and second image using local spatial correlation of image windows. A homography is computed from the feature tracks. The homography is then used to warp each of the selected features to the second image and spatial correlation is applied again on the warped templates to improve the feature tracks. A final homography is then computed from these tracks.

To track between the first and third image, this homography is applied to the feature locations in the first image to map them into the second image. A template around each location in the second image is then spatially correlated with the third image. This step compares second image data to third image data, but it also defines a correspondence between the first and third image. A homography is computed from these correspondences, and it is used to warp the initial feature templates into the third image. Spatial correlation is then applied between these warped templates and the third image which is a direct comparison of image data from the first image to the third image. This step is critical to prevent accumulation of errors and generate truly persistent feature tracks. This process is then repeated to generate tracks between the initial and each subsequent image.

This new algorithm has been implemented in C++. Timing experiments show that it can track image features at rates up to 30 Hz on a 400 MHz R12000 SGI

O2 processor. This processing rate is fast enough to support real-time feature tracking during descent and landing. The algorithm has been tested with numerous data sets including image streams collected from an autonomous remote controlled helicopter (imaging from 50 m), a manned helicopter (imaging from 1000 m), a parachute drop platform (imaging from 40 km), and a sounding rocket (imaging from 4km). These tests with real and relevant data cover a large range of altitudes and attitude dynamics.

The feature tracker was applied to all of the images acquired from the sounding rocket test in April 2006. These results show tracking through complete 360° roll about the optical axis with up to 12° of roll between consecutive images. The tracker performs well with up to 2x change in scale between images, which is important when tracking features near touchdown. It is also able to reacquire features after they have left the field of view, which is important when there are off-nadir oscillations during the parachute phase. Some representative results from this test are shown in Figure 8.

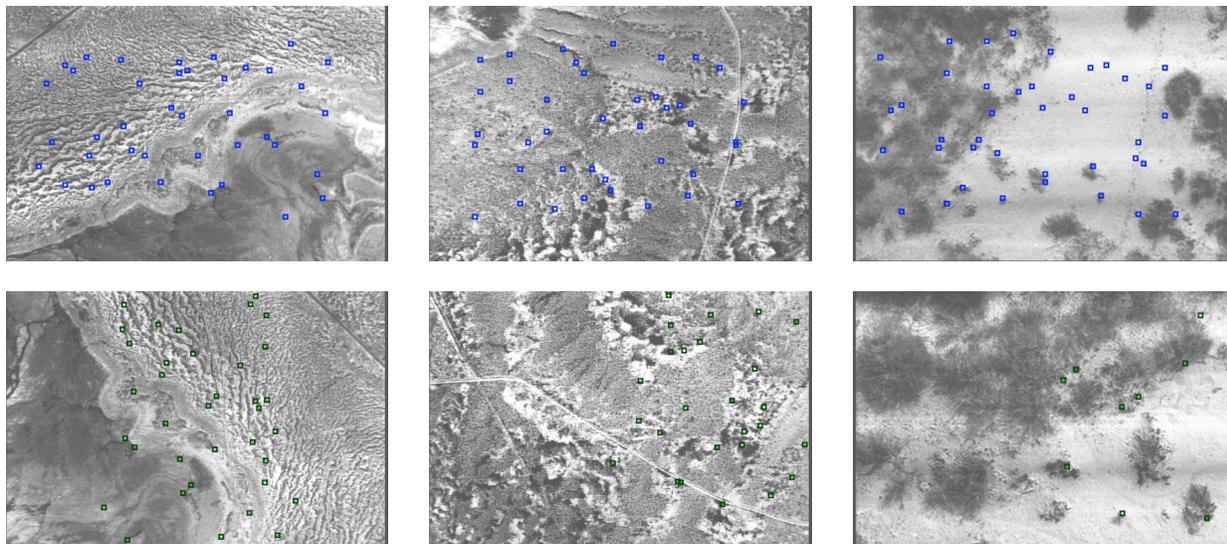


Figure 8. Three feature tracking results from the sounding rocket test. Selected features are shown in the top row (in blue) and tracked features are shown in the bottom row (in dark green). The left pair shows tracking under large roll, the middle shows tracking with a large off-nadir angle and the right result shows tracking with a 2x change in scale near touch down.

V. Vision Aided Inertial Navigation Filter

The Vision Aided Inertial Navigation (VISINAV) filter estimates current vehicle position, velocity and attitude states (in Earth Centered Earth Fixed coordinates) by statistical reduction of onboard measurements from an IMU, radar altimeter, and descent imagery. The filter also estimates 3 axis accelerometer and gyro measurement biases, which is critical to obtain the accuracy desired for pin-point landing. Sub-functions within the filter are integration models for the state vector and the covariance matrix, an observations model, and a state update module. The filter is initialized with an a priori state vector and state covariance before launch. The integrator propagates the state and covariance in time, incorporating all non-gravitational accelerations experienced by the vehicle by directly adding the IMU accelerometer measurements to the modeled gravitational acceleration while at the same time accounting for the effects of earth rotation. When each new observation (ML, PF, OF) becomes available, the filter produces an update to the state estimate that combines the observation with the propagated state at that point in time.

The observation model for map landmarks is a direction vector from the camera to the mapped landmark; these update the filter's position, velocity, and attitude states. The map coordinates of persistent features are added to the state vector, initialized by triangulation from the first two measurements; these also update position and velocity states. Opportunistic features are not added to the state vector; measurements consist of pairs of corresponding image coordinates of features tracked between two consecutive images. These directly update just velocity state. Further details on the VISINAV filter can be found in other recent papers^{14,15,19}.

The sounding rocket data was processed though the VISINAV filter. Between launch and touchdown the filter propagated angular rates and accelerations from the IMU. The filter incorporated ML and OF measurements as

follows. After main parachute deployment a payload door opened giving the camera a view of the ground. Between 3800m and 3100m AGL mapped landmarks were extracted at 3Hz from the imagery and then again between 1600m and 230m AGL at 1Hz. The gap in ML's was used to illustrate the error growth in the filter when these measurements are not available. Although features were extracted from the entire image sequence, the filter only processed OFs from 330m AGL to the ground. The filter did not process altimetry because it was not available. The filter did not add the depth to features into the filter state; all features tracks were processed as opportunistic features.

The results of the VISINAV output were compared to GPS; plots of resulting errors for position and velocity are shown in Figure 9. The errors are shown in blue and the 3 sigma covariance bounds from the filter are shown in red. No ground truth estimates of attitude are available, so only the 3-sigma covariance bounds are shown in Figure 9. The errors stay well within the covariance bounds indicating a well tuned filter. The errors and covariance bounds decrease quickly after the first set of MLs becomes available. The errors grow in the interval between ML sets (open loop integration) and then when the second set of MLs becomes available the errors and covariance bounds decrease again. Table 1 shows the errors at touchdown. The RMS errors are 6.4 m in position and 0.16 m/s in velocity. While these results are already very good, the camera, IMU, and telemetry channel contained several sources of noise that will be improved in subsequent development, so future performance should improve by at least a factor of 2.

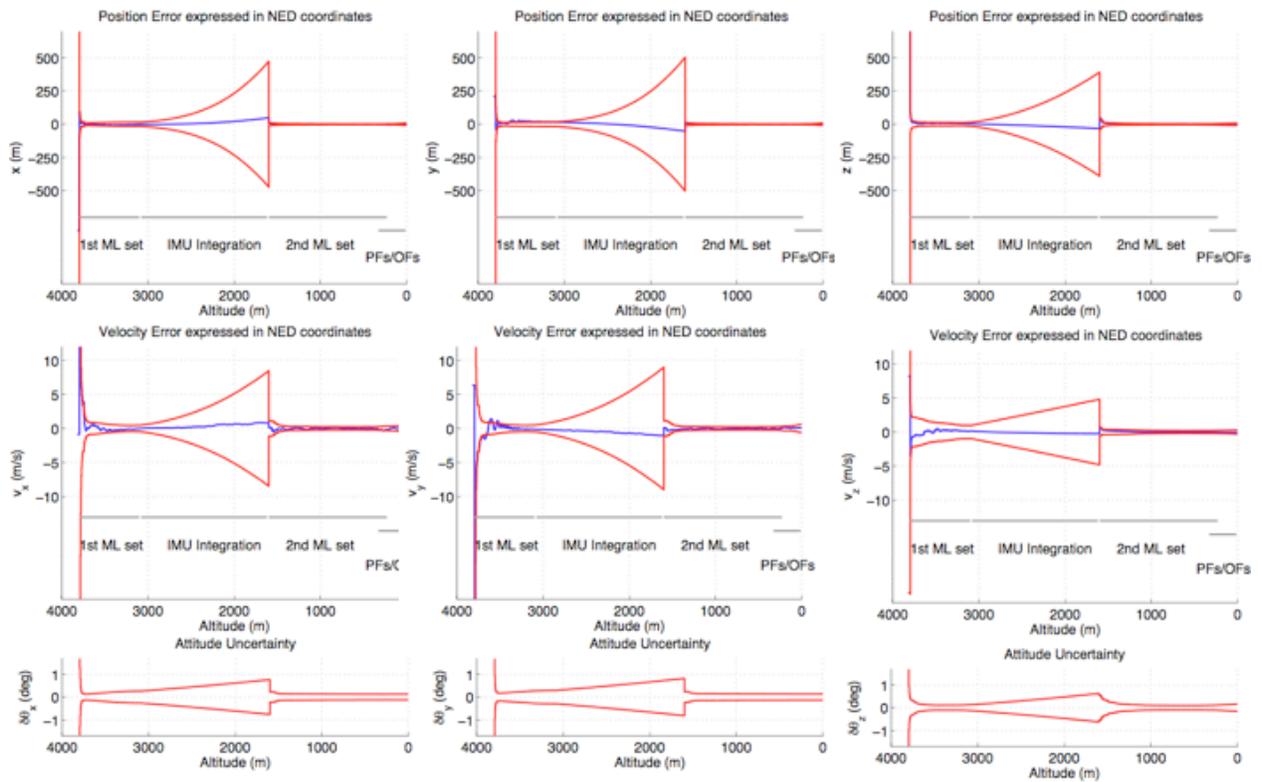


Figure 9. VISINAV filter errors (blue) and covariance bounds (red) during parachute descent.

Table 1. VISINAV errors at touchdown.

	North	East	Down	RSS
Position Error at Touch Down (m)	-1.7	0.1	-6.2	6.4
Velocity Error at Touch Down (m/s)	0.13	0.06	-0.07	0.16
Attitude 3 Sigma Covariance Bounds at Touch Down (degrees)	0.2	0.2	0.2	0.28

VI. Conclusion

In this paper, we have presented the analysis and experimental validation of a vision-aided inertial navigation algorithm for planetary landing applications. The algorithm combines mapped landmarks, persistent features and inertial measurements to generate high rate estimates of lander position, velocity and attitude. Results from a sounding rocket test, covering the dynamic profile of typical planetary EDL scenarios, showed estimation errors of magnitude 0,16 m/s in velocity and 7,6 m in position at touchdown. This level of accuracy is very significant as it enables a whole new class of space science missions. For example, landers can be placed on top small geological features that can be investigated directly from the landed platform using deep drills or robotic arms. 10 m precision landing also enables landing near the hazardous terrain that is typical of the most scientifically rich sites (e.g., layered cliff walls on Mars and ice fissures on Europa).

Acknowledgments

The Mars Technology Program funded development of the described algorithms and the New Millennium Program made the sounding rocket test possible. The work described in this publication was performed at the Jet Propulsion Laboratory, California Institute of Technology and the Department of Computer Science and Engineering, University of Minnesota, under contract from the National Aeronautics and Space Administration.

References

- ¹ 2006 Solar System Exploration Roadmap for NASA's Science Mission Directorate http://solarsystem.nasa.gov/multimedia/downloads/SSE_RoadMap_2006_Report_FC-A_opt.pdf.
- ² Bank, T., Frazier, W., Blume, W., Kubitschek, D., Null, G., Mastrodemos, N., and Synnot, S., "Deep Impact: 19 Gigajoules Can Make Quite an Impression", Proc. 24th Annual AAS Guidance and Control Conference, Feb 2001.
- ³ Bayard, D. S. and Brugarolas, P. B., "An Estimation Algorithm for Vision Based Exploration of Small Bodies in Space," *Proc. 2005 American Control Conference*, Vol. 7, Portland, Oregon, June 2005, pp. 4589–4595.
- ⁴ Carr, J.R. and Sobek, J.S., "Digital Scene Matching Area Correlator (DSMAC)," *Image Processing For Missile Guidance*, Proc. of the Society of Photo-Optical Instrumentation Engineers, Vol. 238, 1980.
- ⁵ Cheng, Y., Johnson, A. and Matthies, L. "MER-DIMES: A Planetary Landing Application of Computer Vision," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005.
- ⁶ Cheng, Y. and Ansar, A., "Landmark Based Position Estimation for Pinpoint Landing on Mars," *Proc. IEEE Int'l Conf. on Robotics*, 2005.
- ⁷ Crassidis, J. L., Alonso, R. and Junkins, J. L., "Optimal Attitude and Position Determination from Line-of-Sight measurements," *Jour. Astronautical Sciences*, Vol. 48, no. 2–3, 2000, pp. 391–408.
- ⁸ Ivey, G. F. and Johnson, E., "Investigation of Methods for Simultaneous Localization and Mapping Using Vision Sensors," *Proc. AIAA Guidance, Navigation, and Control Conference*, no. AIAA 2006-6578, Keystone, CO, Aug. 2006.
- ⁹ Kim, J. and Sukkarieh, S., "Autonomous Airborne Navigation in Unknown Terrain Environments," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 40, No. 3, July 2004, pp. 1031–1045.
- ¹⁰ Langelaan, J. W., "State Estimation for Autonomous Flight in Cluttered Environments," Ph.D. dissertation, Stanford University, Department of Aeronautics and Astronautics, 2006.
- ¹¹ Lewis, J. P. "Fast Template Matching," *Vision Interface*, 1995, pp. 120-123.
- ¹² Lowe, D., "Distinctive Image Features from Scale-Invariant Keypoints," *Int'l Jour. Computer Vision*, Vol. 60, No. 2, 2004, pp. 91-110.
- ¹³ Lucas, B. and Kanade, T., "An Iterative Image Registration Technique with an Application to Stereo Vision," *Proc. 7th Int'l Joint Conf. on Artificial Intelligence*, 1981.
- ¹⁴ Mourikis, A. I. and Roumeliotis, S. I., "A Multi-State Constraint Kalman Filter for Vision-Aided Inertial Navigation," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA '07)*, 2007, pp. 3565-3572.
- ¹⁵ Mourikis, A. I., Trawny, N., Roumeliotis, S. I., Johnson, A. and Matthies, L., "Vision Aided Inertial Navigation for Precise Planetary Landing: Analysis and Experiments," *Proc. Robotics Systems and Science Conference*, June 2007.
- ¹⁶ Roumeliotis, S., Johnson, A., and Montgomery, J., "Augmenting Inertial Navigation with Image-Based Motion Estimation," *IEEE Int'l Conf. Robotics and Automation (ICRA 2002)*, Washington D.C., 2002, pp. 4326–33.
- ¹⁷ Seybold, C., Chen, G., Bellutta, P., Johnson, A., Matthies, L. and Thurman, S. "Suborbital Flight Test of a Prototype Terrain- Relative Navigation System," *Proc. AIAA Infotech@Aerospace Conference*, AIAA-2007-02896, 2007.
- ¹⁸ Thurman, S., Matthies, L., Corliss, J. and Johnson, R. K., "Space Flight Test of Vision- Guided Planetary Landing System," *Proc. AIAA Infotech@Aerospace Conference*, AIAA-2007-02767, 2007.
- ¹⁹ Trawny, N., Mourikis, A., Roumeliotis, S., Johnson, A. and Montgomery, J., "Vision-Aided Inertial Navigation for Pin-Point Landing using Observations of Mapped Landmarks," *Journal of Field Robotics*, 2006.
- ²⁰ Willson, R., Johnson, A. and Goguen, J., "MOC2DIMES: A Camera Simulator for the Mars Exploration Rover Descent Image Motion Estimation System," *Proc. 8th Int'l Symp. Artificial Intelligence, Robotics and Automation in Space*, 2005.
- ²¹ Wu, A., Johnson, E. and Proctor, A., "Vision-Sided Inertial Navigation for Flight Control," *Proc. AIAA Guidance, Navigation, and Control Conference*, no. AIAA 2005-5998, San Francisco, CA, Aug. 2005.