Real-Time Terrain Relative Navigation Test Results from a Relevant Environment for Mars Landing

Andrew E. Johnson[1], Yang Cheng[2], James Montgomery[3], Nikolas Trawny[4], Brent Tweddle[5], and Jason Zheng[6]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

Terrain Relative Navigation (TRN) is an on-board GN&C function that generates a position estimate of a spacecraft relative to a map of a planetary surface. When coupled with a divert, the position estimate enables access to more challenging landing sites through pin-point landing or large hazard avoidance. The Lander Vision System (LVS) is a smart sensor system that performs terrain relative navigation by matching descent camera imagery to a map of the landing site and then fusing this with inertial measurements to obtain high rate map relative position, velocity and attitude estimates. A prototype of the LVS was recently tested in a helicopter field test over Mars analog terrain at altitudes representative of Mars Entry Descent and Landing conditions. TRN ran in real-time on the LVS during the flights without human intervention or tuning. The system was able to compute estimates accurate to 40m (3 sigma) in 10 seconds on a flight like processing system. This paper describes the Mars operational test space definition, how the field test was designed to cover that operational envelope, the resulting TRN performance across the envelope and an assessment of test space coverage.

I. Introduction

TERRAIN Relative Navigation (TRN) is an on-board spacecraft GN&C function that generates a position estimate of a spacecraft relative to a map of a planetary surface. The TRN measurement tells the vehicle its position relative to the center of the landing ellipse and also where the vehicle is relative to hazards identified in the map prior to landing. On Mars, TRN can be used different ways depending on the amount of fuel available for a divert (see Figure 1). The Mars Science Laboratory (MSL) Entry Descent and Landing (EDL) system carried enough fuel to slow for soft landing and also perform a 300+ meter divert to avoid recontact with the backshell. Given a TRN position measurement, this divert can also be used to avoid 100m scale landing hazards thereby enabling selection of landing sites that have scientifically interesting terrain relief and were not selectable in the original implementation of MSL EDL. Mars scientists see great value in adding this capability to the proposed Mars 2020 lander mission [9]. In future Mars landing architectures when the lander has enough fuel to fly out aerodynamic dispersions TRN can be used to target the center of the landing ellipse and enable pin-point landing. This capability would be useful for retrieving a cache of samples during a potential Mars Sample return mission or for deploying resources for an eventual human mission to Mars.

TRN technology development has been ongoing for the past decade. At JPL, TRN research started shortly after the development of the Descent Image Motion Estimation System (DIMES) for the Mars Exploration Rovers mission [4]. The DIMES software was first extended to match between a descent image and a large orbital image and co-registered digital elevation map (DEM) [5]. This software was

[1] Principal Member of Technical Staff, Guidance and Control Section, MS 198-235, 4800 Oak Grove Dr., AIAA Member.
[2] Member of Technical Staff, Mobility and Robotic System Section, MS 198-235, 4800 Oak Grove Dr.
[3] Member of Technical Staff, Guidance and Control Section, MS 198-235, 4800 Oak Grove Dr.
[4] Member of Technical Staff, Guidance and Control Section, MS 198-235, 4800 Oak Grove Dr., AIAA Member.
[5] Member of Technical Staff, Guidance and Control Section, MS 198-235, 4800 Oak Grove Dr., AIAA Member.
[6] Member of Technical Staff, Flight Electronics and Software Systems Section, MS 198-326, 4800 Oak Grove Dr.

then integrated with an extended Kalman filter that fused image measurements with inertial data and demonstrated ten meter landing precision when processing of a sounding rocket data set [1]. In the Autonomous Landing and Hazard Avoidance Technology Project, the TRN software was tested off-line using extensive data sets collected over different terrain types and illumination conditions [2]. This testing led to a modularization of the TRN software and the ability to handle illumination changes and primitive body applications [1].

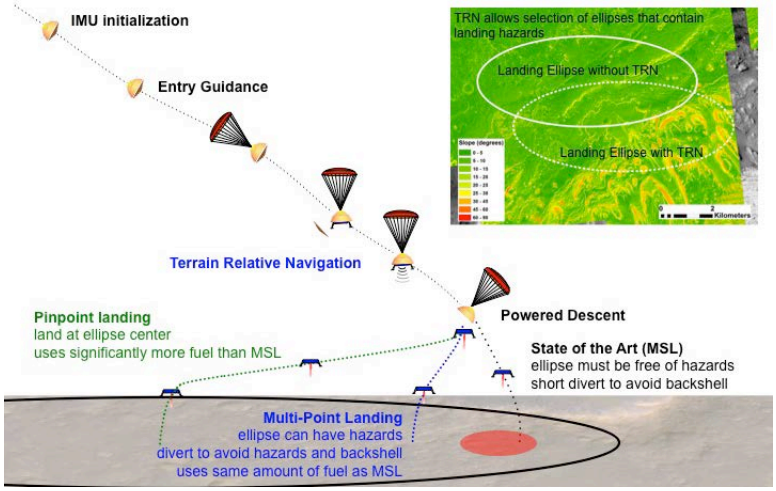

Figure 1. Next Generation Mars EDL capabilities.

These initial studies with real and simulated data led to the development of the initial Lander Vision System concept [6], a bolt on TRN and hazard detection system for Mars EDL. Initial studies by the Mars 2020 Science Definition Team (SDT) [9] indicated that when TRN is added to the MSL EDL architecture multiple specific landing sites (e.g., East Margaritafer, North East Syrtis) that were too hazardous for MSL could potentially be targeted by M2020. On-board active hazard detection was also stated to have benefits, but the report indicated that TRN had a greater impact on terrain accessibility. Given the clear infusion path laid out by the SDT, the work on LVS was redirected to focus on TRN. Working with the Mars 2020 (M2020) EDL team, the driving requirements for LVS were established and a hardware prototype composed of a camera, Inertial Measurement Unit (IMU) and a dedicated compute element was designed and built [7]. Real-time software and FPGA firmware was developed on the hardware to perform TRN at a 1Hz image update rate. The LVS prototype was then flown in a real-time helicopter field test conducted in February and March 2014. This paper describes Mars EDL operational requirements and test space definition, how the field test was designed to cover that operational envelope, the resulting TRN performance across the envelope and an assessment of test space coverage.

II. Terrain Relative Navigation Processing

For the Mars landing application, images are taken during parachute descent and processed to extract matches between descent images and the map. The matched landmarks are integrated with IMU data to generate high rate position, velocity and attitude estimates. A block diagram of the LVS is shown on the left in Figure 2. An IMU and camera are interfaced to a dedicated compute element composed of a general purpose flight processor (LEON3) and a Field Programmable Gate Array (FPGA). The FPGA times the sensor data and performs the highly parallel image processing algorithms. The flight processor coordinates the data flow, eliminates any spurious landmark matches and estimates the vehicle state [7]. The processor also interfaces to the spacecraft to obtain the spacecraft state to initialize TRN. The TRN estimation software on the flight processor and the image processing algorithms implemented on the FPGA are the component technologies in the LVS, while the camera, IMU and high performance compute element are considered novel but standard engineering developments.

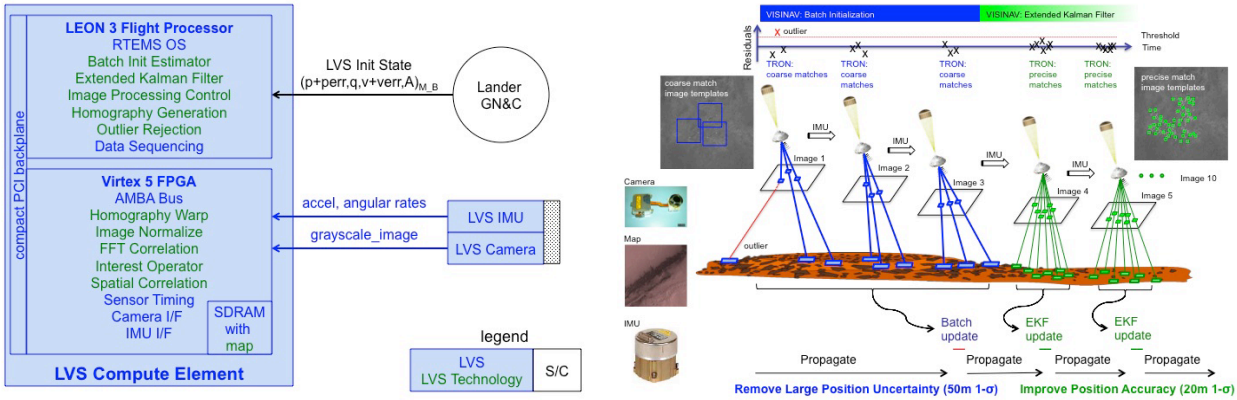

Figure 2. The LVS is composed of a camera and IMU sensors and a dedicated high performance compute element (left). The TRN image processing occurs in in a coarse to fine fashion (right). First large position errors are taken out and then accurate position estimates are generated.

As shown on the right in Figure 2, TRN operates in a coarse and then fine image matching mode. For the first three images, large image patches (3 per image) are matched to the map using frequency based correlation (left of Figure 3). All nine of these matches are input at once into a batch estimator that combines them with IMU data to estimate the large initial position error (2 to 3 km). This is followed by a refinement mode where, for each image, many small image patches (at most 100) are matched by spatial correlation to the map (right of Figure 3). These are fused with IMU data in an extended Kalman filter [8] that reduces the position estimate down to a few 10s of meters. The TRN process takes a total of 10 seconds with a total of nine to ten images being processed. Robust statistics are used though-out the image matching process to eliminate outliers that could cause large position errors. More details on the TRN processing are given in [7].

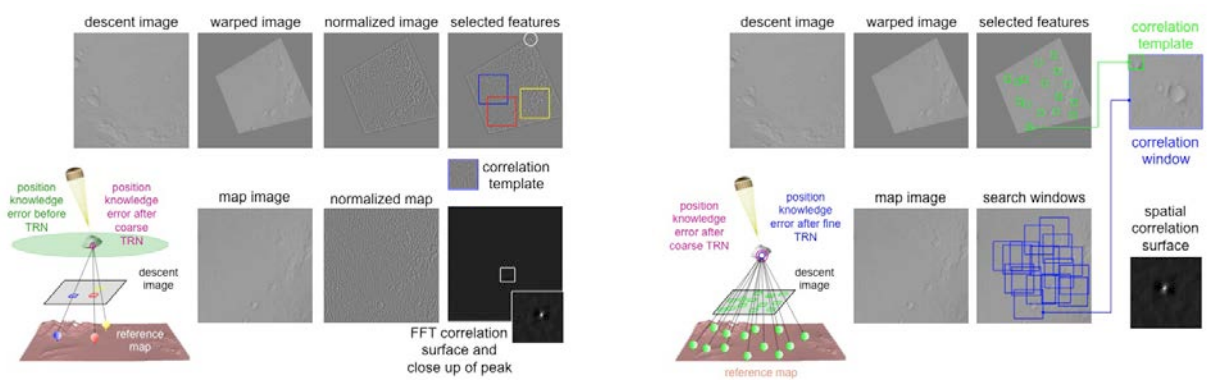

Figure 3. TRN image processing first removes large position uncertainty using frequency based correlation of a few landmarks (left) and then improves accuracy using spatial correlation of many landmarks (right).

III. Field Test Implementation

During initial development the LVS prototype (hereafter called the LVS) was tested in the lab by taking images of a flat picture of Mars. A translation stage and gimbal were used to prove that the LVS could estimate dynamic position, attitude and velocity. A single 1024x1024 pixel map was used, and the initial state of the LVS was always the same and known from metrology. Significant developments were required to convert the lab setup into a field test system that supports constantly changing vehicle state and broad terrain coverage.

Figure 4 shows the field test system block diagram, and pictures of the LVS sensor head and computing electronics. The LVS camera and IMU were mounted in a high speed gimbal along with a flash lidar (FL), extended range laser range finder (ELRF) and a short range laser range finder (SLRF).

The gimbal was used to simulate parachute attitude dynamics and is not part of the LVS. Metrology measurements were taken in the lab to compute the alignment of each sensor to the gimbal platform frame (B-frame). The gimbal has encoders that measure the angular position and rates between the moving gimbal frame (B-frame) and the fixed gimbal frame (A-frame). A commercial INS/GPS system that provides real-time position, attitude and velocity by integrating GPS and IMU measurements was mounted to the vehicle fixed side of the gimbal (A-frame). During operation the INS/GPS data and gimbal encoder values were passed to the field test avionics where they were timed and then combined with the sensor metrology to provide real-time ground truth for the position, attitude and velocity of the sensors on the gimbal. The truth data is recorded for post flight analysis of LVS performance.

When the operator commands the start of TRN processing, the initial estimate for TRN is generated from the truth data by adding noise to all of the state variables. This initial noisy estimate is passed to the LVS Compute Element (LVS CE) to initialize the TRN processing. It is also provided to the map data base (MAP DB) which uses the initial position to crop the current TRN map from a much larger map. The cropped map is large enough to contain the expected position errors and the foot print of the camera image. This cropped map is passed to the LVS CE for TRN landmark matching. TRN executes automatically and when it completes the TRN state estimates, data and intermediate data products and log files are stored back in the data logger in the field test avionics.

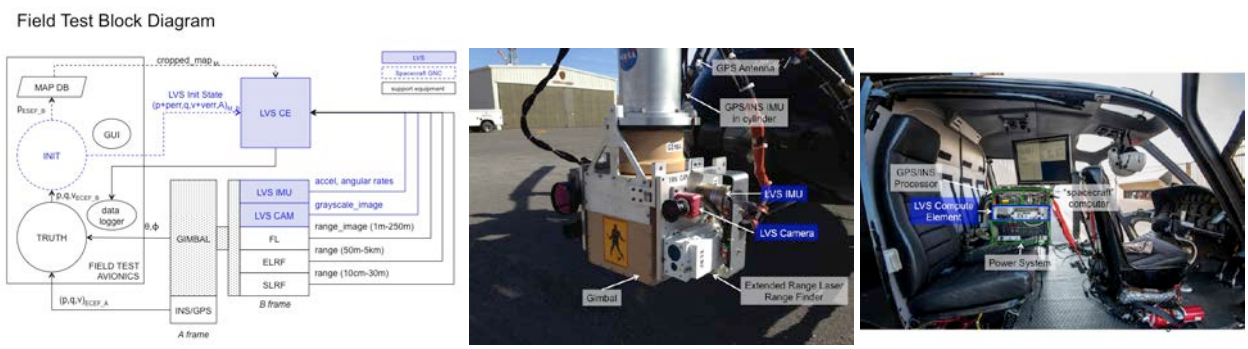

Figure 4. For the field test the prototype LVS is integrated with support equipment to collect ground truth and record test results. Block diagram of field test system (left), sensor head on helicopter (center), and cockpit with field test avionics including LVS compute element.

The field test avionics required GPS for timing and truth position and as such could not be fully tested indoors. As an intermediate step to the expensive helicopter testing, a two degree of freedom translational gantry was built above the JPL Mars Yard (Figure 5). The LVS sensor head and electronics were mounted on the gantry and moved horizontally and vertically in descending profiles. The gantry provided an affordable and controllable outside test platform for debugging and initial checkout of the ground truth system as well as testing of TRN with real illumination and terrain. During the final days of testing on the gantry the LVS was translated vertically at 15 cm/s and horizontally at 30 cm/s while the gimbal panned back and forth looking 10 degrees off of nadir. This test was the first time that the LVS was successfully operated with changing position and attitude and was a huge step toward the helicopter field test because it verified that the LVS calibrations, alignments and data timing were all correct.

Following the successful gantry tests the LVS was integrated into a Eurocopter AS350 AStar helicopter (Figure 4 and Figure 5). The sensor head was mounted to the nose of the helicopter and the electronics were installed in the cockpit. The operator sat in the back seat behind the pilot with access to the LVS power switches, command GUI and camera feeds. Initial shakeout flights were conducted over terrain west of Fox Field in Lancaster, CA. The first time TRN was run over the Fox Field map it operated successfully generating 3 coarse landmark matches in the each of the first three images and 10s of fine landmark matches in the subsequent images. This was the first time TRN was operated at Mars like altitudes in an uncontrolled environment. Subsequent runs within the same flight and during the

following day showed great promise for TRN but they also pointed out some issues that needed to be fixed. The biggest issue was that with the 60 degree field of view camera, TRN had trouble matching at the lower altitudes and when it did match there were not very many landmarks. The eventual solution was to increase the camera field of view to 90 degrees. Other problems with sensor data drop outs, misalignments and map cropping were also identified and fixed as a result of the shake out flights. After three weeks of rework, the LVS was reintegrated with the helicopter and the field testing began in earnest.

Figure 5. LVS sensors mounted on gantry for initial ground testing (left) and the LVS field test helicopter for flying over Mars like terrain at expected operational altitudes (right).

IV. Coverage of the Mars EDL Operational Envelope

The primary purpose of the field test was to prove that TRN works well across the range of expected Mars EDL conditions. These conditions are dictated by the EDL system design for vehicle position and attitude dynamics and by the proposed landing sites for terrain relief and appearance. Landing time of day constrains the illumination conditions while EDL navigation performance dictates the expected position uncertainties that drive map size and TRN timeline. A detailed description of how the Mars conditions were emulated in the field test is given below.

A. Vehicle Dynamics

1. Altitude

Although TRN can operate at any altitude, for the Mars 2020 application, TRN would operate during parachute descent and complete just before the start of powered descent. This approach achieves the best possible position knowledge error (which decreases during descent) while maximizing the hazard avoidance divert by not cutting into the powered descent fuel or timeline. On MSL the navigation filter was required to converge 10 seconds before powered descent initiation to leave enough time to prime the landing engines. In order to not impact the EDL timeline, TRN will also begin and complete in this 10 second window which corresponds to an altitude range of 3km to 2km where coarse matching occurs at higher altitude and fine matching occurs lower. The Eurocopter can easily cover these altitude ranges

2. Velocity

MSL EDL simulations show that the maximum horizontal velocity is 60 m/s for the extreme case where the lander had residual entry velocity and a strong tail wind. The minimum horizontal velocity is 0 m/s while vertical velocity can range from 80 m/s to 105 m/s. The Eurocopter can cover the horizontal velocity range when flying with a strong tail wind, but it can only descend and ascend at a few m/s. The effect of vertical motion on TRN will be determined from off line analysis or a future field test that can drop the LVS at high velocity.

3. Attitude

Vehicle attitude changes in a complicated fashion during parachute descent -- there is typically a off nadir tilt and a low frequency roll about nadir superimposed with smaller coning motions. Instead of attempting to replicate the motion of specific EDL trajectories with the pan/tilt gimbal, a simplified model was developed for the attitude that was tailored to the gimbal kinematics and performance.

On the left in Figure 6 is a plot of the camera boresight (the vehicle z-axis) in the ground frame. The black line is the trajectory of the boresight while the blue line corresponds to the edge of the image and gives an indication of roll about the boresight. The gimbal cannot replicate the roll about boresight, so in the simple attitude model that angle is thrown away. This is acceptable because the roll is slow and therefore will have little effect on landmark matching. What remains is the azimuth and elevation of the boresight which are shown in the center of Figure 6 (red and blue lines) for a nominal EDL trajectory. These plots show that the angles have a single dominate sinusoid superimposed with a linear ramp, and that is how they are modeled for the field test. For model fitting, first a linear function is fit to the actual EDL angles and then removed. This is followed by detecting of the dominate amplitude, frequency and phase using a Fourier Transform of the sinusoid. The modeled boresight angles are shown in the center of Figure 6 (green and black lines) and the boresight motion that results from these angles is shown on the right in Figure 6. Comparison of the left an right plots shows good qualitative agreement between the actual and modeled boresight.

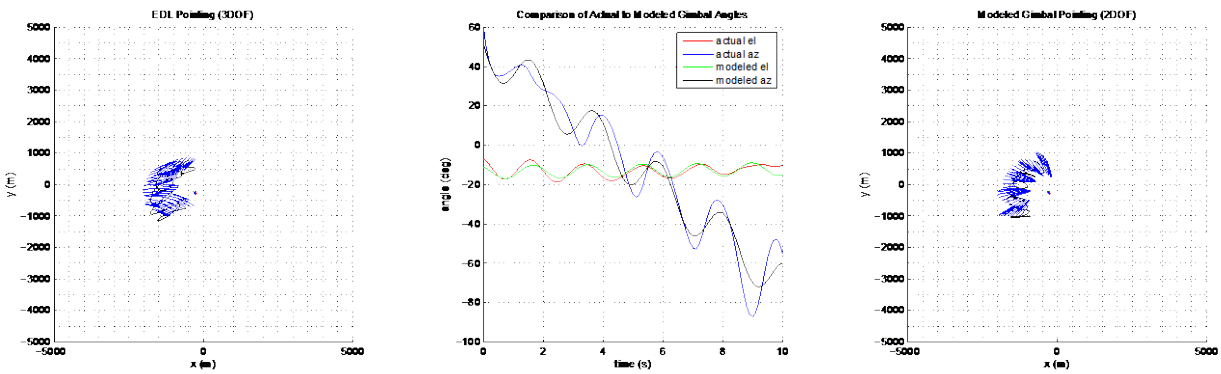

Figure 6. Attitude oscillations during parachute descent. Actual boresight oscillations (left), boresight converted to azimuth and elevation and modeled with a simple function (center), and modeled boresight oscillations.

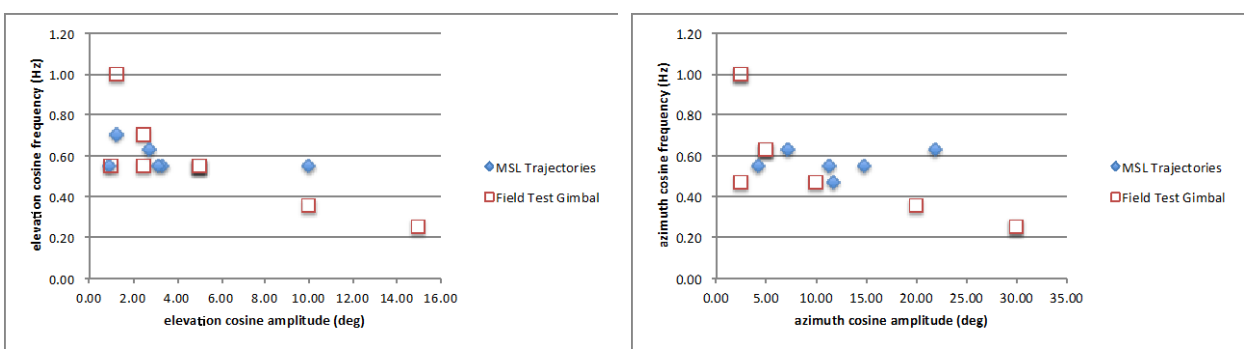

Figure 7. Gimbal profile comparison to EDL attitude trajectories.

Six EDL trajectories that spanned the attitude space were identified and an attitude model fit to each. The sinusoid frequency and amplitude for elevation and azimuth angles are plotted in Figure 7. Given limits on gimbal attitude rates and accelerations, it was not possible to replicate all of the EDL profiles exactly. Instead gimbal profiles were generated that spanned the attitude space while not violating the gimbal constraints. As shown in Figure 7 these profiles cover the attitude space fairly well except for

simultaneous large amplitude and high frequency motion in azimuth and elevation. Although the attitude space coverage is not complete (including the lack of roll about boresight) it was the maximum coverage possible given the available equipment. In the future a high fidelity simulation will be used to obtain complete coverage.

B. Mars Environment

1. Terrain Appearance

The accuracy and number of landmark matches is dependent on the contrast and uniqueness of the map and descent images of the landing site. These properties are difficult to quantify, but it turns out image entropy is a good proxy. Image entropy measures the information content in an image and is defined as $E = -p*log(p)$ where p is the probability of a brightness value in the image. Entropy is high for images with lots of variation in brightness and low for those with low variation. Figure 8 shows the results of a simulation that performed landmark matching on images from Mars and Earth. Since correlation coefficient indicates the quality of a landmark match, the plots show that landmark matching improves as entropy increases. This confirms that entropy is a good indicator for TRN performance.

Figure 9 contains images of potential Mars 2020 landing sites and Figure 8 shows that these sites have entropy that vary from 4 to 7. The terrestrial sites, also shown in Figure 9, were selected to span a greater range of entropies; Cadiz Desert is very bland with an entropy around 3 while Cadiz Rough has entropies as large as 7.5.

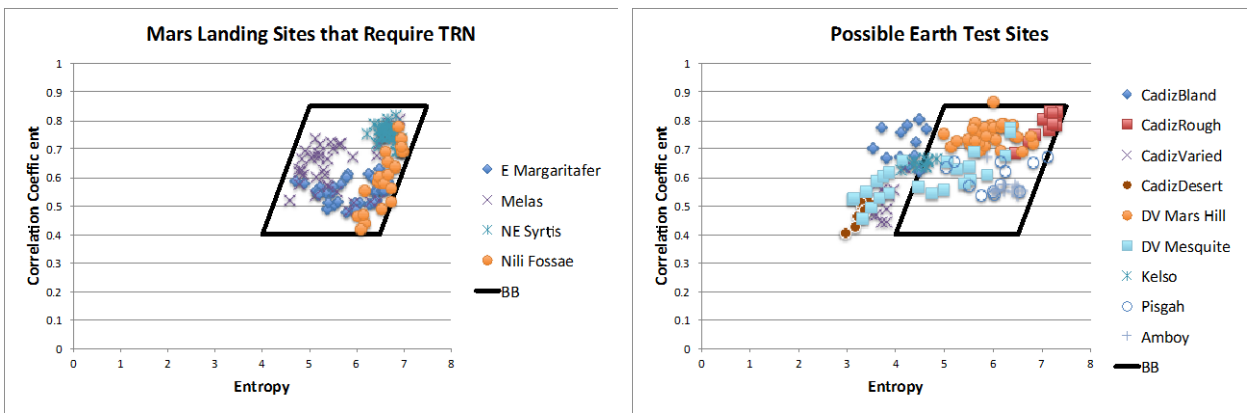

Figure 8. Comparison of image entropy between Mars sites that may require TRN and potential test sites.

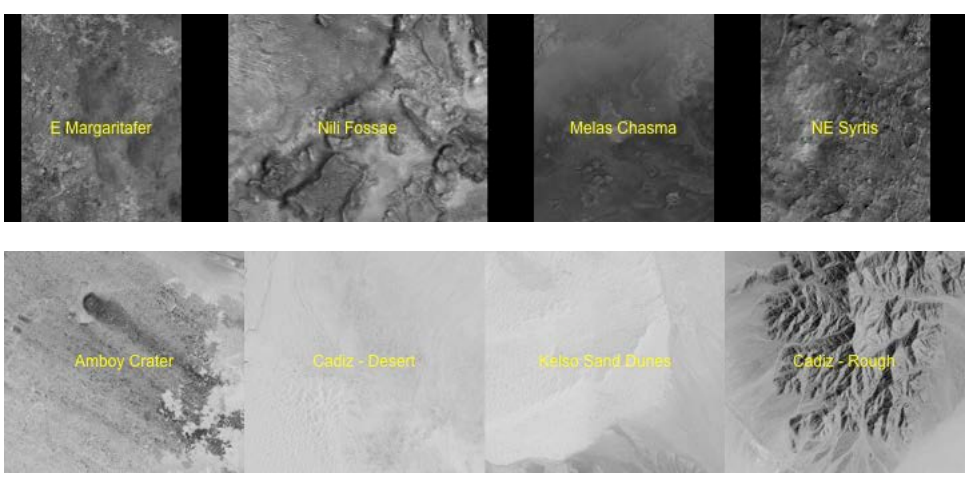

Figure 9. Imagery of landing sites on Mars that may require TRN (top) and Earth test sites (bottom).

2. Terrain Relief

Landmark matching occurs in the ortho-rectified map space. If there is significant terrain relief it could effect landmark matching performance by distorting the image patches that are correlated. There are two different metrics for terrain relief: the slope of the principle plane and the standard deviation of the elevation from this plane. For TRN the relevant area for these metrics is the image patch that is correlated and not the whole image. For the current TRN parameter settings, the coarse landmark patch is around 1.5 km on a side.

Digital Elevation Maps (DEMs) were obtained for three potential M2020 landing sites and three potential earth test sites. The slope and elevation variation terrain relief metrics were computed across each DEM, and histograms of the terrain metrics were generated. Table 1 compares the maximum slope and elevation variation metrics and shows that the test sites bound the terrain relief of the Mars sites.

Table 1. Comparison of terrain relief metrics between Mars landing sites (red) and earth test sites (green). Terrestrial sites have more terrain relief than the Mars sites.

Site	DEM Source	Max Elevation Variation (m)	Max Slope (degrees)
North East Syrtis	CTX	43	10.2
Nili Fossae	CTX	66	7.3
Holden	CTX	36	4.1
Kelso Sand Dunes	SRTM	38	4.3
Cadiz Area	SRTM	115	12.4
Mars Hill	SRTM	120	13

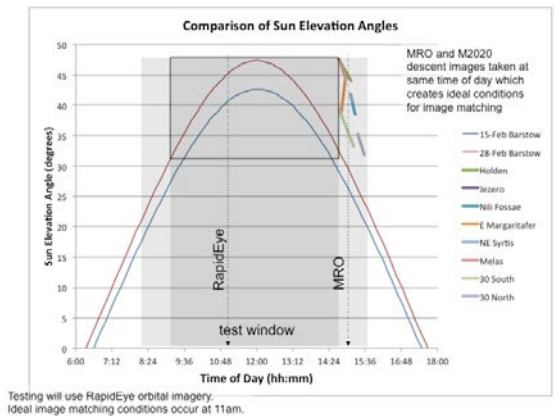

Figure 10. Comparison of illuminations conditions between M2020 landing sites and Earth test sites.

3. Illumination

The simplest implementation of landmark matching, and the one used in the LVS, is to directly correlate the descent image and an orbital image. A potential issue with this approach is the difference in illumination and consequently shading and shadows between the two images. To assess the impact of illumination differences the field test was designed to cover a broad range of sun elevation angles. Given a sun elevation, sun azimuth will vary with the day of year. Since the field test dates were fixed, sun azimuth was not controllable.

Figure 10 shows a comparison of sun elevation angles at different times of day. The orbit of Mars Reconnaissance Orbiter (MRO) is designed to image at 3pm local time. Because MRO is also used for telecom during landing the landing time of day will also be 3pm. Therefore the orbital images taken by MRO for the TRN map and the descent images taken during landing will have very similar illumination conditions. This is fortuitous for LVS since it will make landmark matching that much easier. Because of the variations in landing days and latitudes the sun elevation angles for the M2020 landing sites will vary between 30 and 50 degrees. These sun elevation angles are achieved at the field test sites between the hours of 9pm and 3pm during the field test campaign, so that is when the flights were conducted. The field test maps were generate from RapidEye images which images at 11am, so this is the time during the field test when the map and LVS images have the most similar illumination.

C. TRN Inputs

1. Sensors

Commercial off the shelf sensors were used in the LVS prototype. The camera was a commercial PhotonFocus DS1-D1024 CMOS visible imager with global shutter. The lens was a Kowa 6M6HC 6mm

lens which when paired with the camera obtained a 85 degree FOV. The camera exposure times were on order 1ms and the images were collected on demand at a rate of about 1 Hz with an image transfer time of around 30ms. All of these parameters are achievable with a flight camera designed for LVS. The IMU used in the field test was a commercial Northrop Grumman LN-200 fiber optic gyro with an angle random walk of 0.15˚/√Hr. The flight IMU for LVS will have this performance or better.

2. Initial State Estimate

During the field test the initial state came from the ground truth system with the majority of the error coming from the Applanix PosAV 310 INS/GPS which had position errors less than 6 m, velocity errors less than 0.05 m/s and attitude errors less than 0.1˚. The INS/GPS performance is much better than the MSL EDL performance, so additional errors were added to the LVS initialization. Based on Monte Carlo analysis the errors at LVS initialization were +/- 3km in horizontal position, +/- 120m in altitude, +/- 0.9 m/s in all three velocity directions and +/- 0.15˚ in all three attitude angles.

3. Maps

The TRN maps must cover the entire landing ellipse and the terrain outside the ellipse that could be viewed by the camera. For M2020 this can result in maps as big as 30x30km, however the entire large map does not need to be searched by TRN. Instead this large map is cropped around the on-board position estimate so that it covers the position errors (+/- 3km) and the image field of view (another +/- 3km). This 12km x 12km map is the one used in the coarse matching phase of TRN. Since the FPGA logic has fixed the TRN maps at 1024x1024 pixels, the coarse image pixels are roughly 12m/pixel. For fine matching, the coarse position estimate is used to crop the map to 6km x 6km at 6m/pixel. This increases the numbers of landmarks matched and improves accuracy.

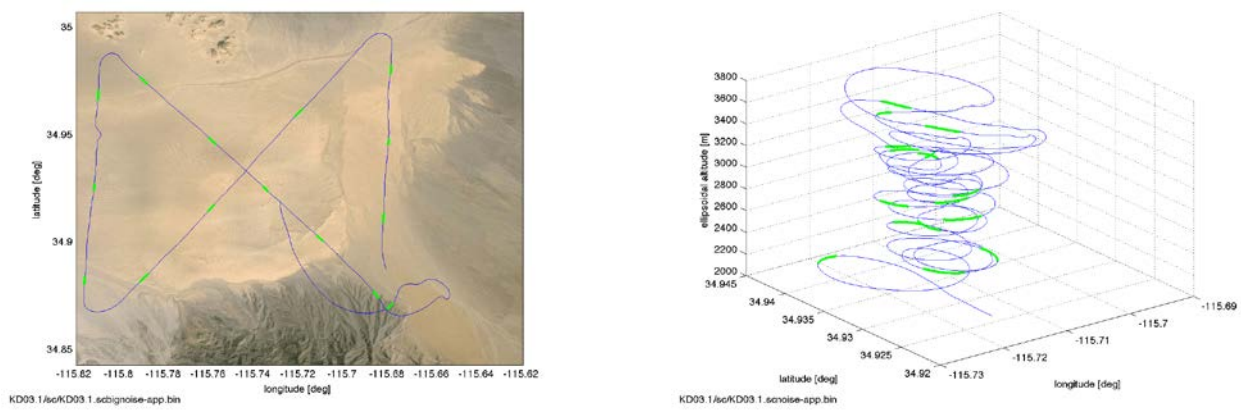

Figure 11. Example flight profiles with green segments indicating time when TRN was operated. A figure 8 over Kelso Sand Dunes (left) and a spiral over Kelso Sand Dunes (right).

Given these map resolutions and sizes for M2020, the TRN maps can be made from MRO CTX imagery. Two overlapping CTX stereo images can be processed to generate a DEM and one of the images can be used as the appearance map. The native resolution of CTX is about 5m/pixel and the stereo DEMs have postings every 25 m. For the field test, the DEM for the TRN map came from the Shuttle Radar Topography Mission (SRTM) data set and had 30m postings, which is slightly worse than the CTX DEMs. The imagery came from a RapidEye, a commercial provider of orbital imagery. The RapidEye images were 5 m/pixel and were generated with a push-broom camera, just like CTX.

V. Field Test Results

During the test, the helicopter flew between 2km and 3km altitude over six different tests sites in Death Valley National Park and the Mojave Desert. These sites were selected to span the full range of terrain appearances and relief expected for landing sites being considered for the M2020 lander mission. The tests were conducted between 9am and 4pm to obtain a full range of illumination conditions. The

gimbal was operated in modes that simulated lander attitude profiles during parachute descent including profiles with large off nadir angles and high angular rates. TRN was run on-board by the LVS with no operator intervention and no tuning of parameters happened between or during fight.

Each flight day had on average 2 flights. During each flight 30-50 TRN runs were conducted. Due to data storage bandwidth limitations from the LVS compute element, TRN could run about once every two minutes. Figure 11 shows two typical flight profiles with the green segments indicating when TRN was operating in flight. As can be seen from the images a variety of terrains (rough, bland, flat) and altitudes were covered by the flights.

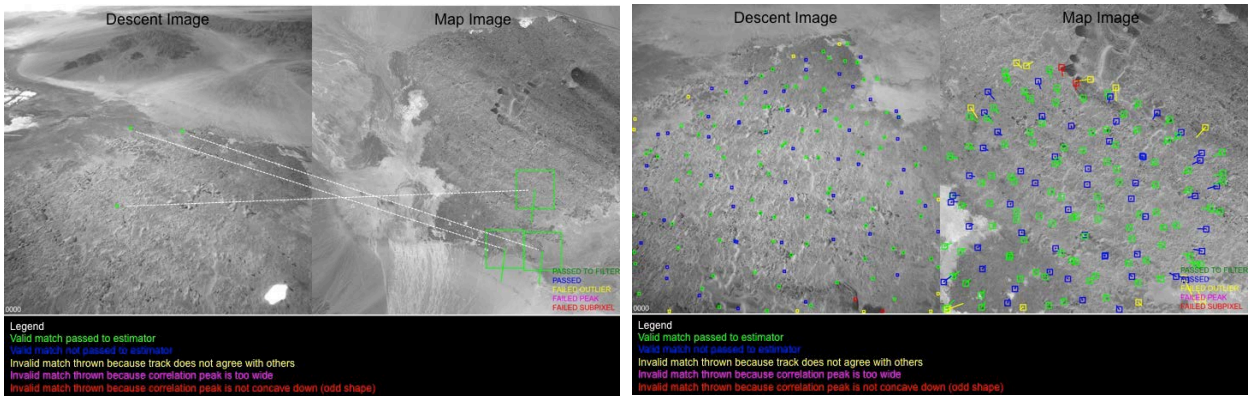

Figure 12. Example landmark matching has coarse (left) and fine (right) phases.

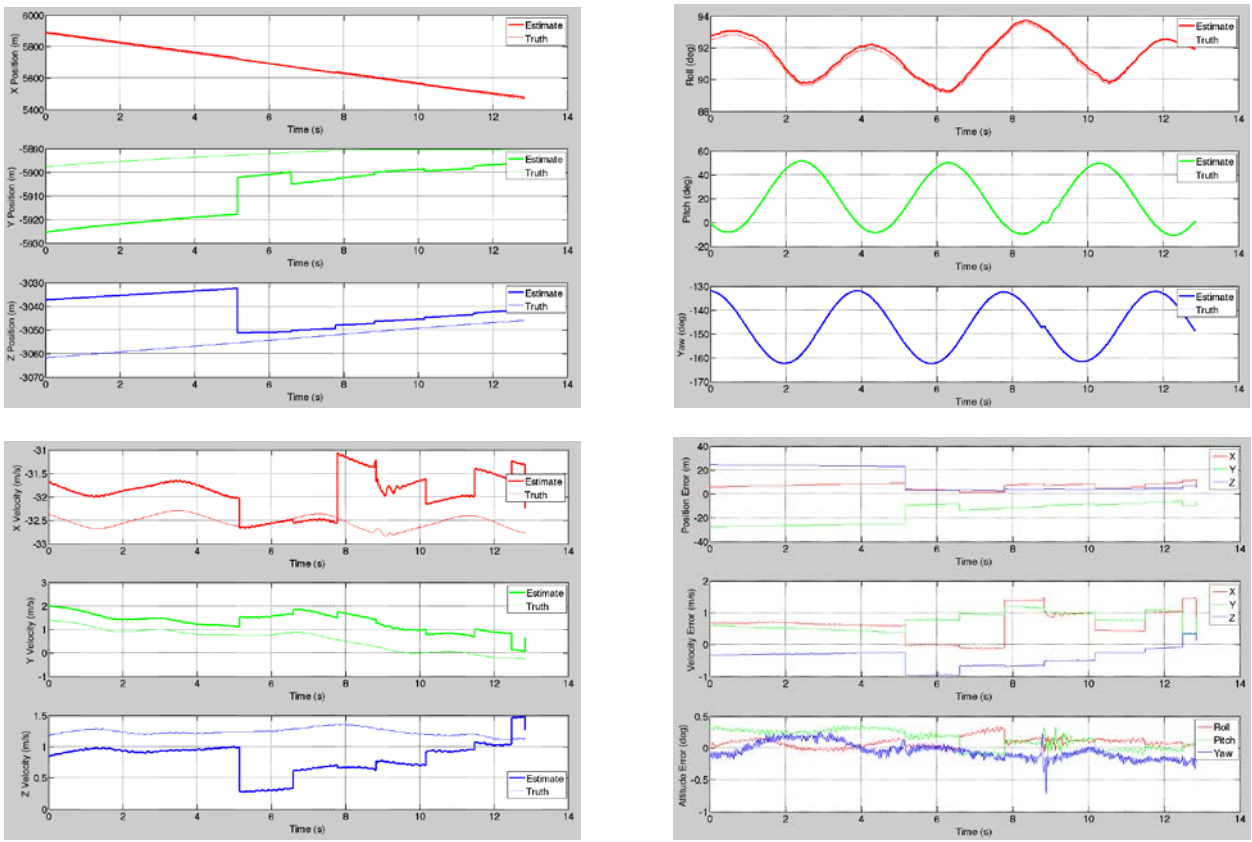

Figure 13. LVS state estimates compared to ground truth: position (top left), attitude (top right), velocity (bottom left) and state errors (bottom right).

For each TRN run, LVS processed 3 coarse and 7 fine images. Example coarse and fine landmark matches for a run over Amboy lava flow are shown in Figure 12. The coarse matching uses a 12km x

12km map image at 12 m/pixel. The precision matching switches to a 6km x 6km map at 6m/pixel; many landmarks are matched but only 100 are sent to the estimator to limit the processing time of the EKF update. Figure 13 shows the position, attitude and velocity state estimates and errors compared to ground truth. The first large jump in the estimates happens after the 3x3 three coarse landmarks are processed in the batch initialization estimator; the smaller jumps that follow occur at EKF landmark updates. The state estimates show that for this run the helicopter was flying north (X) at about 32 m/s and the gimbal was oscillating dramatically up to 30˚/s. In fact, because the IMU was offset from the gimbal center of rotation, some of the gimbal motion shows up in the velocity. The error plots show that initially large errors are taken out and the EKF converges over time to accurate estimates. For this run at 3km altitude, the position errors are less than 20m, the velocity errors are less than 1 m/s (0.3% of altitude) and the attitude errors are less than 0.5˚.

Some of the TRN runs were intentionally performed outside of the operational envelope. Figure 14 shows a case were the terrain has an entropy around 3 which is more bland than any potential Mars landing sites that require TRN. Landmark matching still succeeds although there are fewer fine matches than is desired. Figure 14 also shows a result with a landmark matching performing very well even in the presence of over 300m of elevation change viewed 45˚ off from nadir.

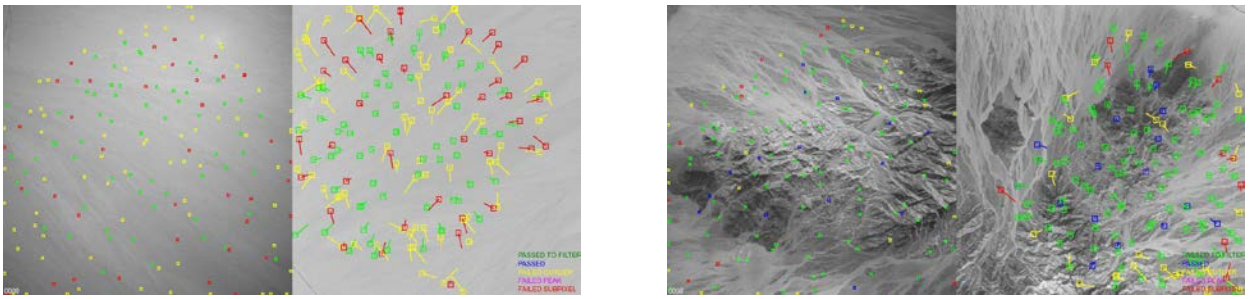

Figure 14. Examples of extreme conditions where TRN succeeded. Very bland terrain with the zero phase spot in the image (left) and 300m of elevation change viewed 45˚ off nadir (right).

During the field test there were a series of flights with no changes to the LVS software and hardware. The flights were over the Cadiz, Amboy Pisgah and Kelso maps. During these flights 128 of the TRN runs were within the operational envelope and 301 were not in the envelope for a total of 429 runs. Scatter plot of the horizontal errors after coarse matching and after the final fine match for runs inside the operational envelope are shown in Figure 15. The coarse matching reduces the initial errors that are up to +/- 3km down to +/- 200m. However, it took the fine matching phase with numerous landmark matches and higher resolution maps to bring the errors down within the 40m error bounds. These plots are the as flown performance of the LVS without modification. Although the 40m requirement is met the fine matching performance does have a bias of approximately 10m in the south (-X) and east (+Y) directions. The bias is currently under investigation.

The number of landmark matches is a very powerful indicator of the likelihood that the position estimate is correct. Some of the TRN runs matched more than 700 landmarks total across the 10 images and in all these cases the position errors met the 40m requirement. However, 7 of the 128 TRN runs in the operational envelope matched a small number of landmarks and had poor position estimation performance. Figure 16 shows the total number of landmarks matched against the total horizontal position errors. The 7 outliers are at the bottom of the plot. These “invalid” estimates can easily be eliminated by throwing out estimates that have to few landmark matches. In the plot, a threshold of 200 landmark matches easily separates the matches that meet the 40m requirement from those that don’t. The total number of landmarks metric clearly separates the space of correct and incorrect estimates and can be used to prevent the LVS from providing a spurious answer.

Figure 15. The errors in horizontal position after coarse match (left) and fine match (right) compared to the 40m error requirement.

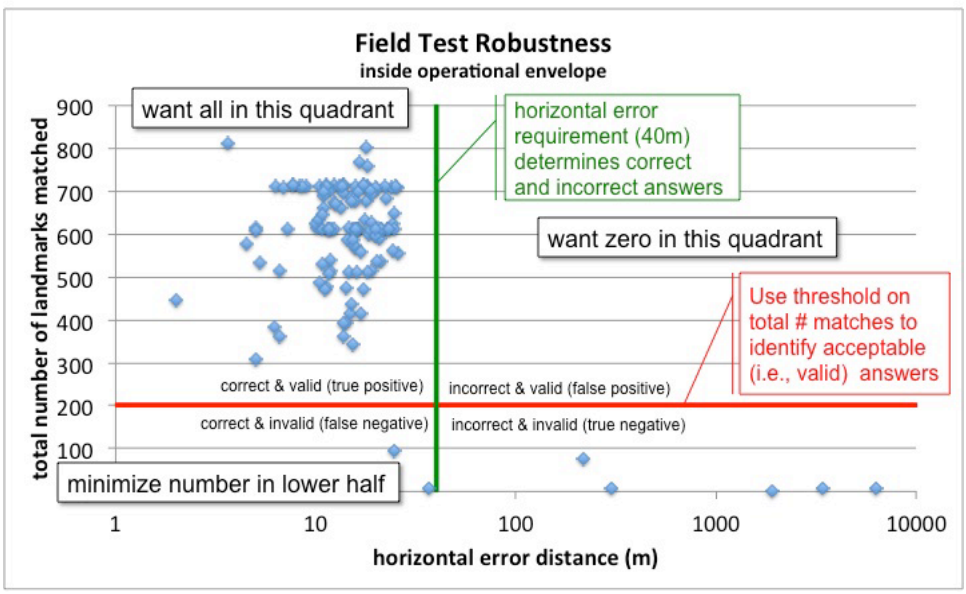

Figure 16. The total number of landmarks easily separates correct answers from incorrect answers.

Although 128 is not very many runs to cover the large operational test space, care was taken in the planning of field test flights to maximize coverage and hit the extremes of the envelope. Table 2 compares the M2020 operational environmental bounds (as defined in Section IV) to the maximum and minimum values obtained in the field test runs. Image entropy, elevation variation, slope, sun elevation, off nadir angle and initial state errors from the field test all exceeded the M2020 envelope (albeit not all at once). The minimum horizontal speed was not achieved because the helicopter required forward motion to operate at high altitude, however low horizontal speed is not expected to have an adverse effect on TRN performance. The vertical speed could not be met with the helicopter and as discussed future analysis and testing will cover this part of the space. Finally the maximum boresight angular rate of 60˚/s was not quite achieved and this parameter will be investigated further in simulation and will be covered in a future test.

Table 2. Comparison of the Mars 2020 operational envelope to the field test coverage.

Metric	Mars 2020 Min	Mars 2020 Max	Test Min	Test Max
Image Entropy	4.0	7.5	3.0	7.5
Elevation Variation	0 m	50 m	0 m	70 m
Slope	0˚	16˚	0˚	37˚
Sun Elevation	30˚	50˚	18˚	50˚
Horizontal Speed	0 m/s	60 m/s	15 m/s	70 m/s
Vertical Speed	80 m/s	105 m/s	0 m/s	5 m/s
Altitude Range of TRN	2300m	3000m	1400m	3400m
Boresight Off Nadir Angle	0˚	45˚	0˚	70˚
Boresight Angular Rate	0˚/s	60˚/s	0˚/s	45˚/s
Initial State Errors	+/- 3km horz, +/- 60m vert		+/-3km horz, +/- 120m vert	

VI. Conclusion

The Mars LVS helicopter field test showed that the LVS met the 40m accuracy in 10 seconds TRN performance requirements across the entire Mars 2020 operational envelope. TRN worked well with large terrain relief, large differences between descent image and map illumination, bland terrain and significant attitude dynamics. The results also showed that the LVS has a very strong metric for distinguishing incorrect position estimates from correct ones that will prevent LVS from reporting spurious measurements. Current work is focused on eliminating the small bias in the position estimates and explaining the 7 invalid TRN runs in the operational envelope. Future work will investigate the performance of LVS with EDL like vertical motion and very large attitude rates.

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Mars Program in the NASA Science Mission Directorate. The field test was implemented by the LVS Field Test Team including: Charles Bergh, James Bodie, Charles Fisher, Kim Gostelow, Dean Holt, John Koehler, Garen Khanoyan, Ari Mahpour, James Roberts, Aaron Stehura, Phoebe Sulzen, Geoffrey Vaughan and Gregory Villar. The helicopter was provided by Wolfe Air and operated by John and Mike Tamburro.

References

[1] J. Alexander, Y. Cheng, W. Zheng, N. Trawny and A. Johnson, "A Terrain Relative Navigation Sensor Enabled by Multi-Core Processing," *Proc. IEEE Aerospace Conference (AEROCONF 2012),* March 2012.

[2] Y. Cheng, D. Clouse, A. Johnson, W. Owen and A. Vaughan, "Evaluation and Improvement of Passive Optical Terrain Relative Navigation Algorithms for Pin-Point Landing," *Proc. AAS Space Flight Mechanics Meeting*, January 2011.

[3] C. Epp and T. Smith, "The Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," *Proc. of the Aerospace Conference*, Big Sky, MT, March 2008, pp. 1-7.

[4] A. Johnson, R. Willson, Y. Cheng, J. Goguen, C. Leger, M. SanMartin and L. Matthies, "Design Through Operation of an Image-Based Velocity Estimation System for Mars Landing," *International Journal of Computer Vision*, 74(3), pp. 319-41, January 2007.

[5] A. Johnson, A. Ansar, L. Matthies, N. Trawny, A.I. Mourikis, and S.I. Roumeliotis, "A General Approach to Terrain Relative Navigation for Planetary Landing", *Proc. 2007 AIAA Infotech at Aerospace Conference*, Rohnert Park, CA, May 7-10, 2007.

[6] A. Johnson and M. Golombek, Lander Vision System for Safe and Precise Entry Descent and Landing, *Concepts and Approaches for Mars Exploration,* June 12–14, 2012, Houston, Texas. LPI Contribution No. 1679.

[7] Johnson, A., Bergh, C., Cheng, Y., et al, "Design and Ground Test Results for the Lander Vision System," *36th Annual AAS Guidance and Control Conference*, Breckenridge, CO, 2013, AAS 13-042.

[8] Mourikis, A., Trawny, N., Roumeliotis, S., Johnson, A., Ansar A., and Matthies, L., "Vision Aided Inertial Navigation for Spacecraft Entry Descent and Landing", *IEEE Transactions on Robotics and Automation* Vol. 25, No. 2, 2009, pp. 264-280.

[9] Mustard, J.F., et al., *Report of the Mars 2020 Science Definition Team* by the Mars Exploration Program Analysis Group (MEPAG), URL: http://mepag.jpl.nasa.gov/reports/MEP/Mars_2020_SDT_Report_Final.pdf. [cited 23 May 2014].

[10] Wolf, A., Acikmese, B., Cheng, Y. Casilova, J., Carson, J., Ivanov, M., "Toward improved landing precision on Mars," IEEE Aerospace Conference, Big Sky MT, 2011.