HIGH-PRECISION TERRAIN RELATIVE NAVIGATION FOR TERMINAL DESCENT AND LANDING

Ike R. Witte,[1*] Lindsey A. Marinello,[1] and Steve N. Jenkins[1]; [1]Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723 USA [*][Isaac.Witte@jhuapl.edu]

Abstract. *Terrain Relative Navigation (TRN) enables autonomous precision landings, but typical algorithms either stop or switch to a visual odometry mode for terminal descent and landing due to the coarseness of the reference maps. Both options lead to a drift in spacecraft state knowledge, necessitating larger safe-landing ellipses. In this paper, we present a landing-site-anchored optical navigation algorithm that provides high-precision TRN with little drift and without the need for higher-resolution reference maps. We compare our algorithm with current methods, and demonstrate algorithm performance both qualitatively on available landing videos as well as quantitatively through an array of lunar landing simulations.*

Introduction. Increasingly, robotic exploration of small and planetary bodies requires accurate and autonomous navigation for increased scientific gain of the mission. This navigation must be done real-time and done autonomously on the spacecraft. One approach to this problem is to employ TRN, which uses onboard camera imagery to solve for the spacecraft's position and/or attitude. TRN can be supplemented with additional sensors such as star trackers and LIDARs, and typically rely on reference maps to navigate relative to.

When the reference maps become too coarse relative to the camera GSD, one typical approach switches to relative odometry TRN measurements. Rather than comparing the camera image to a reference map, relative TRN algorithms compare the camera image to another recently taken image and estimate the change in pose. Reference maps are not relied upon during this phase of TRN so relative measurements can be achieved at a lower altitude than absolute measurements, but because the measurements are more of a velocity measurement and are not anchored to the surface, the position uncertainty will grow over time similar to the navigation drift during IMU-only navigation, leading to larger landing error ellipses.

We look to improve the performance of TRN during this terminal descent and landing phase of the entry, descent, and landing (EDL) trajectory by performing persistent feature tracking and consensus-based matching in order to accurately track and navigate relative to the desired landing site. We call this algorithm the Landing-site Anchored Navigation for Descent, or LAND. This algorithm could complement existing higher-altitude algorithms by providing an opportunity to continue performing accurate TRN measurements down to the surface. LAND is able to robustly track a landing site and surrounding features in order to estimate lateral offset, scale, and rotation about the camera boresight, easing the resolution requirement on reference maps and allowing for more precise landings.

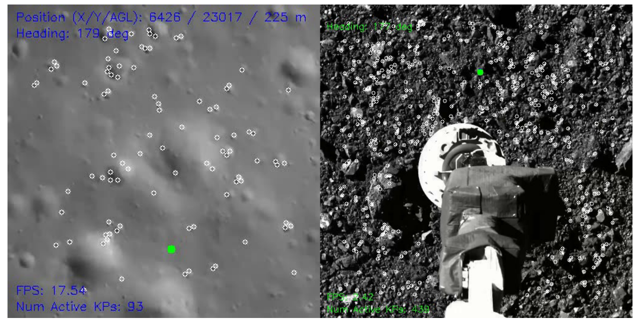

Figure 1. LAND tracking the landing site (green) during the Chang'e 4 lunar landing (left) and the OSIRIS REx touch-and-go (TAG) maneuver to sample Bennu (right). Tracked and matched features shown as white circles.

Methedology. The LAND algorithm is built on top of the Consensus-based Matching and Tracking (CMT) object tracking algorithm designed by Nebehay and Pfugfelder,[1] modified in several key ways to enable the tracking and interpretation of the landing site point relative to the terrain rather than a bounding box within the frame. The CMT algorithm tracks an object across frames by separating features into the background or the desired object, and matching the features from the first frame to those of subsequent frames. ORB feature detection and matching is used for both detector and descriptor in the LAND algorithm, followed a brute-force matcher.[2] To add further robustness to the feature tracking, the CMT algorithm also computes the displacement of each keypoint by employing the pyramid variant of the Lucas and Kanade optical flow algorithm.[3] Matched keypoints supersede tracked keypoints, as they do not rely on recursive estimation and are therefore generally more robust. By matching back to the original frame, the CMT tracker keeps the tracker from drifting even across many frames. As long as the features are able to be matched, there is very little error accumulation in the localization of the object in the frame. This is an advantage for precision landing: the terrain can be treated as rigid so as the persistent features are matched to the initial frame, and the error does not accumulate like it would for sequential image-to-image algorithms.

In order to estimate the state, each tracked keypoint casts a vote $h(a, m)$ for the object center, where a is the keypoint position in absolute image coordinates and m is the index of the corresponding keypoint in the original frame O. The general form of the vote takes the form:

$$h(a, m) = a - sRr_m$$

where s is the estimated scale factor, r_m is the relative position of the corresponding keypoint in O, and R is the 2D rotation matrix around the boresight. s is computed with ehte pairwise Euclidean distance between the current feature a_i and every other feature a_j and comparing the distances to those of the corresponding keypoint r_{mi} and r_{mj} in O. If $a^{i,j} = a_i - a_j$ and $r^{i,j} = r_{mi} - r_{mj}$ then s is estimated from the median of the distribution in scale changes:

$$D_s = \left\{ \frac{\left\|a^{i,j}\right\|}{\left\|r^{i,j}\right\|}, i \neq j \right\}$$

Similarly, the rotation estimate α is estimated from the median of the distribution $D_\alpha = \left\{\alpha_{i,j}, i \neq j\right\}$, where

$$\alpha_{i,j} = atan2(a_y^{i,j}, a_x^{i,j}) - atan2(r_y^{i,j}, r_x^{i,j})$$

CMT has been modified for LAND to search for features over the entire image rather than a local neighborhood in order to track the terrain, include an exclusion zone mask to ignore parts of the spacecraft in the field of view (FOV), and to replenish features as features are no longer tracked due to traversal or large scale changes. Additionally, prior to the voting and state estimation the feature locations are warped into a nadir frame using homographic warping H, which is valid even with terrain relief due to the pure rotational warping.

$$H = K R_{curr} (I_{3x3} - \vec{T}_{12} v^T) R_{des} K^{-1}$$

where K is the camera intrinsic matrix, R_{curr} is the rotation matrix of the spacecraft in a NED frame, $\vec{T}_{12}$ is the translation vector from the current to the desired frame and in this case is 0, v^T is the normal vector of the terrain and is ignored due to the lack of translation, and R_{des} is the nadir rotation matrix of the spacecraft in a NED frame. Given an initial range estimate prior to LAND starting (frame 0), the altitude of the spacecraft for frame k is estimated from the scale using $\hat{z}_k = \hat{z}_0/s$. Finally, the x and y position estimate of the spacecraft is computed with

$$\begin{bmatrix} \hat{x}_k \\ \hat{y}_k \end{bmatrix} = \vec{L} + GSD * \begin{bmatrix} -cos(\alpha) & sin(\alpha) \\ sin(\alpha) & cos(\alpha) \end{bmatrix} \begin{bmatrix} \delta_i \\ \delta_j \end{bmatrix}$$

where $\vec{(L)}$ is the x and y location of the landing site, GSD is the ground sample distance of the current image which can be estimated from the altitude estimate, and δ_i and δ_j are the row and column pixel offsets between the camera boresight and the tracked landing site.

Initial Results. In addition to the OSIRIS-REx and Chang'e 4 mission data, LAND has also been tested extensively in simulation in order to determine estiamtion capability in an environment where absolute truth is known. Simulated data included relevant descent trajectories, as well as stressing cases where the landing site temporarily leaves the FOV. While the full paper will include array of tests and results performed in this analysis, for the sake of brevity in this abstract a single test case is shown that is from a representative descent trajectory. Figure 2 shows the North, East, and Down position estimation error as well as the normed 3D estimation error vs time and total 3D distance traveled in the trajectory. LAND performs well, and is able to retain track of the landing site throughout, ending with a 3D position error of just over 1m at the end of the trajectory.

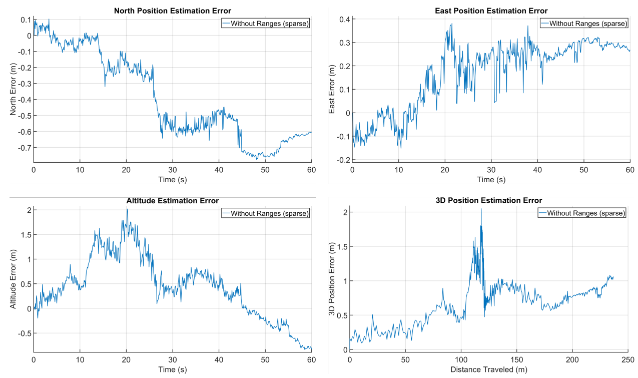

Figure 2. Results using simulated imagery from a terminal descent test trajectory. LAND is able to keep the position estimate error to 1-2 meters.

Conclusion. In this paper we propose a high-precision landing-site-tracking TRN algorithm to be used during terminal planetary descents and landings. Built around a robust CMT object tracker, the LAND algorithm enables precision landings on the order of 1 meter without the need for extremely high resolution maps. By employing persistent feature matching and tracking, LAND accumulates very little error during descent even though it is not a map-based algorithm, resulting in landing accuracies similar to position accuracy of the spacecraft's final map-based measurement. We have demonstrated robust landing-site tracking on historical flight data, and have shown accurate position estimation performance in simulation, outperforming current image-to-imagealgorithms. LAND will continue to be developed and improved, with future endeavors focusing on closed-loop Monte Carlo analysis and implementation on relevant hardware.

References.

[1] G. Nebehay and R. P. Pflugfelder, "Consensus-based matching and tracking of keypoints for object tracking," *IEEE Winter Conference on Applications of Computer Vision*, pp. 862–869, 2014.

[2] M. Muja and D. G. Lowe, "Fast matching of binary features," in *2012 Ninth Conference on Computer and Robot Vision*, pp. 404–410, 2012.

[3] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *International Joint Conference on Artificial Intelligence*, 1981.