# IMAGE PROCESSING FOR VISION BASED NAVIGATION

N. Rowell[1], S.M. Parkes[1], M.N. Dunstan[1], I. Martin[1],
O. Dubois-Matra[2]
[1]University of Dundee, United Kingdom [2]ESA (ESTEC), Netherlands

**ABSTRACT**

*Optical imagery provides a powerful means of autonomous navigation for spacecraft in the vicinity of Solar system bodies, where the communication delay may be much longer than the dynamical timescales. Cameras are now routinely included among the suite of sensors used to guide exploratory spacecraft, with information extracted from the images fused with that of other sensors to estimate the spacecraft state and plan orbital maneuvers. The information content of a typical image is orders of magnitude larger than that supplied by other sensors, and sophisticated processing algorithms are required to extract navigation information and reduce the raw pixel data to a manageable size. It is common to offload this task to a separate image processing system rather than integrate this directly with the GNC system.*

*We have developed an image processing algorithm that identifies and tracks surface features through a sequence of images, in a manner robust to changes in the viewing angle, illumination and scale such as might occur during descent and landing onto a planetary surface or orbit around a small body. It does not rely on the presence of any particular type of morphological feature, and may be used for absolute navigation provided a suitable map of the terrain is available. We will describe the operation of this algorithm, present results based on simulated image data generated by the PANGU tool, and discuss further ways in which imaging data can be used to support navigation.*

## 1. INTRODUCTION

Extracting structure and motion information from monocular image streams is a well-studied area of computer vision. It is widely used on Earth, for example, in fields as diverse as archaeology, medicine, CGI [1] and robot navigation. In the latter case, research has focused on achieving real-time performance of vision-based simultaneous localization and mapping (SLAM) systems, e.g. [2], where the camera motion and environment structure is measured simultaneously from imaging data alone. In a typical application, there are two distinct problems that must be solved in order to extract geometrical information from a set of images. The first is known as the *correspondence problem*, and consists of finding matching points among the images. This can be a major challenge depending on the situation: rapid motion can result in large disparity between frames, which means large sections of the images must be searched for features that may have changed significantly in appearance. The second problem is that of *motion estimation*, which involves using the corresponding points to measure the relative position and orientation of the camera when each image was taken. The calculations involved are extremely ill-conditioned for small motions, which results in an unfavourable trade off between rapid image capture for improved point correspondences and larger disparity for improved conditioning of the motion estimation.

Space missions now routinely include cameras among the set of onboard sensors. These provide accurate surface-relative navigation, hazard detection and topographic mapping for planetary and asteroid rendezvous craft [3][4]. The computational load of typical image processing algorithms is extremely high, and it is common to offload parts of the algorithm to dedicated hardware. In Section 2 we review related work done at the University of Dundee, and outline enhancements to an established feature tracking algorithm designed to improve the performance of higher level navigation functions. In Section 3, we describe efforts to integrate this enhanced algorithm with an absolute navigation system, by identifying suitably defined known landmarks in descent imagery.

## 2. FEATURE EXTRACTION AND TRACKING INTEGRATED CIRCUIT (FEIC)

The University of Dundee has developed the Feature Extraction Integrated Circuit (FEIC) [5] as part of the NPAL project [6] under ESA contract 15618/01/NL/FM. It is designed to extract and track feature points through a monocular image stream. The algorithm uses the Harris detector [7] for feature point selection. Tracking is achieved by extracting a 7×7 pixel region about each feature point, and searching for this in later images using intensity correlation in a restricted search window about either the last known position of the feature or the predicted new location. Many of the window-based operations can be done in a highly parallelized manner when implemented in hardware, which enables much faster processing. The original algorithm runs at 20Hz on 1024×1024 pixel images. It was designed for planetary descent and landing applications, where the approach speed is substantial and the image processing must be done rapidly.

### A. FEIC ALGORITHM ENHANCEMENTS

In vision-based guidance system terms, the FEIC is designed to solve the correspondence problem. We have been developing improvements to the FEIC algorithm intended for applications with longer dynamical timescales, such as NEO rendezvous missions, when more time can be spent processing each image to achieve more robust point correspondences. Note that the improvements are to the algorithm only, and are currently only implemented in software. These improvements are specifically designed to improve the conditioning of the higher level operations. There are two main enhancements to the basic algorithm. The first enhancement allows the 7×7 image sections extracted for each feature to be periodically updated (‘texture updating’), which corrects for gradual changes in the scale and orientation of a feature that occur during tracking and would otherwise cause the point to be lost. The second enhancement is designed to reduce the number of outlying point correspondences – that is, points that wander during tracking and do not correspond to the same surface feature between frames. The main cause of these points is the existence of multiple possible matches within the search region, indicated by several strong correlation peaks. The original algorithm takes the point of strongest correlation to be the new location of the feature (provided it is above a threshold), whereas the enhanced algorithm simply rejects any points for which there are multiple thresholded correlation peaks.

### B. HIGHER LEVEL NAVIGATION FUNCTIONS

In order to demonstrate the improvements to higher level vision-based navigation functions that result from the improved point correspondences obtained by the enhanced FEIC algorithm, we have processed a 100 frame image sequence generated using the PANGU tool [8]. In the test scenario, the camera flies over a simulated planetary surface at an altitude of 200m, and moves roughly 10m along the camera boresight between each frame. The tenth frame in the sequence is shown in Fig. 1. In the left image the original FEIC algorithm has been used to track points, and in the right image the enhanced FEIC algorithm has been used. Only points that have tracked for all ten frames have been plotted for clarity. The enhanced algorithm has tracked more points overall, because the use of texture updating (every 3 frames in this case) has prevented points being lost due to gradual changes in appearance. It also results in fewer outliers (wandering points evident in the lower section of the image), due to the checks made on secondary correlation maxima during tracking – it simply rejects points if there is a risk they have been misidentified.

We use point correspondences ten frames apart (like those depicted in Fig. 1) to measure the epipolar geometry between pairs of images (see [9]), by calculating the Essential Matrix using a linear technique [10] with RANSAC outlier rejection. The Essential Matrix can be used to recover the rotation and translation *direction* of the camera between frames, and requires only the point correspondences and the calibration properties of the

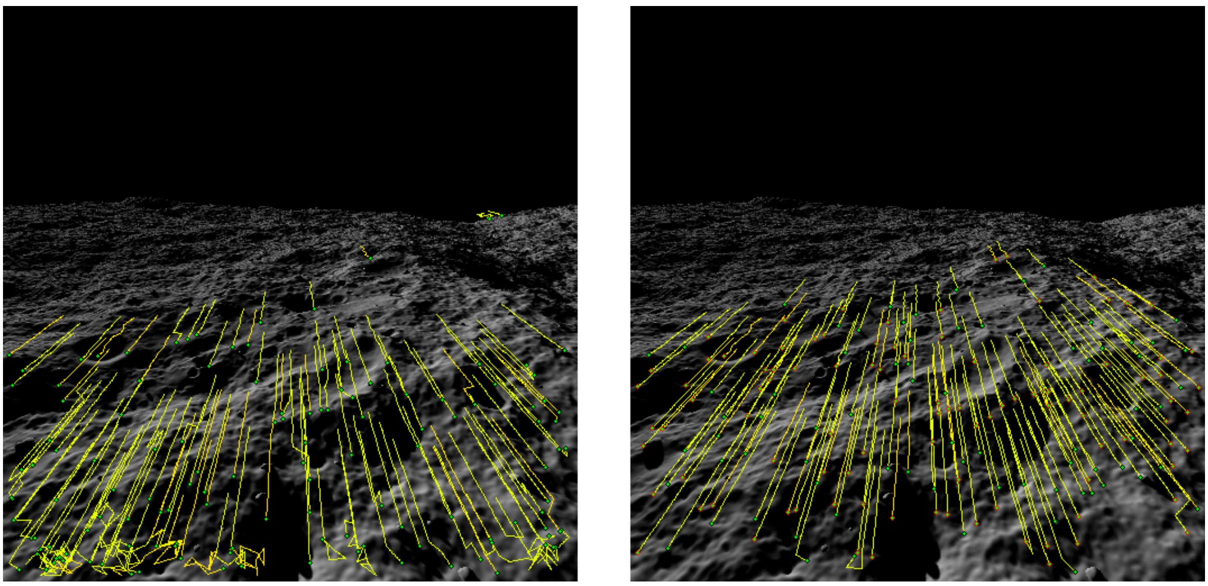

***Fig. 1:*** *Tracking using the original (left) and enhanced (right) FEIC algorithm*

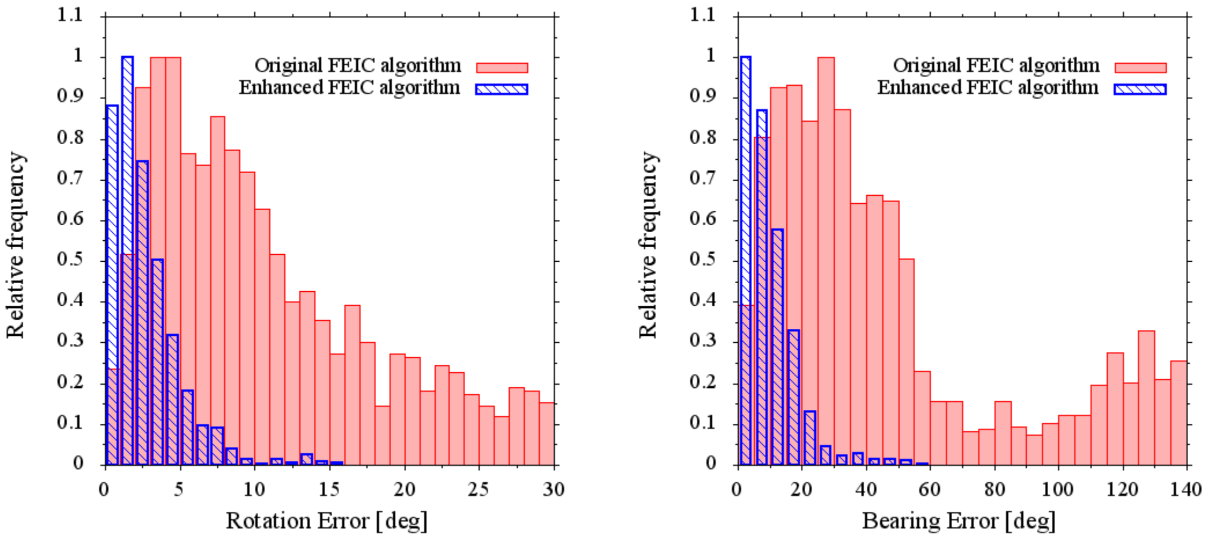


***Fig. 2:*** *Performance of the original and enhanced FEIC algorithms at measuring the epipolar geometry of image pairs*

camera. The magnitude of the translation cannot be recovered due to an overall scale ambiguity – the geometry cannot distinguish between a large object far from the camera and a small object nearby. In order to compare the solutions obtained by each algorithm, we look at the error in the rotation and the direction of the translation (the *bearing*). The results for all frames over the sequence are presented in the histograms in Fig. 2. These have been normalized so that the maximum value for each histogram is one. It is clear that the enhanced algorithm performs significantly better: with more point correspondences and fewer outliers to deal with, the conditioning of the calculation is improved and the solutions are more accurate.

An estimate of the Essential Matrix between pairs of frames is the first step towards structure from motion (SFM) measurements. Once the transformation between the camera frames has been recovered, the camera frame position vectors of all the tracked points can be found by triangulation. By applying suitable non-linear geometric constraints, the motion and structure can be refined further. From here, SFM algorithms differ according to whether the focus is on measuring the scene structure or the instantaneous camera motion. In the first case, batch processing of many images is used to optimize over all the inter-frame transformations and triangulated points simultaneously in a step called *bundle adjustment*, and in the second case an extended Kalman filter is used to sequentially update the camera position and locations of tracked points as more images arrive. Both types have applications in celestial vision-based navigation; in the first case, autonomous mapping and surface topography using imaging data, and in the second case, real-time navigation.

## 3. ABSOLUTE NAVIGATION USING THE FEIC

We have been investigating the possible integration of the FEIC algorithm with vision-based absolute navigation systems, in which the position and orientation of the spacecraft relative to the target body coordinate frame is determined by triangulation from a set of pre-defined 'known landmarks'. The main image processing task that is involved is the ability to re-recognize known landmarks on the target that have previously been identified and localized. This corrects the gradual drift that occurs in a vision-based guidance system when only short-lived transient features are used to measure the position via dead reckoning. The ability to identify known landmarks in an image whenever they are in the field of view is very challenging, and no well-established system exists. This is of course due to the fact that the same feature may have a very different appearance when seen from two different views, due to scale, orientation, illumination, affine and perspective changes. In the past, craters have provided suitable landmarks (e.g. [11],[12]), but the recent experience of the Hayabusa mission demonstrates that not all rocky Solar system bodies have craters. This motivates an approach that does not rely on any particular type of morphological feature for use as known landmarks, and instead uses generalized interest points that place no restriction on the type of surface feature they are attached to, other than the general roughness required to make a point distinct from its immediate environment. Perhaps the most common approach for this is to use SURF features [13][4], which aim to extract a scale-, orientation- and affine-invariant feature descriptor for interest points, which would then be stored as a known landmark and searched for in later imagery. All SURFs that are found in a new image would be checked against the entire known landmark database to see if they correspond to a known landmark, based on the similarity of their feature descriptors. However, it is not clear that the surface terrain common to Solar system bodies provides sufficiently distinctive SURF points for reliable matching in this manner.

An alternative and conceptually different interest-point-based approach is to use simpler types of image feature, and match them according to their geometrical arrangement rather than their individual appearance. Reference [14] extracted scale- and orientation-invariant descriptions of known landmarks based on the distribution of Harris corners in their vicinity. By rectifying later images to match the homography of the images used to construct the database, they were also able to make the descriptors invariant to affine distortions. They also found that certain surface features produce Harris corners consistently as the scale, orientation and projection varies, while the appearance of the feature changes drastically. Note that this approach does not rely solely on the imagery, as an estimate of the spacecraft orientation is required in order to rectify the images.

The approach we have been developing is similar to [14] in that simple image features (Harris corners) are used to provide known landmarks. This allows much of the FEIC algorithm to be re-used for absolute navigation. However, the image rectification that is crucial to [14] relies on the terrain in view being flat. As we have been focusing on NEO applications, where the surface in view is often extremely varied in depth and curvature, we have aimed to remove this requirement. In our system, the database of known landmarks consists simply of the target body frame positions of prominent Harris corners, and the associated uncertainty. Because the known landmarks are selected as prominent Harris corners, there is a chance that during later imaging they will be extracted and tracked as transient points, and we therefore expect some fraction of transient points to correspond to known landmarks. In order to make identifications between transient tracking points and known landmarks, we use an estimate of the target body frame position and orientation of the spacecraft to project all the known landmarks onto the image. The matching is then done simply by looking for transient points that lie close to projected known landmarks in the image plane. This system is naturally sensitive to errors in the target body frame transformation used to project the known landmarks onto the image, and we are currently investigating the noise degradation of the algorithm.

Fig. 3 depicts a model of the Lunar south pole roughly 1000km along each side, viewed using the PANGU tool [8]. The cyan diamonds mark 203 known landmarks; these have been selected as sites providing prominent Harris corners when viewed from an altitude of 300km and at a wide range of view directions, such as might be the case for orbital imagery. The close up images on the right demonstrate that these known landmarks are not restricted to any particular type of feature; in the top image the corners of shadows lying inside craters provide known landmarks, and in the lower image we find known landmarks on undulating terrain of indistinct morphology.

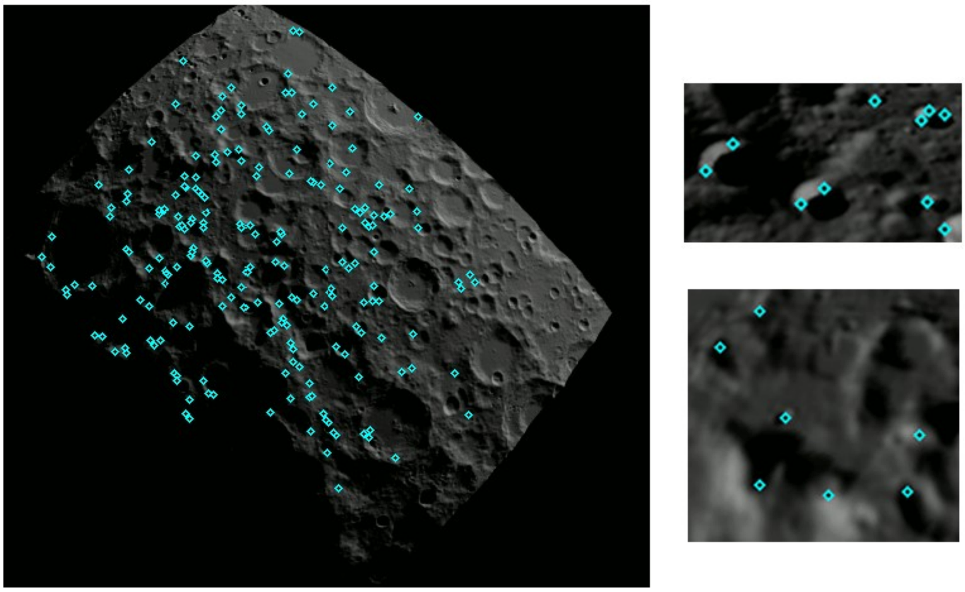

***Fig. 3:*** *Lunar south pole model viewed in PANGU, with known landmarks overlaid*

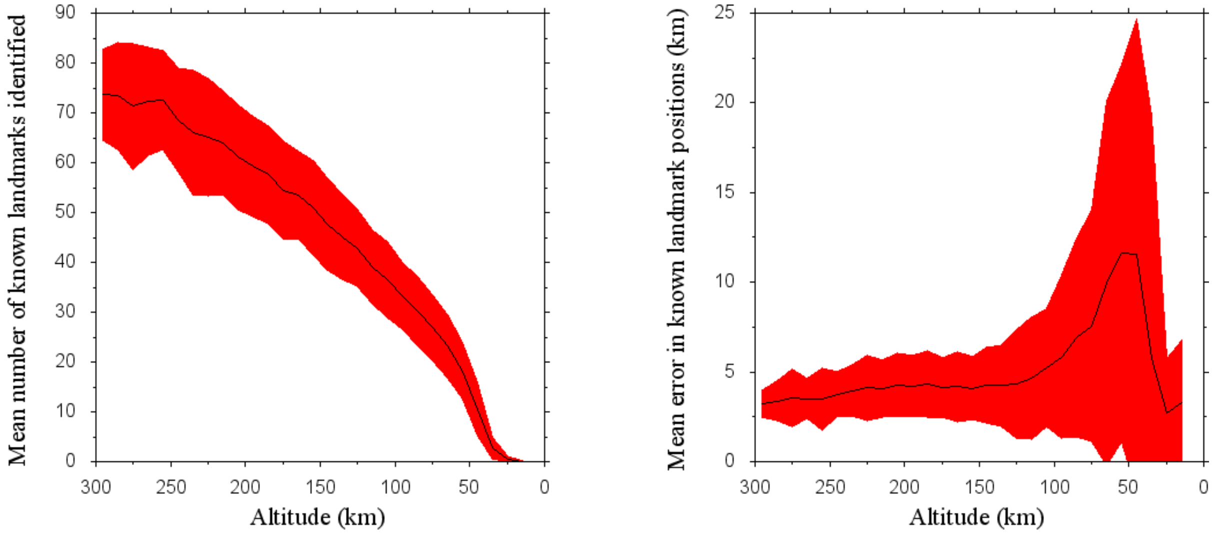


***Fig. 4:*** *Known landmarks identified during descent and landing simulations*

Using the Lunar south pole model and this database, we flew a set of fifty simulated descent trajectories from 300km to 0km, and investigated the performance of the known landmark detection algorithm during flight. We ran the FEIC algorithm on the simulated images, and matched transient tracking points with known landmarks based on proximity in the image. The spacecraft position and orientation used by the algorithm were degraded with Gaussian noise of $\sigma$=0.8 degrees in attitude and $\sigma$=200m in position to simulate realistic errors - these are the typical magnitude of Mars Express errors with respect to the Mars body frame [15] and are assumed to be representative of a Lunar scenario. One hundred and fifty frames were processed during each flight, and the aggregated results are presented in Fig. 4. In the graph on the left, we show the mean number of known landmarks identified during descent, along with the one-sigma sample variance. In the graph on the right, we show the mean 3D offset between the known landmark position and the position of the transient tracking point that it is matched to (measured using PANGU). This quantity is a proxy for the accuracy of the absolute navigation achievable using the identified known landmarks. It is clear that a significant number of known landmarks that were selected at 300km are reliably identified down to around 100km, after which noise starts to dominate the system as incorrect matches are made. These results validate the main assumption of this algorithm; that known landmarks selected as Harris corners from high altitude are reliably identifiable at much lower altitudes, given a reasonable estimate of the spacecraft pose.

We are working on improvements to this algorithm that build on these results and aim to improve the performance close to landing and the matching process overall. For example, the results of Section 3.B could be used to triangulate the target body frame positions of all points tracked during descent using the known landmarks to fix the scale, thus upgrading all transient points to known landmarks as they are based on the same type of feature point. We are also working on a lost-in-space version of the algorithm that would not rely on an initial estimate of the spacecraft pose, but would instead use randomized matching trials and possibly projective-invariant signatures for sets of known landmarks to determine this in a statistical manner.

## 4. CONCLUSIONS

In this paper we have demonstrated that relatively simple enhancements to the basic FEIC feature tracking algorithm can significantly improve the performance. This has been shown by analyzing the results of a typical vision-based navigation calculation that uses the improved point correspondences, namely the solution for the Essential Matrix and the recovery of the rotation and translation direction of the craft. With these enhancements, the problem of extracting sufficiently robust point correspondences is largely solved.

We have also demonstrated a conceptually new approach to known landmark identification, for the purposes of supporting vision-based absolute navigation systems. This algorithm is still in development, and we are working on enhancements to the basic system presented here that aim to improve the transient point/known landmark matching process and the performance close to landing. We are also developing a lost-in-space version of the algorithm that would not rely on an initial estimate of the spacecraft pose in the target body frame, but would instead determine this independently.

## 5. REFERENCES

[1] M. Pollefeys, *et al.* “Visual modeling with a hand-held camera”, *International Journal of Computer Vision* 59(3), 207–232, 2004

[2] A. Davidson, “Real-Time Simultaneous Localisation and Mapping with a Single Camera”, *IEEE International Conference on Computer Vision*, 1403–1410, Oct 2003

[3] A. Johnson, Y. Cheng, L.H. Matthies, "Machine vision for autonomous small body navigation," *IEEE Aerospace Conference Proceedings*, vol.7, 661–671 vol.7, 2000

[4] C. Cocaud, T. Kubota, “Probabilistic Motion Estimation for Near Real-Time Navigation and Landing on Small Celestial Bodies”, *In Proceedings of the 22nd International Symposium on Space Flight Dynamics*, 2011

[5] M. Dunstan, S. Parkes, S. Mancuso, “Visual Navigation Chip for Planetary Landers”, *In Proceedings of the DAta Systems In Aerospace conference (DASIA),* 2005

[6] B. Frapard, *et al.* “Navigation for Planetary Approach and Landing”, *In Proceedings of the 5th International ESA Conference on GNC*, 2002

[7] C. Harris, M.J. Stephens. “A combined corner and edge detector”, *Alvey Vision Conferences*, pages 147–152, 1988

[8] S. Parkes, M. Dunstan, I. Martin, M. McCrum, O. Dubois-Matra, “Testing advanced navigation systems for planetary landers and rovers”, *60th International Astronautical Congress* (IAC 2009), 2009

[9] R. Hartley, A. Zisserman, *Multiple View Geometry in computer vision*, Cambridge University Press, 2003

[10] R. Hartley, "In Defense of the Eight-Point Algorithm", *IEEE Transaction on Pattern Recognition and Machine Intelligence* 19 (6): 580–593, 1997

[11] Y. Cheng, A. Johnson, L. Matthies, C. Olson, “Optical landmark detection for spacecraft navigation”, *In Proceedings of the 13th Annual AAS/AIAA Space Flight Mechanics Meeting*, 2003

[12] B. G. Williams, *et al.* “Navigation for NEAR Shoemaker: the First Spacecraft to Orbit an Asteroid”, *In Proceedings of the AAS/AIAA Astrodynamics Specialists Conference*, AAS 01-371, 2001

[13] H. Bay, A. Ess, T. Tuytelaars, L. Van Gool, “SURF: Speeded Up Robust Features”, *Computer Vision and Image Understanding (CVIU)*, Vol. 110, No. 3, pp. 346–359, 2008

[14] B. van Pham, S. Lacroix, M. Devy, M. Drieux, C. Phillipe, “Visual Landmark Constellation matching for spacecraft pinpoint landing”, *In Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2009

[15] R. Maddè, *et al*. “Delta-DOR: A New Technique for ESA’s Deep Space Navigation”, *ESA Bulletin 128*, November 2006