

Neural network tracking and extension of positive tracking periods

Jay C. Hanan,^{1*} Tien-Hsin Chao,¹ Pierre Moreels²

¹Jet Propulsion Laboratory, CA, USA

²California Institute of Technology, CA, USA

ABSTRACT

Feature detectors have been considered for the role of supplying additional information to a neural network tracker. The feature detector focuses on areas of the image with significant information. Basically, if a picture says a thousand words, the feature detectors are looking for the key phrases (keypoints). These keypoints are rotationally invariant and may be matched across frames. Application of these advanced feature detectors to the neural network tracking system at JPL has promising potential. As part of an ongoing program, an advanced feature detector was tested for augmentation of a neural network based tracker. The advanced feature detector extended tracking periods in test sequences including aircraft tracking, rover tracking, and simulated Martian landing. Future directions of research are also discussed.

Keywords: autonomous tracking, feature extraction, feature detector, neural network, target recognition, data reduction.

1. INTRODUCTION

Automatic target tracking technology has a wide variety of applications. Tracking technology has been researched for use with navigation of Martian landers, rovers, or on Earth for avionics testing or various military applications. Domestic applications abound and a successful technology which is capable of robust application to a variety of scenarios with little expert technical attention is still under investigation. Specific applications often have solutions where a peculiar feature may be available, but that feature is not expected applicable to the general tracking problem. In order to obtain a general purpose tracking tool, a neural network method was employed. The goal of such a tool was to simplify the necessary inputs given by the operator; thus, a general robust, yet simple method was sought.

A previous examination of a neural network based tracking system showed the potential for improved tracking periods (the number of frames where tracking remains precisely on target).¹ However, information input to the neural network was limited, so that potentially useful image information was ignored—primarily in the interest of computational efficiency. Here we reexamine this technology with a consideration for conserving potentially useful information available to the system. Advanced feature detectors were applied to the entire frame. While tracking of only a small component in the frame may be sought, this additional information provides data such as camera pose, frame movement, object rotation, and scale. The advanced feature detector focuses on areas of the image with significant information. Basically, if a picture says a thousand words, the feature detectors are looking for the key phrases (keypoints). These keypoints are rotationally invariant and may be matched across frames. Application of these advanced feature detectors to the neural network tracking system at JPL has promising potential. Some tests of this technology with example applications are described below.

2. METHODOLOGY

An analysis of the tracking results reported previously suggests the stable tracking period for the examined sequence can be extended beyond 100 frames, if sufficient information can be input to the system to correct the instabilities.¹ A proof may be sought from the results of human tracking operators, traditionally used to mark object trajectories. The human operator uses additional information beyond the tagged area of interest to maintain awareness of the target. To input this additional information into the automatic tracking software, rotationally invariant features across the image are useful.

* Corresponding author; electronic mail: Jay.C.Hanan@jpl.nasa.gov

Furthermore, these features must be matched between sequential frames so that neurons may be trained based on similar objects. However, such an operation ordinarily would cost excessive computing resources. Thus, a compromise considering resources and accuracy must be determined.

Feature detectors have been considered for the role of supplying additional information to the neural network tracker. Feature detectors have been employed for some time. Lowe presents a recent feature detector that has been used with success.² The key ideas behind features detectors are simple. Of course, the best description of the image is the image itself. However, it takes considerable of time and space, so we want to store only small collections of data about an image. Homogenous zones contain little information, and more important, any zone from a homogenous region looks exactly the same as any other, giving poor localization (if you cut a small patch from a homogenous region, you have no way of telling where this patch has been cut). Many detectors look at places where there is a high gradient. Gradients are typically computed with filters such as:

$$\begin{pmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{pmatrix}$$

This ‘‘Sobel’’ filter does some smoothing to remove noise, and at the same time calculates the gradient at the center point. The gradient in one direction is not enough: if you identify the edge of a table, you have good localization in the direction perpendicular to the edge, but poor localization in the direction of the edge. Thus, we must make use of derivatives in both x and y . One detector that was developed 10 years ago and still used frequently is the Harris detector. The Harris detector takes the matrix

$$M = \begin{bmatrix} \left(\frac{d}{dx}I\right)^2 & \left(\frac{d}{dx}I\right)\cdot\left(\frac{d}{dy}I\right) \\ \left(\frac{d}{dx}I\right)\cdot\left(\frac{d}{dy}I\right) & \left(\frac{d}{dy}I\right)^2 \end{bmatrix}$$

and then takes the local extrema of

$$R = \det(M) - k(\text{tr}(M))^2$$

There are other feature detector methods: the Foerstner detector and the Hessian detector, which are based on the second derivatives. The detector used by Lowe is based on filtering with a difference of gaussians,² similar to a laplacian filter—again second derivatives. Recently, the detections have been applied in multiscale: first one detection is performed, then resample at a lower resolution and perform detection again. This is in order to match features even if there was a change of scale between views.

It is difficult to objectively determine which feature detector is ‘the best’. It is in large part application-dependent.³ Beyond an extension of previous work, in investigating these advanced feature detectors, we applied the Hessian detector to a series of simulated landing images for a Mars entry, decent, and landing sequence. In addition, we tested the neural network tracking system as part of a path planning test bed for extended tracking of rovers and obstacles over terrain. Results from these tests are given below.

3. PRELIMINARY RESULTS

3.1 Aircraft Tracking

Consistent with previously reported results,¹ the additional information provided by advanced feature detectors can extend the stable tracking period the length of the sequences (>300 frames compared to the 100+ reported earlier). However, optimal automatic selection of scale and feature density is still under investigation. The system was tested on the sequences analyzed previously along with a rocket sequence depicted in Figure 1. The previously tested sequences provide a test of rotation and small changes in scale. The rocket sequence provided a test of more extreme changes in scale along with some cloud obscurations. The addition of occlusions provided a good opportunity to test the advanced feature detector. Advantages in target reacquisition and frame movement tracking were observed. As with the previously tested sequences, tracking periods improved with the input of additional information. The rocket sequence was positively tracked with periodic loss of target associated with cloud cover.

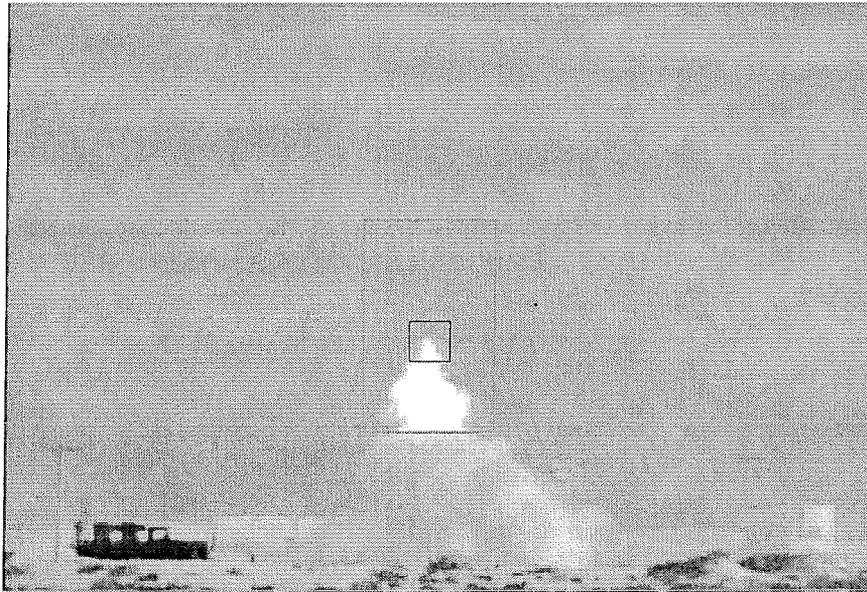


Figure 1 An early frame in a sequence where successful tracking of a rocket through large changes in scale and with the presence of obscurations such as clouds was demonstrated.

3.2 Rover Tracking

The advanced feature detectors were unnecessary for positive indefinite tracking of the rover in the test bed set up (Figure 2 A). Tracking of the rover across the test bed was successful on all trials.⁴ Variation in lighting and rover travel rate were tested against the system. Rover tracking was consistent and continuous up to the maximum rate of travel available by the system (~ 0.15 m/s).

Along with motion tracking of the rover science craft vehicle, science targets (represented by blocks in Figure 2 A) must also be identified. While this is a simple segmentation task for the blocks, a more advanced recognition system will be employed for field tests. Feature detectors were tested and found useful for applications requiring identification of natural science targets (Figure 2 B), such as the rocks expected in field application. The features on the rock in Figure 2 B were tracked through a 360 degree rotation. Incorporation of the feature detector to the rover system is part of the ongoing work of this project.

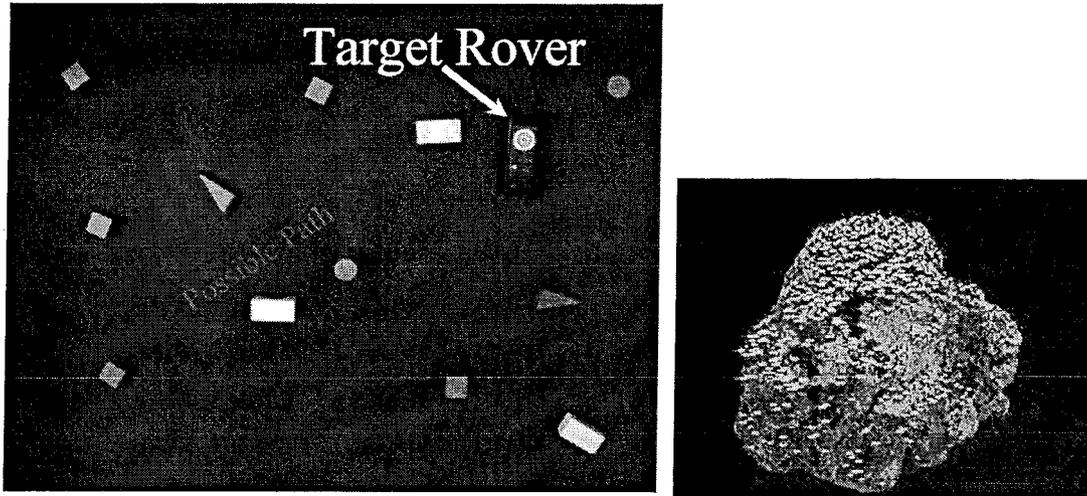


Figure 2 A) Test bed for rover path planning and obstacle avoidance tests. B) Feature detector used to identify "rock" target.

3.3 Mars Landing

The landing sequence provided a good test case for use of the advanced feature detector. The change of scale from high-altitude to surface provided a challenging case for the neural network tracker. Tracking stability decreases since the landing site is of low contrast and the scale and rotation changes are significant throughout the sequence. Here, since the object of interest is a safe landing site of low contrast, tracking of "off-target" objects throughout the image was necessary to maintain consistent tracking. Consistent recognition of craters between images was achieved using the advanced feature detector (see Figure 3). The simulation indicates the tool could perform tracking throughout the landing sequence. Testing of the optimal interconnection between the advanced feature detector and the neural net target tracking system remains in progress.

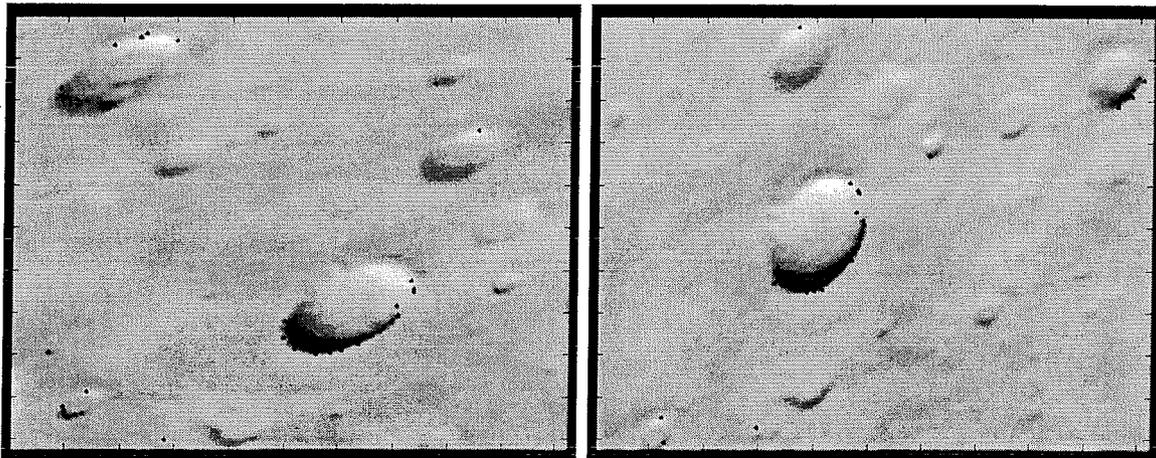


Figure 3 Extraction of two frame showing rotation and translation of the simulated Martian surface.⁵ The craters provide sufficient contrast to be recognized across frames using the advanced feature detector alone. Small blue spots mark areas where the advanced feature detector found a match.

4. SUMMARY

An advanced feature detector has been tested for augmentation of the neural network tracker. It is clear a second level of information extends tracking periods as evidenced by the tracking of the entire aircraft flight sequence as compared to the portion of the sequence reported previously.¹ However, the best method to incorporate this second level of information is still under investigation. Additional results will be presented in Optical Pattern Recognition XV.

Potential improvements may be realized by extending the number of frames over which features are compared, particularly where clouds may temporarily obscure the target. An automated connection between the confidence level given by the neural network and the level to which dependence on the advanced feature detector is a necessary first step. Field application to a longer sequence of up to 1 million frames will provide opportunities for testing further extensions of the tracking periods.

A case of indefinite positive tracking was confirmed on a simple target consisting of a rover for a path planning application at JPL. In this case the scale remained constant, but rotation was often extreme. Future application of the advanced feature detector to recognition of science targets and tracking of "natural" rock targets over multiple frames is also in progress. A final application of the advanced feature detector to simulated Mars landing images provided an efficient use of the method. The addition of the advanced feature detector significantly increased the tracking period for this sequence over that which was possible with solely the neural network tracking method. Additional tests of the method to real Mars landing data will also be presented.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Portions of the research leading to this publication were supported by the White Sands Missile Range (WSMR) and the Director's Research and Development Fund (DRDF).

REFERENCES

-
- ¹ J. C. Hanan, H. Zhou, T. Chao, "Precision of a radial basis function neural network tracking method," *SPIE, Optical Pattern Recognition XIV*, 5106 (2003).
 - ² D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," accepted for publication: *International Journal of Computer Vision*, (2004).
 - ³ C. Schmid, R. Mohr, and C. Bauckhage, "Comparing and Evaluating Interest Points," *IEEE Proc. Int'l Conf. Computer Vision*, pp. 230-235, (1998).
 - ⁴ W. Fink, J. C. Hanan, T.-H. Chao, M. A. Tarbell, "Field-Deployable Integrated Air-Ground Multi-Agent Autonomous Remote Planetary Surface Exploration." *NASA NTR-40428* (2003).
 - ⁵ N. J. Villaume, B. J. Martin "Mars entry, decent, and landing image sequence simulation" *unpublished data* (2003).